IBM Tivoli Storage Productivity Center
Version 5.2

User's Guide

SC27-4060-01

IBM Tivoli Storage Productivity Center
Version 5.2

User's Guide

SC27-4060-01

Note:
Before using this information and the product it supports, read the information in "Notices" on page 1775.

This edition applies to version 5, release 2, modification 0 of IBM Tivoli Storage Productivity Center (product numbers 5725-F92, 5725-F93, and 5725-G33) and to all subsequent releases and modifications until otherwise indicated in new editions. This edition replaces SC27-4060-00.

© Copyright IBM Corporation 2001, 2013.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures ix

Tables xi

Preface xv

Who should read this guide xv
Publications xv
Accessing publications online xv
IBM Redbooks xvii
Translation xvii
Downloading publications xvii
Providing feedback about publications xvii
Tivoli Storage Productivity Center Service Management Connect community xviii
Conventions used in this guide xviii
Contacting IBM Software Support xix
Reporting a problem xix

New for Tivoli Storage Productivity Center Version 5.2 xxi

Chapter 1. Getting started 1

Managing resources 6
Configuring and administering 12
Navigation 13

Chapter 2. Product overview 19

Architecture 19
Product licenses 21
Functions available in the web-based GUI based on role and license 23
Functions available in the stand-alone GUI based on role and license 28
Interfaces for Tivoli Storage Productivity Center . . 33
Available functions in the interfaces 34
Functions that were moved to the web-based GUI in V5.2 38
Starting Tivoli Storage Productivity Center 41
Starting the Tivoli Storage Productivity Center stand-alone GUI 41
Starting the Tivoli Storage Productivity Center web-based GUI 43
Starting Tivoli Storage Productivity Center for Replication 44
Tivoli Storage Productivity Center for Replication . 45
Key components of Tivoli Storage Productivity Center for Replication for System z 47
Interfaces for Tivoli Storage Productivity Center for Replication 48
Tivoli Storage Productivity Center for Replication integration with Tivoli System Automation Application Manager 48
Management server. 49
Storage systems 49
Host systems 49
Ports 50
Sessions 52
Session types 53
Metro Mirror heartbeat 60
Site awareness 60
Users and groups 61
User roles 61
Key concepts 61
Data collection 61
Performance monitoring and troubleshooting . . 62
Storage optimization 64
Cloud configuration 65
Provisioning storage 71
Tivoli Storage Productivity Center VASA provider 73
vSphere Web Client extension for Tivoli Storage Productivity Center. 74
Alerts 75
Reporting 77
Role-based authorization 77
Fabrics and zones 79
Agents 82
Single sign-on 89
Launch and launch in context 91

Chapter 3. Scenarios. 93

Optimizing the performance of storage virtualizers 93
Redistributing volume workloads to balance pools 94
Re-tiering volumes to balance pools 96
Tiering volumes by I/O density and file usage 97
Creating a performance threshold and viewing violations. 100
Scheduling data collection 100
Defining a performance threshold 101
Viewing threshold violations 102
Archiving and deleting files 103
Checking the status, date, and time of the last probe 103
Starting a probe 103
Defining the criteria to identify the files to be archived and deleted 104
Scanning the server to find the files that meet the criteria 104
Creating a report that shows the files that meet the criteria 105
Archiving and deleting the files that meet the criteria 105
Limiting a set of storage to provision in the vSphere Web Client extension 106
Creating a profile with a filter 107
Replication scenarios 109
Performing planned maintenance scenarios . . 110
Recovering from a disaster scenarios 113

© Copyright IBM Corp. 2001, 2013

Practicing disaster recovery scenarios 116

Chapter 4. Administering 121

Starting and stopping the Tivoli Storage Productivity Center servers 121
Stopping the Tivoli Storage Productivity Center servers 121
Starting the Tivoli Storage Productivity Center servers 122
Default locations of log files 124
Overview of required user names for initial logon to the Tivoli Storage Productivity Center family of products 125
Role-based authorization 126
Functions available in the web-based GUI based on role and license 128
Functions available in the stand-alone GUI based on role and license 133
Administering Tivoli Storage Productivity Center 138
Checking the version and license of Tivoli Storage Productivity Center 138
Viewing and managing services 139
Changing the user authentication configuration 146
Changing the LTPA token expiration for single sign-on 173
Checking Tivoli Storage Productivity Center status 174
Increasing the memory allocation for the Data server 174
Setting timeout values for the Device server . . 176
Changing passwords 177
Checking for a fully qualified host name . . . 188
Granting local administrative privileges to a domain account 190
SQL access to Tivoli Storage Productivity Center views 191
Repository copy tool 194
Collecting diagnostic information about Tivoli Storage Productivity Center 197
Administering resources and data sources 203
Checking the status data sources 205
CIM agents 205
Servers and Storage Resource agents 210
Hypervisors and VMware data sources 229
Storage systems 232
Switches and fabrics 241
Out-of-band fabric agents 244
Tivoli Storage Productivity Center servers . . . 246
Administering DB2 251
Backing up the Tivoli Storage Productivity Center database 251
Using the command line on UNIX and Linux 252
Manually starting DB2 on Windows 253
Manually stopping DB2 on Windows 253
Starting the IBM Data Studio full client. . . . 253
Monitoring DB2 254

Chapter 5. Managing resources . . . 257

Adding resources 257
Resources that you can monitor in the web-based GUI. 258
Adding servers 262
Adding hypervisors 271
Adding storage systems 272
Adding fabrics and switches 293
Collecting data 295
Collecting asset and status data 296
Collecting performance data 299
Collecting information about shares on storage systems 306
Collecting storage usage and trending information 306
Determining the availability of storage resources 327
Collect storage statistics from Tivoli Storage Productivity Center servers. 331
Data collection groups 334
Working with groups of resources 337
Alerting 341
Alert severities 343
Creating alerts 344
Configuring alert destinations 355
Viewing all alerts 356
Viewing alerts for a type of resource 358
Viewing alerts for a specific resource 358
Acknowledging alerts 359
Disabling or enabling alerts 360
Modifying alerts 361
Renaming alerts 362
Removing alerts 362
Viewing information about resources 363
How information is organized. 363
Viewing overall storage information on the dashboard 365
Viewing information about top-level resources 371
Viewing information about internal and related resources 371
Display differences in the interfaces 373
Monitoring the status of resources 374
How status is determined 375
Viewing the overall status of resources 378
Viewing the status of top-level resources . . . 379
Viewing the status of internal resources . . . 379
Acknowledging the status of resources 380
Monitoring the performance of resources 383
Viewing performance information 384
Defining performance thresholds 397
Viewing threshold violations and alerts. . . . 400
Viewing the data paths of resources 402
Resources that you can view in a data path . . 403
Navigating in a data path 406
Viewing information about top-level resources 406
Viewing information about internal resources 406
Zooming in 406
Zooming out 407
Changing the layout 407
Saving the display settings 407
Printing a view. 407
Exporting a data path to a file. 408
Configuring for storage cloud provisioning . . . 408
Service class configuration 409

Capacity pool configuration 414
Provisioning storage 417
Setting the zoning policy 418
Opening the Provision Storage wizard 418
How available space for a block-storage service class is calculated and how the placement of volumes is determined 421
How available space for a file-storage service class is calculated and how placement of shares is determined 425
Optimizing storage tiering 428
Tiering storage overview 429
Setting the tier level of storage pools 433
Analyzing tiering by servers 433
Analyzing tiering by hypervisors 434
Analyzing tiering by storage virtualizers . . . 434
Analyzing tiering by storage pools 434
Analyzing tiering by volumes 435
Modifying the period for analyzing performance data 435
Balancing pools 436
Transforming and migrating volumes 437
Modifying the properties of resources 437
Opening the management GUI for a resource. . . 440
Opening the DS Storage Manager GUI for a DS8000 441
Exporting information about resources, jobs, and alerts 441
Customizing lists of resources, jobs, and alerts . . 443
Filtering lists 443
Sorting lists 444
Showing, hiding, and reordering columns in lists 445
Managing tasks. 445
Viewing tasks 446
Managing tasks for provisioning 449
Managing tasks for tiering storage, balancing pools, and transforming storage 452
Renaming tasks 456
Deleting tasks 457
Managing resources in a vSphere environment . . 458
Provisioning block storage by using the vSphere Web Client extension for Tivoli Storage Productivity Center 458
Provisioning file storage by using the vSphere Web Client extension for Tivoli Storage Productivity Center 461
Managing resources in the stand-alone GUI . . . 464
GUI overview 464
Discovering storage resources 489
Checking storage configurations 490
Enforce storage policies 505
Working with storage resource groups 552
Working with scripts 555
Using the topology viewer 561
Managing storage systems 598
Managing schedules, jobs, and runs 605
Launch and launch in context 612
Working with invoices 620

Chapter 6. Managing replication . . . 629

Administering 629
Starting and stopping Tivoli Storage Productivity Center for Replication 629
Verifying that the Tivoli Storage Productivity Center for Replication server is running . . . 630
Identifying the version of IBM Tivoli Storage Productivity Center for Replication 631
Backing up and restoring IBM Tivoli Storage Productivity Center for Replication configuration data 631
Verifying the Java heap size 634
Managing management servers 634
Management servers 634
SNMP alerts 636
Setting up a standby management server . . . 639
Applying maintenance to an active management server 640
Reconnecting the active and standby management servers 641
Performing a takeover on the standby management server 641
Configuring SNMP 642
Adding SNMP managers 642
Changing the standby management server port number 642
Changing the client port number 643
Managing storage systems 643
Storage systems 643
Storage connections 645
Protected volumes 649
Site awareness 649
Adding a storage connection 650
Removing a storage system. 653
Modifying the location of storage systems . . . 654
Modifying storage connection properties . . . 654
Refreshing the storage system configuration . . 655
Setting volume protection 655
Restoring data from a journal volume 656
Configuring the use of the Easy Tier heat map transfer function 656
Managing host systems 659
Adding a host system connection. 659
Modifying a host system connection. 660
Removing a host system connection 661
Removing a session from an AIX host system connection 662
Managing logical paths 662
Viewing logical paths. 663
Adding logical paths 663
Adding logical paths by using a CSV file . . . 663
Removing logical paths 664
Setting up data replication 665
Sessions 665
Session types 673
Session commands 706
Site awareness 726
Preserve Mirror option 726
Creating sessions and adding copy sets. . . . 728
Using the Metro Mirror heartbeat 736
Exporting copy set data 739

Importing copy set data 739
Exporting historical data for Global Mirror role pairs 740
Modifying the location of session sites 740
Removing sessions 741
Removing copy sets 741
Migrating an existing configuration to Tivoli Storage Productivity Center for Replication . . 742
Practicing disaster recovery. 745
Practice volumes 745
Practicing disaster recovery for a Metro Mirror Failover/Failback with Practice session. 745
Practicing disaster recovery for a Global Mirror either Direction with two-site Practice session . 746
Practicing disaster recovery for a Global Mirror Failover/Failback with Practice session. 747
Practicing disaster recovery for a Metro Global Mirror Failover/Failback with Practice session . 747
Monitoring health and status 748
Viewing the health summary 748
Viewing SNMP alerts. 748
Viewing sessions 748
Viewing session properties 756
Viewing copy sets in a session. 780
Viewing storage system details 781
Viewing storage connection details 781
Viewing volume details 781
Viewing logical paths. 781
Viewing host connection details 782
Viewing console messages 782
Security 782
Users and groups 782
User roles 783
Granting access privileges for a user. 785
Viewing access privileges for a user 786
Modifying access privileges for a user 786
Removing access privileges for a user 787

Chapter 7. Reporting 789

Getting started 789
Predefined reports. 789
Predefined reports listed by package 790
Custom reports. 797
Viewing predefined reports. 798
Creating custom reports 798
Viewing and creating reports in Tivoli Common Reporting 799
Searching the predefined reports for a property or performance metric 800
Scenarios 800
Analyzing the volumes on the most active hypervisors 801
Monitoring daily the performance of volumes 802
Investigating a degradation in the performance of a storage pool 802
Identifying the relationships between a storage system and other resources. 803
Configuring 804
Setting options for reports 804
Scheduling a report to run with predefined options 804
Creating a copy of a predefined report 805
Starting the IBM Cognos Configuration GUI . . 806
Sending reports by email 806
Specifying locations for saving reports 807
Setting up multiple connections to the TPCDB data source 808
Customizing the logo and title of reports . . . 809
Creating a search index of reports, report views, and folders 811
Configuring security for Tivoli Common Reporting with the default authentication configuration 812
Storage virtualizer reports that are available when you upgrade 820
Storage virtualizer resource reports 820
Storage virtualizer performance reports. . . . 834
Storage virtualizer performance details reports 843
Predefined reports about historical information 858
Storage virtualizer custom reports about performance. 874
Predefined reports about resource relationships . . 930
Running the Storage Resource Relationships Summary report 930
Storage Resource Relationships Summary report 931
Running the Storage Resource Relationships Summary (Configurable) report 932
Storage Resource Relationships Summary (Configurable) report. 933
Running the File System to Volume Relationships report 935
File System to Volume Relationships report . . 935
Predefined reports about switches and switch ports 937
Predefined reports about switches 937
Predefined reports about switch ports 942
Predefined reports about hypervisors 949
Running the Hypervisors Capacity report . . . 949
Hypervisors Capacity report 950
Running the Most Active Hypervisors report 951
Most Active Hypervisors report 951
Running the Summarized Performance of Volumes by Hypervisor report. 952
Summarized Performance of Volumes by Hypervisor report 953
Running the Performance of Volumes by Hypervisor report 954
Performance of Volumes by Hypervisor report 955
Running the Hypervisor Data Stores Capacity report 956
Hypervisor Data Stores Capacity report . . . 956
Running the Hypervisor Disks Capacity report 957
Hypervisor Disks Capacity report 958
Predefined reports about servers and file systems on servers 959
Predefined reports about servers 959
Predefined reports about file systems on servers 967
Predefined reports about storage systems and components 969
Predefined reports about storage systems . . . 969
Predefined reports about controllers, modules, and nodes 982
Predefined reports about disks 993

Predefined reports about host connections . . 1002
Predefined reports about I/O groups 1008
Predefined reports about managed disks . . . 1013
Predefined reports about storage pools . . . 1021
Predefined reports about ports 1035
Predefined reports about RAID arrays. . . . 1040
Predefined reports about volumes 1046
Predefined reports about storage tiering 1058
Running the Storage Resource Group - VDisk Workload Activity Details report 1059
Storage Resource Group - VDisk Workload Activity Details report 1060
Running the VDisk Details report 1064
VDisk Details report. 1064
Running the MDisk Groups - VDisk Workload Activity report 1067
MDisk Groups - VDisk Workload Activity report 1067
Running the Storage Virtualizer - VDisk Workload Activity report 1071
Storage Virtualizer - VDisk Workload Activity report 1071
Running the Storage Resource Group - VDisk Workload Activity report 1075
Storage Resource Group - VDisk Workload Activity report 1076
Running the MDisk Group - VDisk Workload Activity Details report 1080
MDisk Group - VDisk Workload Activity Details report 1080
Running the MDisk Group Details report . . 1084
MDisk Group Details report 1085
Running the MDisk Groups - Workload Activity report 1087
MDisk Groups - Workload Activity report 1087
Read I/O capability formula 1092
Changing threshold values in storage tier reports 1094
Threshold values in storage tier reports . . . 1094
Custom reports about performance. 1095
Creating custom performance reports 1096
Data and properties in performance reports 1097
Custom reports about capacity and relationships 1218
Creating custom capacity and relationship reports 1219
Adding filters to reports 1220
Data and properties in capacity and relationship reports 1224
Troubleshooting reports 1289
Showing the package version number and build ID. 1294
Reporting in the stand-alone GUI 1295
Choosing a reporting type. 1295
Choosing a report category 1297
What can I find out using reports? 1303
Report considerations and limitations 1307
Working with reports 1311
Working with user-defined online reports . . 1320
System reports 1510
Batch reports 1562
User ID (saved) reports. 1574
Rollup reports. 1575
Data source reports - overview 1633
Data source reports - CIMOM agents 1633
Tivoli Storage Productivity Center Monitoring Agent reports 1636
Sharing Tivoli Storage Productivity Center alerts. 1641
Viewing Tivoli Storage Productivity Center VASA provider reports 1642
Reporting in a vSphere environment 1643
Viewing vSphere storage mapping to Tivoli Storage Productivity Center 1644
Viewing storage system metrics in vSphere Web Client extension 1647
Viewing Tivoli Storage Productivity Center fabric connections in vSphere Web Client extension 1654
Viewing block storage reports in the vSphere Web Client 1654
Viewing file system reports in the vSphere Web Client 1656
Viewing Tivoli Storage Productivity Center VASA provider reports 1656

Appendix A. Reference 1659

Return codes used by Storage Resource agent 1659
Agent types for monitoring fabrics and switches 1662
Supported storage systems providing full disk encryption and solid-state drives 1663
Alerts 1663
Triggering conditions for alerts 1665
Alert suppression actions 1681
Triggered actions for alerts 1682
Alert display colors 1683
Alerts and monitoring groups 1685
Performance data. 1690
Performance thresholds 1690
Performance metrics. 1701
Arguments and window locations available with the inbound launch in context feature 1729
Launch in context parameters 1729
Launch-in-context task parameters 1731
Available windows 1731
Opening Tivoli Storage Productivity Center on Windows operating systems 1733
Opening Tivoli Storage Productivity Center CLIs and GUIs 1733
Accessing administration tools 1735
Frequently Asked Questions 1737
General and Migration information. 1737
Data Manager 1739
Fabric Manager 1739
Disk Manager 1740
Protocols and standards 1741
Web Based Enterprise Management 1741
Storage Management Initiative Specification 1742
Service Location Protocol 1743
Simple Network Management Protocol . . . 1743
Fibre Channel Methodologies of Interconnects 1744

Appendix B. Accessibility features for Tivoli Storage Productivity Center 1745

Appendix C. Accessibility features for Tivoli Storage Productivity Center for Replication 1747

Glossary 1751
A 1751
B 1751
C 1751
D 1752
E 1752
F 1752
G 1752
H 1752
I 1753
J 1753
K 1753
L 1753
M 1753
N 1754
O 1754
P 1754
R 1754
S 1755
T 1755
V 1755
Z 1755

Index 1757

Notices 1775
Privacy policy considerations. 1777
Trademarks 1777

Figures

1. Tivoli Storage Productivity Center components 21
2. Global security page, configure federated repositories 153
3. Federated repositories page, Manage repositories 154
4. Manage repositories page, Add a new repository. 155
5. Manage repositories > New page 156
6. Global security > Federated repositories page, Add base entry to realm 157
7. Global security > Federated repositories > Repository reference page, Add base entry to realm 159
8. Opening the Global security page. 168
9. Configuring the available realm definitions 169
10. Removing the LDAP repository 170
11. Managing the repositories 171
12. Deleting the LDAP repository 172
13. Master and subordinate server architecture 246
14. File Filters tab, pop-up menu 318
15. Create Condition window 319
16. Create Group Window 319
17. File Filter window 320
18. Navigation tree 465
19. Content pane 466
20. Snapshot selection panel 492
21. Icons and colors of the change overlay 494
22. The Create Analyzer window 497
23. Topology overview 563
24. Minimap 564
25. Computer group 565
26. Data Path Explorer view 586
27. Connections 594
28. Highlighted entity 595
29. Pinned entity 596
30. Sample Invoice — OS Storage Usage By User 625
31. Sample Invoice — OS Disk Capacity by Computer. 626
32. Sample Invoice — Storage Usage by Database User 627
33. Sample Invoice — Total Size by Database-Tablespace 628
34. Chart Generation 1314
35. Tabs 1317
36. Trending with Historical Data. 1318
37. History Chart for File System Used Space 1320
38. Selection Resources window 1325
39. Usage Reporting—Access Load—Access Time—Chart Network Wide 1421
40. Usage Reporting—Access Load—Modification Time—Chart Network Wide 1421
41. Usage Reporting—Files—File Size Distribution—File Count Chart Tab 1425
42. Managed SAN overview report 1637
43. Managed SAN Health report 1638
44. TPC Server overview report 1639
45. TPC Server overview report 1640

© Copyright IBM Corp. 2001, 2013

Tables

1. Locations of publications for Tivoli Storage Productivity Center and related products . . xvi
2. Tasks for getting started with Tivoli Storage Productivity Center 1
3. Tasks for monitoring and managing resources 7
4. Tasks for configuring and administering Tivoli Storage Productivity Center 12
5. Tivoli Storage Productivity Center licenses when you upgrade from Version 4.2 to Version 5.2 23
6. Role and license restrictions in the web-based GUI 24
7. Functions with each Tivoli Storage Productivity Center license 29
8. Available reports with each IBM Tivoli Storage Productivity Center license 31
9. Available configuring and administering tasks in each GUI 34
10. Available monitoring tasks in each GUI 36
11. Available alerting tasks in each GUI 37
12. Available reporting tasks in each GUI 38
13. Available reporting information in each GUI 38
14. Functions that were moved from the stand-alone GUI to the web-based GUI . . . 38
15. Where to find Tivoli Storage Productivity Center for Replication port numbers 50
16. Storage system default ports 51
17. Basic HyperSwap 54
18. FlashCopy 54
19. Snapshot 55
20. Metro Mirror Single Direction 55
21. Metro Mirror Failover/Failback 56
22. Metro Mirror Failover/Failback with Practice 56
23. Global Mirror Single Direction 57
24. Global Mirror Either Direction with Two-Site Practice 57
25. Global Mirror Failover/Failback 58
26. Global Mirror Failover/Failback with Change Volumes 58
27. Global Mirror Failover/Failback with Practice 59
28. Metro Global Mirror 59
29. Metro Global Mirror with Practice 59
30. How roles in previous versions of Tivoli Storage Productivity Center are mapped to the roles in version 5.2 and later 78
31. Utilization of the pools on the storage virtualizer `svc_payment_records`. 94
32. Utilization of the pools on the `svc_sales_records` storage virtualizer. 96
33. Minimum threshold values for tier 1, tier 2, and tier 3 98
34. Current I/O density and file usage data for the volumes in tier 1 pools. 98
35. Role-pair states at the completion of each step (participating and nonparticipating role pairs) 117
36. Role-pair states at the completion of each step (participating and nonparticipating role pairs) 118
37. Role-pair states at the completion of each step (participating and nonparticipating role pairs) 120
38. Required user name for initial logon after installation on multiple servers with a remote database schema 126
39. Required user name for initial logon after installation on multiple servers and Tivoli Storage Productivity Center reports are remote 126
40. How roles in previous versions of Tivoli Storage Productivity Center are mapped to the roles in version 5.2 and later 127
41. Role and license restrictions in the web-based GUI. 129
42. Functions with each Tivoli Storage Productivity Center license 133
43. Available reports with each IBM Tivoli Storage Productivity Center license 136
44. Tasks for administering resources in each GUI 204
45. Fabric resources that you can monitor in the web-based GUI 259
46. Hypervisor resources that you can monitor in the web-based GUI 259
47. Server resources that you can monitor in the web-based GUI 260
48. Storage system resources that you can monitor in the web-based GUI. 261
49. Switch resources that you can monitor in the web-based GUI 262
50. 262
51. Agent types for switch and fabric functions 294
52. Data collection jobs, resources, and the GUI where you define the jobs 295
53. Possible severities of alerts 344
54. Possible statuses of monitored resources 374
55. Internal resources that are used to determine the status of top-level resources 375
56. Propagation of the statuses for internal resources in the web-based GUI 377
57. Internal resources and locations where you access the performance view 388
58. Controls for the performance view 390
59. Resources that you can view for DS6000, DS8000, and ESS storage systems in the chart legend 395
60. Resources that you can view for SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified storage systems in the chart legend 395
61. Resources that you can view for XIV systems in the chart legend. 396
62. Resources that you can view for SMI-S BSP systems in the chart legend 396

© Copyright IBM Corp. 2001, 2013

63. Resources that you can view for switches in the chart legend. 396
64. Resources that you can view for servers in the chart legend. 396
65. Resource icons in a data path 403
66. Filtering candidates for provisioning based on service class properties 422
67. Filtering candidates for provisioning based on service class properties 427
68. Source pools selection. 430
69. Target pool selection 430
70. Minimum thresholds for the I/O density of the storage tiers. 431
71. I/O densities of the volumes in the source pool 431
72. Minimum thresholds for file access frequency for each storage tier 432
73. Last accessed statistics for the source pool on tier 1 432
74. Icons and colors of the change overlay 493
75. Descriptions of policy types and their triggering conditions 506
76. Group locations in the topology viewer 566
77. Topology status indicators and their icons 574
78. Internal resources that are used to determine the status of top-level resources 578
79. Propagation of the statuses for internal resources in the stand-alone GUI 579
80. Categories of operational status for storage entities. 580
81. Session state change traps 637
82. Configuration change traps 638
83. Suspending-event notification traps 638
84. Communication-failure traps 638
85. Management Servers state-change traps 639
86. Number of role pairs and volumes per copy set for each session type 666
87. Basic HyperSwap 674
88. FlashCopy 674
89. Snapshot 674
90. Metro Mirror Single Direction 675
91. Metro Mirror Failover/Failback 676
92. Metro Mirror Failover/Failback with Practice 676
93. Global Mirror Single Direction 677
94. Global Mirror Either Direction with Two-Site Practice 677
95. Global Mirror Failover/Failback 678
96. Global Mirror Failover/Failback with Change Volumes 678
97. Global Mirror Failover/Failback with Practice 679
98. Metro Global Mirror 679
99. Metro Global Mirror with Practice 679
100. Basic HyperSwap commands 707
101. FlashCopy commands. 707
102. Snapshot session commands 708
103. Snapshot group commands 708
104. Metro Mirror commands. 709
105. Metro Mirror with Practice commands 711
106. Global Mirror commands 712
107. Global Mirror with Practice commands 715
108. Metro Global Mirror commands 718
109. Metro Global Mirror with Practice commands 722
110. Session status icons 749
111. Volume role symbols 749
112. Data copying symbols 750
113. Session states 752
114. Detailed status messages for Participating and Non-Participating role pairs. 754
115. Information about clusters associated with a Storwize V7000 Unified storage system or an IBM SONAS system 1342
116. Information about nodes associated with a Storwize V7000 Unified storage system or an IBM SONAS system 1343
117. Information about pools associated with a Storwize V7000 Unified storage system or an IBM SONAS system 1344
118. Information about a selected NSD associated with a Storwize V7000 Unified storage system or an IBM SONAS system 1344
119. Information about a selected NSD associated with a Storwize V7000 Unified storage system 1345
120. Information about the relationship between a selected NSD and an associated virtual disk or volume on a Storwize V7000 Unified storage system. 1346
121. Information about a selected file system associated with a Storwize V7000 Unified storage system or an IBM SONAS system. . 1347
122. Information about a selected file set associated with a Storwize V7000 Unified storage system or an IBM SONAS system. . 1348
123. Information about a selected export associated with a Storwize V7000 Unified storage system or an IBM SONAS system. . 1348
124. NFS protocols for a selected export associated with a Storwize V7000 Unified storage system or an IBM SONAS system. . 1349
125. CIFS protocols for a selected export associated with a Storwize V7000 Unified storage system or an IBM SONAS system. . 1350
126. Information about a directory associated with a Storwize V7000 Unified storage system or an IBM SONAS system 1350
127. Tivoli Storage Productivity Center alerts in vCenter 1641
128. vSphere Storage mapping to Tivoli Storage Productivity Center 1645
129. Storage resource agent return codes 1659
130. Supported agent types for Switch and Fabric Functions 1662
131. Triggering conditions and related error messages for computer alerts 1665
132. Triggering conditions and related error messages for file system alerts 1666
133. Triggering conditions and related error messages for directory alerts 1668
134. Triggering conditions and related error messages for instance alerts 1669
135. Triggering conditions and related error messages for database-table space alerts . . 1669

136. Triggering conditions and related error messages for table alerts 1670
137. Triggering conditions and related error messages for storage-system alerts 1671
138. Triggering conditions and related error messages for fabric alerts 1675
139. Triggering conditions and related error messages for switch alerts 1676
140. Triggering conditions and related error messages for endpoint device alerts. . . . 1677
141. Triggering conditions and related error messages for hypervisor alerts 1678
142. Triggering conditions and related error messages for NAS alerts 1678
143. Triggering conditions and related error messages for replication alerts 1680
144. Alert colors 1683
145. Predefined monitoring groups and alerts for storage subsystems 1686
146. Predefined monitoring groups and alerts for fabrics 1688
147. Predefined monitoring groups alerts for computers 1689
148. Array thresholds 1692
149. Controller thresholds 1696
150. Port thresholds 1697
151. Available windows for inbound launch-in-context 1731
152. Interop Namespaces for CIM Agents for switches and storage systems 1738
153. Agent types for switch and fabric functions 1739

Preface

Who should read this guide

This publication is intended for administrators and users who want to learn how to use Tivoli® Storage Productivity Center to monitor and manage storage resources and infrastructure.

Administrators and users must be familiar, or become familiar, with the following concepts:

- Storage Area Network (SAN)
- Simple Network Management Protocol (SNMP)
- Storage provisioning, optimization, performance, tiers, and cloud configuration
- Tivoli Storage Productivity Center for Replication
- IBM® DB2®
- Jazz™ for Service Management
- IBM Tivoli Common Reporting
- vSphere Web Client extension and VASA

Publications

You are provided with a list of publications in the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication libraries and other related publications. The following section also describes how to access publications online, how to order publications, and how to submit comments about publications.

The publications are available from the IBM Publications Center at http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss.

Accessing publications online

Publications for this product and other related products are available online.

Information about installing, configuring, upgrading, and uninstalling Tivoli Storage Productivity Center and related products is available online. Use Table 1 on page xvi to view and download these publications. Translated documents are available for some products.

IBM Tivoli Documentation Central also provides links to the information centers for all Tivoli products. For information about referenced Tivoli products, such as Tivoli Storage Manager and Tivoli Enterprise Console®, go to the home page for Tivoli Documentation Central.

© Copyright IBM Corp. 2001, 2013

Table 1. Locations of publications for Tivoli Storage Productivity Center and related products

Product	Online location
IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication	Tivoli Storage Productivity Center Information Center. In the navigation pane, click **IBM Tivoli Storage Productivity Center**. To obtain PDF documents, click **IBM Tivoli Storage Productivity Center** > **Printable documentation**. To view previous versions of the IBM Tivoli Storage Productivity Center Information Center, go to the Tivoli Storage Productivity Center page in Tivoli Documentation Central.
IBM Tivoli Storage Productivity Center for Replication for System z®	Tivoli Storage Productivity Center Information Center. In the navigation pane, click **IBM Tivoli Storage Productivity Center for Replication for System z**. To obtain PDF documents, click **IBM Tivoli Storage Productivity Center for Replication for System z** > **Printable documentation**. To view previous versions of the IBM Tivoli Storage Productivity Center for Replication Information Center, go to the Tivoli Storage Productivity Center page in Tivoli Documentation Central.
IBM WebSphere® Application Server	http://publib.boulder.ibm.com/infocenter/wasinfo/v8r0/index.jsp
Jazz for Service Management	http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.psc.doc_1.1.0/psc_ic-homepage.html
IBM System Storage® DS3000, IBM System Storage DS4000®, or IBM System Storage DS5000	http://www.ibm.com/support/entry/portal/
IBM System Storage DS6000™	http://publib.boulder.ibm.com/infocenter/dsichelp/ds6000ic/index.jsp
IBM System Storage DS8000®	http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp
IBM System Storage DS® Open Application Programming Interface publications	http://www.ibm.com/support/entry/portal/Troubleshooting/Hardware/System_Storage/Storage_software/Other_software_products/CIM_Agent_for_DS_Open_(API)/ Use these publications for information about how to install, configure, and use the CIM agent.
IBM System Storage SAN Volume Controller	http://pic.dhe.ibm.com/infocenter/svc/ic/index.jsp
IBM Storwize® V3500	http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp
IBM Storwize V3700	http://pic.dhe.ibm.com/infocenter/storwize/v3700_ic/index.jsp
IBM Storwize V7000	http://pic.dhe.ibm.com/infocenter/storwize/ic/index.jsp
IBM Storwize V7000 Unified	http://pic.dhe.ibm.com/infocenter/storwize/unified_ic/index.jsp

Table 1. Locations of publications for Tivoli Storage Productivity Center and related products (continued)

Product	Online location
IBM Scale Out Network Attached Storage (IBM SONAS)	http://pic.dhe.ibm.com/infocenter/sonasic/sonas1ic/index.jsp
IBM XIV® Storage System	http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp
IBM DB2 Database for Linux, AIX®, and Windows	http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/index.jsp
IBM System Storage N series	http://www.ibm.com/systems/storage/network/redbooks.html For more information about IBM System Storage N series, see http://www.ibm.com/systems/storage/network/software/.
IBM Systems Director	http://www.ibm.com/systems/software/director/index.html
VMware vSphere	http://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html

IBM Redbooks

The IBM Redbooks® are publications about specialized topics.

You can order publications through your IBM representative or the IBM branch office serving your locality. You can also search for and order books of interest to you by visiting the IBM Redbooks home page at http://www.redbooks.ibm.com.

Translation

Translated publications are available from the information center which is available in certain translated languages. It is displayed in the language that is appropriate for the browser locale setting.

When a locale does not have a translated version, the information center is displayed in English, which is the default language. Translations of the PDFs are available when the information center is translated.

Contact your service representative for more information about the translated publications and whether translations are available in your language.

Downloading publications

IBM publications are available in electronic format to be viewed or downloaded free of charge.

You can download IBM publications from http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss.

Providing feedback about publications

Your feedback is important to help IBM provide the highest quality information. You can provide comments or suggestions about the documentation from the IBM Tivoli Storage Productivity Center Information Center.

Go to the information center at http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp and click **Feedback** on the information center Welcome page or at the bottom of the individual topic pages.

Tivoli Storage Productivity Center Service Management Connect community

Connect, learn, and share with Service Management professionals: product support technical experts who provide their perspectives and expertise.

Access Service Management Connect at https://www.ibm.com/developerworks/servicemanagement/. Use Service Management Connect in the following ways:

- Become involved with transparent development, an ongoing, open engagement between other users and IBM developers of Tivoli products. You can access early designs, sprint demonstrations, product roadmaps, and prerelease code.
- Connect one-on-one with the experts to collaborate and network about Tivoli and the Storage Management community.
- Read blogs to benefit from the expertise and experience of others.
- Use wikis and forums to collaborate with the broader user community.

Conventions used in this guide

Information is given about the conventions that are used in this publication.

This publication uses several conventions for special terms and actions, and for operating system-dependent commands and paths.

The following typeface conventions are used in this publication:

Bold

- Flags that display with text
- Graphical user interface (GUI) elements (except for titles of windows and dialogs)
- Names of keys

Italic

- Variables
- Values that you must provide
- New terms
- Words and phrases that are emphasized
- Titles of documents

monospace

- Commands and command options
- Flags that display on a separate line
- Code examples and output
- Message text
- Names of files and directories
- Text strings that you must type, when they display within text
- Names of Oracle Java™ methods and classes
- HTML and XML tags that display `like this`, in monospace type

For syntax notations, remember the following details.

- In AIX, the prompt for the root user is #.
- In AIX and Linux, the commands are case-sensitive, so you must type commands exactly as they are shown.

Contacting IBM Software Support

You can contact IBM Software Support by phone, and you can register for support notifications at the technical support website.

- Go to the IBM Tivoli Storage Productivity Center technical support website at http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center.

  To receive future support notifications, sign in under **Notifications**. You are required to enter your IBM ID and password. After you are authenticated, you can configure your subscription for Tivoli Storage Productivity Center technical support website updates.
- Customers in the United States can call 1-800-IBM-SERV (1-800-426-7378).
- For international customers, go to the Tivoli Storage Productivity Center technical support website to find support by country. Expand **Contact support** and click **Directory of worldwide contacts**.

You can also review the *IBM Software Support Handbook*, which is available at http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html.

The support website offers extensive information, including a guide to support services; frequently asked questions (FAQs); and documentation for all IBM Software products, including Redbooks and white papers. Translated documents are also available for some products.

When you contact IBM Software Support, be prepared to provide identification information for your company so that support personnel can readily assist you. Company identification information might also be needed to access various online services available on the website. See “Reporting a problem.”

Reporting a problem

Provide the IBM Support Center with information about the problems that you report.

Have the following information ready when you report a problem:

- The IBM Tivoli Storage Productivity Center version, release, modification, and service level number.
- The communication protocol (for example, TCP/IP), version, and release number that you are using.
- The activity that you were doing when the problem occurred, listing the steps that you followed before the problem occurred.
- The exact text of any error messages.

New for Tivoli Storage Productivity Center Version 5.2

New features and enhancements are available in IBM Tivoli Storage Productivity Center Version 5.2.

Product component, command, and license updates

Read about updates to Tivoli Storage Productivity Center components, commands, and licensing.

Product component updates

The following updates were made to the Tivoli Storage Productivity Center product components:

- The web-based GUI now runs on the embedded IBM WebSphere Application Server, instead of Tivoli Integrated Portal. This new server is called the *web server*.
- The Device server and Replication servers now run on WebSphere Application Server Liberty Profile, instead of the embedded WebSphere Application Server.
- Tivoli Storage Productivity Center reports now require Jazz for Service Management Version 1.1.0.1 and IBM Tivoli Common Reporting Version 3.1.0.1. You must install Jazz for Service Management and Tivoli Common Reporting from the Tivoli Storage Productivity Center installation program.

  See *Installing Tivoli Storage Productivity Center in a single-server environment* or *Installing Tivoli Storage Productivity Center in a multiple-server environment* in the Tivoli Storage Productivity Center Information Center at http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp.

Command updates

The commands to start or stop Tivoli Storage Productivity Center services changed. See "Starting the Tivoli Storage Productivity Center servers" on page 122 and "Stopping the Tivoli Storage Productivity Center servers" on page 121.

License updates

The following updates were made to the Tivoli Storage Productivity Center licenses:

- The license that was previously called IBM Tivoli Storage Productivity Center Advanced is now the IBM SmartCloud® Virtual Storage Center (VSC) Storage Analytics Engine license. The user interfaces and messages still refer to this license as IBM Tivoli Storage Productivity Center Advanced.

  The new name of this license reflects its inclusion in the VSC solution, which bundles Tivoli Storage Productivity Center with IBM System Storage SAN Volume Controller and Tivoli Storage FlashCopy® Manager. By bundling these separate products, VSC provides a combined storage virtualization platform and storage management solution.

  The VSC Storage Analytics Engine contributes the advanced capabilities of Tivoli Storage Productivity Center, such as storage tier optimization and volume workload distribution. The VSC Storage Analytics Engine license is only available as part of the VSC solution.

© Copyright IBM Corp. 2001, 2013

For more information about IBM SmartCloud Virtual Storage Center, see http://www.ibm.com/software/products/us/en/vsc/.

Enhancements to the Tivoli Storage Productivity Center web-based GUI

Many of the functions for monitoring and managing storage were moved from the stand-alone GUI to the web-based GUI. With this migration, you can complete the following storage-management tasks in the web-based GUI:

- Adding resources for monitoring
- Administering monitored resources
- Monitoring the performance of resources, including server-centric monitoring of SAN resources without requiring an agent
- Monitoring violations of performance thresholds
- Configuring for storage cloud provisioning
- Implementing policy-based storage provisioning
- Identifying hot spots on storage virtualizers
- Balancing the workload of volumes across storage tiers
- Re-tiering volumes automatically or by using criteria such as file usage to place volumes on storage tiers
- Transforming storage by moving volumes between pools, converting volumes to thin-provisioned volumes, and converting volumes to compressed volumes

These and other tasks that you can complete from the web-based GUI are discussed in the following sections.

For more information about the functions that were moved from the stand-alone GUI to the web-based GUI, see "Functions that were moved to the web-based GUI in V5.2" on page 38.

For information about all the functions that are available in the web-based GUI, see "Available functions in the interfaces" on page 34.

Adding resources

Tivoli Storage Productivity Center provides wizards that guide you through the steps for discovering resources, adding resources for monitoring, and scheduling data collection. You can add the following resources for monitoring:

- Storage systems (other than IBM SONAS storage systems)
- Servers (without a Storage Resource agent)
- Hypervisors
- Fabrics
- Switches

See "Adding resources" on page 257.

Administering monitored resources

You can complete many administrative tasks for resources that are monitored by Tivoli Storage Productivity Center. These tasks include setting up data collection (probes and performance monitors), administering Storage Resource agents, testing connections, viewing logs, opening management GUIs, updating user credentials, and removing resources from Tivoli Storage Productivity Center.

See “Administering resources and data sources” on page 203.

Monitoring the performance of resources

You can collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations.

Metrics can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage. Threshold violations alert you when the performance of a monitored resource falls outside of a specified range.

See “Monitoring the performance of resources” on page 383.

Configuring for storage cloud provisioning

You can create service classes and capacity pools to classify and separate storage resources.

Service classes

You create service classes to describe capabilities and characteristics of storage resources for future provisioning requests.

Capacity pools

You create capacity pools to optionally divide storage resources in any way that is appropriate for your environment or your business needs. For example, storage resources for different departments within an organization can be separated into different capacity pools. Future provisioning and optimization requests can be restricted to storage within a capacity pool.

See “Service classes” on page 66 and “Capacity pools” on page 70.

Provisioning storage

Use the Provision Storage wizard to assign storage to servers and hypervisors. The wizard guides you through the steps for provisioning volumes or shares. Using the wizard, you request the amount of storage you need and a service class. The service class acts as a set of requirements for the placement of the volume or share. Tivoli Storage Productivity Center identifies the storage resources that can provide the requested capacity and that can also satisfy the requirements of the service class.

See “Provisioning storage” on page 71.

Identifying hot spots for storage pools

You can scan the utilization values for pools on a storage virtualizer, such as SAN Volume Controller, and identify the hot spots or the pools that have high utilization values on the storage virtualizer. You can also determine the course of action that you take. For example, if some of the pools on a tier have high utilization values and some have low utilization values, you can use the **Balance Pools** wizard to redistribute volumes across the storage tier. However, if all of the pools on the tier have high utilization values, you can use the **Analyze Tiering** wizard to balance the pools by re-tiering the most active volumes.

See “Redistributing volume workloads to balance pools” on page 94, “Re-tiering volumes to balance pools” on page 96, and “Tiering volumes by I/O density and file usage” on page 97.

Balancing the workload of volumes across storage pools

Use Balance Pools Analysis wizard to balance the workload of volumes across pools on the same tier. The pools are analyzed and

recommendations are generated to move volumes from pools with high utilization values to pools with low utilization values.

See “Balancing pools” on page 436 and “Managing tasks for tiering storage, balancing pools, and transforming storage” on page 452.

Re-tiering volumes

Use the Analyze Tiering wizard to tier volumes automatically or based on the criteria that you set in tiering policies. For example, you can tier volumes based on the volume workload or on file usage, or both. Depending on the conditions that are set in the tiering policy, recommendations are generated. For example, you can reduce storage costs by moving volumes with low workloads to lower or less expensive tiers. You can also improve performance and use storage more efficiently by moving volumes with heavy workloads to the tiers that best meet their workload requirements.

See “Optimizing storage tiering” on page 428 and “Managing tasks for tiering storage, balancing pools, and transforming storage” on page 452.

Transforming volumes

Use the Transform Storage wizard to complete the following tasks for one or more volumes in storage virtualizer pools:

- Move volumes from one storage virtualizer pool to another pool in the same storage virtualizer pool.
- Move volumes in a storage virtualizer pool to an Easy Tier® enabled pool.
- Convert fully allocated volumes to thin provisioned volumes and convert thin provisioned volumes to fully allocated volumes.
- Convert fully allocated volumes to compressed volumes and convert compressed volumes to fully allocated volumes.

The pools are analyzed and recommendations are generated. For example, analysis occurs to ensure that there is sufficient space in the pool to convert volumes or add volumes.

See “Transforming and migrating volumes” on page 437 and “Managing tasks for tiering storage, balancing pools, and transforming storage” on page 452.

Managing virtual storage in a VMware vSphere environment

The vSphere Web Client extension for Tivoli Storage Productivity Center and VASA provider features support integration with vSphere 5.1.

You can use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision storage and view reports on storage that is monitored by Tivoli Storage Productivity Center directly in the vSphere Web Client. See “Managing resources in a vSphere environment” on page 458 and “Reporting in a vSphere environment” on page 1643.

You can use Tivoli Storage Productivity Center as a VASA provider to publish storage alerts as events or alarms in vSphere, and to enhance vSphere reports with storage data. See “Tivoli Storage Productivity Center VASA provider” on page 73.

Managing data collection jobs

You can schedule data collection jobs such as probes and performance monitors to collect key data and metrics about monitored resources. Scan jobs for hypervisors and servers are still managed in the stand-alone GUI.

In Tivoli Storage Productivity Center Version 5.2, you can complete the following data collection jobs by using the web-based GUI:

- Schedule probes and performance monitors
- View performance monitor and probe logs
- Start a probe run
- Start or stop a performance monitor run

When you upgrade Tivoli Storage Productivity Center from a version that is earlier than V5.2, the existing probe and performance monitor jobs are automatically migrated.

For more information about collecting data, see "Collecting data" on page 295.

Accessing online help

The online help was enhanced to include more information about how to use Tivoli Storage Productivity Center. To access the help, click the help icon on any page in the GUI. For some controls, you can hover the mouse pointer over the hover icon for help.

Enhancements to reporting

The following features are new for reports that are in the web-based GUI.

Performance troubleshooting reports

You can view reports to analyze the performance of resources and to troubleshoot performance issues. For example, to compare up to four performance metrics for multiple switches over time, you can view the Compare Performance of Multiple Switches report. To compare one performance metric on one pool over two time periods, you can view the Compare Performance of One Pool over Time Ranges report.

You can view reports about the capacity of hypervisors and reports that show the most active servers and hypervisors. You can use the Performance Data Export report to export the performance metrics data for one or more resources on a storage system.

The storage virtualizer part of the Performance package was removed. The Performance package now treats all systems as storage systems, even if they are configured as storage virtualizers.

Scheduling and distributing reports

You can schedule reports and create the report output in HTML, PDF, and other formats. You can also configure reports to save the report output to your local file system, and to send reports as email attachments.

Reorganization of reports folders

The predefined reports are now organized by resource. For example, to access reports about storage system pools, go to **IBM Tivoli Storage Productivity Center Predefined Reports** > **Storage Systems** > **Pools**. Each report and folder has hover help and a description that is displayed in the **Details View** of Tivoli Common Reporting.

Changes to the appearance of predefined reports

The appearance of all predefined reports was changed to match more closely with the appearance of the web-based GUI. You can customize the

report templates to change the logo that is displayed in the output of all reports. You can also change the title that is displayed in the output of a report.

See Chapter 7, “Reporting,” on page 789.

New features for Tivoli Storage Productivity Center for Replication

Read about new features and enhancements that are available in Tivoli Storage Productivity Center for Replication.

Connect to z/OS® from open systems and manage HyperSwap® and hardened freeze You can use an IP connection to connect to IBM z/OS host systems from Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed.

You can use this feature to connect to a z/OS host system from a Tivoli Storage Productivity Center for Replication instance that is running on Windows, Linux, AIX, or z/OS. When you are connected to a host system, you can fully manage the HyperSwap and hardened freeze functions for volumes that are attached to the host system. In addition, you can manage these functions for volumes that are attached to other z/OS systems that are connected to the host system through a sysplex. This feature enables one instance of Tivoli Storage Productivity Center for Replication to manage HyperSwap and hardened freeze functions for volumes on multiple z/OS systems.

See “Managing host systems” on page 659, “Enabling HyperSwap or Open HyperSwap for a session” on page 700, and the session properties topics under “Viewing session properties” on page 756.

New session type for SAN Volume Controller and Storwize
The Global Mirror Failover/Failback with Change Volumes session is available for SAN Volume Controller and IBM Storwize storage systems that use change volumes. Change volumes contain point-in-time images that are copied from the host and target volumes. When you use change volumes, you can reduce replication traffic because only data that was changed is replicated.

See “Global Mirror” on page 689 and “Global Mirror Failover/Failback with Change Volumes session properties” on page 777.

WebSphere Application Server Liberty Profile replaced WebSphere Application Server IBM WebSphere Application Server Liberty Profile replaced WebSphere Application Server as the embedded server for Tivoli Storage Productivity Center for Replication. WebSphere Application Server Liberty Profile provides a dynamic, small footprint run time to maximize system resources.

As a result of this change, the directory structure for many Tivoli Storage Productivity Center for Replication files changed to *TPC_install_directory*\wlp\usr\servers\replicationServer\.

Open HyperSwap to be deprecated in future releases of Tivoli Storage Productivity Center for Replication
The Open HyperSwap feature is scheduled to be deprecated for future releases of Tivoli Storage Productivity Center for Replication. More information about this change will be provided in future documentation.

Limitations and known issues for Tivoli Storage Productivity Center

For information about limitations and known issues that you might encounter when you use IBM® Tivoli® Storage Productivity Center V5.2, see the *Limitations and known issues* document on the IBM Support Portal at http://www-947.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center. Search for *limitations and known issues*. This web site also includes other product support documentation, including alerts, APARs, and supported hardware and software.

Chapter 1. Getting started

Learn about the key tasks for setting up Tivoli Storage Productivity Center for managing the resources and infrastructure in a storage environment. The getting started guide provides a list of tasks in sequential order for helping you to learn how to install, configure, and use Tivoli Storage Productivity Center.

Complete the following tasks to install and set up Tivoli Storage Productivity Center for monitoring and managing storage resources. Click the included links to learn more information about each task.

Table 2. Tasks for getting started with Tivoli Storage Productivity Center

Task	Description	Learn more
Determine prerequisites and plan for installation.	The time that you invest in planning directly affects your ease of implementation. Planning includes the following tasks: • Ensuring that your systems meet hardware and software requirements • Becoming familiar with the operating systems, storage resources, agents, communication protocols, and system configurations that can be used with Tivoli Storage Productivity Center • Ensuring that ports are open for use by Tivoli Storage Productivity Center components • Understanding authentication and role-based authorization • Understanding the requirements for specific types of resources	For more information about planning, go to the Tivoli Storage Productivity Center Information Center and search for the following terms: • *hardware requirements* • *software requirements*. • *default TCP ports* • *planning for role-based authorization* • *planning for resources* For more information about the software and resources that are supported by Tivoli Storage Productivity Center, see Tivoli Storage Productivity Center Interoperability Matrix.

© Copyright IBM Corp. 2001, 2013

Table 2. Tasks for getting started with Tivoli Storage Productivity Center (continued)

Task	Description	Learn more
Install Tivoli Storage Productivity Center.	You can install Tivoli Storage Productivity Center in the following environments: **Single-server environment** Install Tivoli Storage Productivity Center and its components on a single server. **Multiple-server environment** A multiple-server environment is ideal if you are monitoring large storage environments, where one server is not sufficient to manage all the Tivoli Storage Productivity Center components. When you install Tivoli Storage Productivity Center, the following components are also installed: **DB2 (required)** Tivoli Storage Productivity Center stores the information that is collects about storage resources in a DB2 database. **Jazz for Service Management and Tivoli Common Reporting (optional)** Install these components if you want to use Tivoli Storage Productivity Center reports in the web-based GUI.	You can run the installation process as a GUI wizard or through a command line in silent mode. For more information about installing Tivoli Storage Productivity Center, go to the Tivoli Storage Productivity Center Information Center and search for the following terms: • *installing DB2* • *single-server environment*. • *multiple-server environment* • *installing Jazz for Service Management*

Table 2. Tasks for getting started with Tivoli Storage Productivity Center (continued)

Task	Description	Learn more
Access the Tivoli Storage Productivity Center GUIs.	Tivoli Storage Productivity Center provides the following GUIs for managing and monitoring storage resources: **Web-based GUI** This interface runs in a web browser and includes the ability to monitor, troubleshoot, and report on the condition of storage resources. In 5.2 and later. It also includes many of the functions for managing storage, disks, and fabrics. You can access this interface from anywhere that you have a web browser and connectivity to a network. **Stand-alone GUI** The stand-alone GUI is the same, Oracle Java based interface that was available in Tivoli Storage Productivity Center v4.2.2.3 and earlier. In version 5.2 and later, many of the functions for monitoring and managing storage in this interface were moved to the web-based GUI. **Tivoli Storage Productivity Center for Replication interfaces** Tivoli Storage Productivity Center for Replication provides a GUI that runs in a web browser and a command-line interface. Use these interfaces to manage data replication and disaster recovery.	For more information about how to access the GUIs and the functions that are available in each, see the following topics: • "Starting Tivoli Storage Productivity Center" on page 41 • "Available functions in the interfaces" on page 34

Table 2. Tasks for getting started with Tivoli Storage Productivity Center (continued)

Task	Description	Learn more
Add resources for monitoring.	Before you can monitor the resources in your environment, you must add them to Tivoli Storage Productivity Center. **Web-based GUI** In the web-based GUI, you can add the following resources for monitoring: • Storage systems • Servers (without Storage Resource agents) • Hypervisors • Switches • Fabrics **Stand-alone GUI** In the stand-alone GUI, you can add the following resources for monitoring: • IBM SONAS systems • Servers (with Storage Resource agents) • NAS filers • Tivoli Storage Productivity Center servers as subordinate servers	For more information about adding resources for monitoring, see "Adding resources" on page 257.
Deploy Storage Resource agents on servers (optional).	In the stand-alone GUI, you can deploy Storage Resource agents on servers that you want to monitor. When you deploy Storage Resource agents, the following data management tasks are enabled: • Collecting storage usage and trending information about servers (scans) • Collecting database application information • Discovering the fabrics that are visible to a server • Enforcing storage policies • Collection zone information	For more information about Storage Resource agents, see "Storage Resource agents" on page 83. For a list of the servers where you can deploy Storage Resource agents, see Tivoli Storage Productivity Center Interoperability Matrix and go to the *Agents, Servers, and Browsers* section.

Table 2. Tasks for getting started with Tivoli Storage Productivity Center (continued)

Task	Description	Learn more
Collect data about resources.	Determining the information that you want to gather about storage resources is critical to helping you implement a storage management strategy. When you determine the types of information you want to collect about your resources, you can select the data collection jobs that you want to run. In the web-based GUI, probes and performance monitors are automatically scheduled when you add resources for monitoring: **Probes** Probes collect status, asset, and capacity information about resources. **Performance Monitors** Performance monitors collect metrics that measure the performance of storage systems and switches. In the stand-alone GUI, schedule pings, scans, and TPC server probes after you add resources: **Scans** Scans collect file usage and trending information about servers or hypervisors. **Pings** Pings collect information about the availability (uptime) of servers. **TPC Server probes** TPC server probes gather the information that was collected by multiple Tivoli Storage Productivity Center servers.	For more information about data collection, see the following topics: • "Adding resources" on page 257 • "Collecting data" on page 295 • "Collect storage statistics from Tivoli Storage Productivity Center servers" on page 331
Configure data retention.	In the stand-alone GUI, configure how long to retain data that Tivoli Storage Productivity Center collects about resources. By specifying a number for days, weeks, or months for each type of resource, you can control the amount of data that is retained and available for historical analysis and charting. The longer you keep the data, the more informative your analysis.	For more information about configuring data retention, go to the Tivoli Storage Productivity Center Information Center and search for the following terms: • *resource history retention* • *removed resource retention*

Table 2. Tasks for getting started with Tivoli Storage Productivity Center (continued)

Task	Description	Learn more
Determine the conditions that must occur before you are alerted to a problem within your environment.	Determining when and how you are alerted to conditions or violations within your storage environment is important to helping you maintain and administer storage resources. Alerts are triggered by conditions that are detected during data collection and event processing. When an alert is triggered, you can choose to be notified through an SNMP trap, IBM Tivoli Enterprise Console event, log in notification, event log (Windows), syslog (UNIX), or email. To set up and manage alerts, complete the following tasks: • In the stand-alone GUI, specify the conditions and violations for each type of resource that you want to monitor through alerts. • In the web-based GUI, view, acknowledge, and remove alerts that are detected. For alerts that are triggered by performance threshold violations, you can view a chart of the historical performance of the resource that violated a threshold.	For more information about alerts, see the following topics: • Creating alerts in the stand-alone GUI • Managing alerts in the web-based GUI • "Triggering conditions for alerts" on page 1665
Configure authentication and user access.	Determine the authentication mechanism and authorization level for users. The authentication mechanism determines the user groups to which you can assign roles. Roles determine the Tivoli Storage Productivity Center functions that are available to the users in a group. In the web-based GUI, you can assign the following roles to groups of users: • Administrator • Monitor • External Application	For more information about authorizing users, go to the Tivoli Storage Productivity Center Information Center and search for *authorizing users*.

Related reference:

"Glossary" on page 1751

Managing resources

Lean about the key tasks for getting started with monitoring and managing storage resources.

After installing Tivoli Storage Productivity Center and adding resources for monitoring, complete the following tasks to manage your storage environment.

Table 3. Tasks for monitoring and managing resources

Task	Description	Learn more
Monitor the performance of resources.	Collect and view key performance metrics about resources and specify alerts to be notified of threshold violations. Metrics can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage. Threshold violations alert you when the performance of a monitored resource falls outside of a specified range. On the resource pages of the web-based GUI, you can view performance information in a chart or table format to help you quickly identify where and when performance issues are occurring. The chart is a visual representation of how the performance of resources trend over time. You can also drill down into resources to view detailed information about the performance of their internal and related resources. In the reporting function of the web-based GUI, you can also view and create performance reports about multiple resources.	For more information about performance monitoring, see the following topics: • "Monitoring the performance of resources" on page 383 • Chapter 7, "Reporting," on page 789
Monitor the capacity of resources.	View detailed information about the capacity of monitored resources in the web-based GUI and its reporting function. You can view capacity information is available for storage systems, servers, hypervisors and their internal resources, such as volumes, pools, disks, and virtual machines. For example, on the Volumes page in the web-based GUI, you can view the total amount of storage space that is committed to a volume, the amount of allocated space that is used by a volume, and the amount of space that is allocated to a volume and is not yet used. In the reporting function of the web-based GUI, you can also view and create capacity reports about multiple resources. Use the information that is collected about storage capacity to help with the capacity planning of the storage within your organization.	For more information about viewing the capacity of resources, see the following topics: • "Viewing information about resources" on page 363 • Chapter 7, "Reporting," on page 789 • Click the Help icon in the web-based GUI to view descriptions of the capacity information that is displayed for resources.

Table 3. Tasks for monitoring and managing resources (continued)

Task	Description	Learn more
Monitor the status of resources.	Monitor the status and condition of resources to identify potential problem areas in a storage environment. The status of a resource is determined by its consolidated status and the status of its internal resources. For example, if a Warning status is detected on a volume in a storage system, that Warning status is used to determine the overall status of the storage system. This propagation of statuses provides you with a comprehensive view of the condition of your storage environment. On the dashboard in the web-based GUI, you can view a high-level summary of the statuses for each type of resource. From the dashboard, you can drill down to view the individual status for each monitored resource.	For more information about viewing the status of resources, see the following topics: • "How status is determined" on page 375 • Viewing the overall status of a storage environment • Viewing the status of individual resources • Viewing the status of internal resources
Viewing the relationships between resources.	In the data path view of the web-based GUI, you can generate a graphical or tabular representation of the relationships between host connections, servers, hypervisors, virtual machines, storage systems, and the fabrics through which they share storage. View data paths to monitor status and pinpoint problem areas in the path of monitored storage throughout your environment. In the reporting tool of the web-based GUI, view reports about the relationships between resources. For example, you can view a report about the end-to-end relationship of all the servers in your environment to back-end storage. In the stand-alone GUI, use the topology viewer to view the relationships between resources.	For more information about viewing the relationships between resources, see the following topics: • Viewing data paths in the web-based GUI • Viewing relationship reports in the web-based GUI • Using the topology viewer in the stand-alone GUI

Table 3. Tasks for monitoring and managing resources (continued)

Task	Description	Learn more
Identifying hot spots of high utilization.	Tivoli Storage Productivity Center collects detailed information about storage system pools, including utilization values. The utilization value is an estimate of the average daily utilization of the physical resources such as controllers, disks, ports, and nodes that are associated with the pool. By scanning the utilization values for pools on a storage virtualizer, such as SAN Volume Controller, you can identify the hot spots or the pools that have high utilization values on the storage virtualizer. You can also determine the course of action that you take. For example, if some of the pools on a tier have high utilization values and some have low utilization values, you can redistribute volumes across the storage tier. However, if all of the pools on tier have high utilization values, you can balance the pools by re-tiering the most active volumes.	For more information about identifying hot spots, see the following topics: • "Redistributing volume workloads to balance pools" on page 94 • Viewing the overall status of a storage environment • Viewing the overall status of a storage environment
Configuring the cloud and provisioning storage.	Configure cloud storage by creating service classes and capacity pools to classify and separate storage resources. After configuration, use Provision Storage wizard to assign storage to servers and hypervisors. The wizard guides you through the steps for provisioning volumes or shares. In the wizard, you can request the amount of storage that is needed and a service class. The service class acts as a set of requirements for the placement of the volume or share. Tivoli Storage Productivity Center identifies the storage resources that can provide the requested capacity and that can also satisfy the requirements of the service class. Tivoli Storage Productivity Center recommends storage placement that is based on storage system free space and performance data. If you agree with the recommendation, Tivoli Storage Productivity Center can implement the provisioning.	For more information about provisioning storage, see "Provisioning storage" on page 71.

Table 3. Tasks for monitoring and managing resources (continued)

Task	Description	Learn more
Manage a virtual storage environment by using the vSphere Web Client extension for Tivoli Storage Productivity Center.	If you maintain virtual storage in a VMware vSphere environment, you can use the vSphere Web Client extension for Tivoli Storage Productivity Center to complete provisioning tasks and to view information about storage resources in the vSphere Web Client. The client includes details about performance and asset, copy services, and fabric information. Tivoli Storage Productivity Center uses the vSphere API for Storage Awareness (VASA) to connect to vSphere. When you register Tivoli Storage Productivity Center as a VASA provider, you can view block and file storage information from Tivoli Storage Productivity Center in vSphere reports.	For more information about the vSphere Web Client extension for Tivoli Storage Productivity Center, see "Managing resources in a vSphere environment" on page 458.
Manage replication.	Use Tivoli Storage Productivity Center for Replication to manage the advanced copy services for storage system through the following tasks: • Setting up management servers • Adding storage connections and paths • Creating sessions and copy sets • Setting up data replication • Performing planned maintenance • Recovering from a disaster	For more information about managing replication, see Chapter 6, "Managing replication," on page 629.
Determine the storage policies that you want to enforce.	Enforcing storage policies within your organization can be critical to ensuring the most cost-effective use of your storage resources. Use the following policy management functions in the stand-alone GUI to set policies for resources that are monitored by Tivoli Storage Productivity Center: • Use quotas to define limits on the amount of storage that a user or a group of users can use. • Use constraints to define the acceptable and unacceptable file types, file sizes, and file owners for a computer or a set of computers in your environment. You must use the stand-alone GUI to deploy Storage Resource agents on the servers where you want to enforce policies.	For more information about enforcing storage policies in the stand-alone GUI, see the following topics: • Creating and enforcing storage policies • "Setting storage usage quotas" on page 507 • "Setting file constraints" on page 532

Table 3. Tasks for monitoring and managing resources (continued)

Task	Description	Learn more
Run reports to view detailed information about the condition, capacity, and performance of resources.	Tivoli Storage Productivity Center provides multiple user interfaces for viewing reports about storage. **Web-based GUI** Access the Tivoli Common Reporting interface in the web-based GUI to view predefined reports and create custom reports about monitored resources. Reports include detailed information about storage capacity, assets, and performance of storage systems, switches, fabrics, servers, hypervisors, and their internal resources. **Stand-alone GUI** View detailed information about the storage resources in your environment in the stand-alone GUI. These reports are organized into different types and categories and provide both summary and detailed information, depending on your needs. Reports are organized into report types and categories that you can expand to reveal extra subcategories of reporting. Many reports are also accessible through the topology viewer. **vSphere Web Client** Use the vSphere Web Client extension for Tivoli® Storage Productivity Center to view reports on your virtual environment and storage devices. You can view reports that are customized to use information from Tivoli Storage Productivity Center. The reports include fabric connections, storage mapping information, and performance metrics for storage systems.	For more information about reporting, see the following topics: • Chapter 7, "Reporting," on page 789 • "Reporting in the stand-alone GUI" on page 1295 • "Reporting in a vSphere environment" on page 1643

Related concepts:

Chapter 5, "Managing resources," on page 257
Use Tivoli Storage Productivity Center to monitor and manage the resources and storage infrastructure in an enterprise environment.

Related reference:

"Glossary" on page 1751

Configuring and administering

Learn about the key tasks for getting started with configuring and administering Tivoli Storage Productivity Center and its related components.

Complete the following tasks to customize Tivoli Storage Productivity Center for your environment and administer its related components.

Table 4. Tasks for configuring and administering Tivoli Storage Productivity Center

Task	Description	Learn more
Configure Tivoli Storage Productivity Center.	Configure Tivoli Storage Productivity Center according to the standards within your organization. The configuring tasks that you want to complete determines which GUI to use: **Web-based GUI** • Authorizing users and specifying an authentication mechanism • Customizing lists of resources, tasks, and alerts • Enabling reports in Tivoli Common Reporting **Stand-alone GUI** • Entering licensing information • Configuring how alerts are communicated • Specifying how long information about resources is retained • Defining rules for generating email addresses of quota violators • Setting up NAS filer access • Analyzing historical data to identify potential problems in a SAN configuration **vSphere client** • Configuring the VASA provider • Configuring the VMware vSphere web client extension	For more information about configuring, see the following topics: • Role-based authorization • "Changing the user authentication configuration" on page 146 • "Customizing lists of resources, jobs, and alerts" on page 443 • If the reporting function is not enabled in the web-based GUI, you must install Jazz for Service Management and IBM Tivoli Common Reporting. For more information, go to the Tivoli Storage Productivity Center Information Center and search for *Jazz for Service Management and Tivoli Common Reporting*. • For information about the configuring options in the stand-alone GUI, go to the Tivoli Storage Productivity Center Information Center and search for *configuration nodes*. • For information about configuring the VASA provider and vSphere web client entension, go to the Tivoli Storage Productivity Center Information Center and search for the following terms: – *configuring the VASA provider* – *configuring the vSphere Web Client*.

Table 4. Tasks for configuring and administering Tivoli Storage Productivity Center (continued)

Task	Description	Learn more
Administer DB2 and Tivoli Storage Productivity Center.	After Tivoli Storage Productivity Center is installed and configured, complete the following administering tasks as needed: • Administering CIM agents, storage systems, servers, Storage Resource agents, switches, fabrics, hypervisors, VMware data sources, and TPC Servers • Starting and stopping Tivoli Storage Productivity Center services and servers • Administering DB2 and backing up the database repository • Increasing the memory allocation and checking the status of Tivoli Storage Productivity Center components • Setting timeout values for the Device server • Changing passwords for Tivoli Storage Productivity Center components • Collecting information with the service tool	

Navigation

The web-based GUI provides many of the functions for managing a storage environment. To access the functions in this GUI, use the navigation pane on the left side of the main window. An interactive breadcrumb at the top of each page traces your location in the GUI.

Navigation pane	Description
 Home	**Dashboard** Use the dashboard to view an overall status of monitored resources and identify potential problem areas in a storage environment. You can learn about the status and capacity of resources, the jobs that collect performance data, and the unacknowledged alert conditions that are detected on resources. **Performance Monitors** Use performance monitors to collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations. **Alerts** Use alerts to be notified when certain conditions or violations are detected on monitored resources. In the stand-alone GUI, define alerts and thresholds. In the web-based GUI, view the alerts that were detected and the thresholds that were violated. **Tasks** Manage the tasks that are created when you complete the following actions: • Assign storage to servers or hypervisors. • Place volumes on the tiers that match the workload requirements of the volume. • Move or convert volumes. • Move volumes from overutilized pools to pools that are less used.
 Storage Resources	Add storage systems for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer storage systems: • Modify the schedules for data collection. • View the data path of connected resources. • Assign storage to servers or hypervisors. • Place volumes on the tiers that match the workload requirements of the volume. • Move or convert volumes. • Move volumes from overutilized pools to pools that are less used. • Add resources to a capacity pool. • Open management GUIs. • Update login credentials. • Acknowledge statuses.

Navigation pane	Description
 Server Resources	Add agentless servers and hypervisors for monitoring and view asset and status information about their internal and related resources. You can also complete the following actions to manage and administer these resources: • Modify the schedules for data collection. • View the data path of connected resources. • Assign storage to servers or hypervisors. • Modify and upgrade Storage Resource agents. • Acknowledge statuses. **Tip:** Use the stand-alone GUI if you want to deploy Storage Resource agents to servers or virtual machines. For more information, see "Adding servers" on page 262.
 Network Resources	Add switches and fabrics for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer these resources: • Modify the schedules for data collection. • View the data path of connected resources. • Setting or modifying the zoning policy. • Open the management GUI for a switch. • Acknowledge statuses.
 Advanced Analytics	Configure Tivoli Storage Productivity Center for provisioning and optimizing storage. Configuration tasks include creating service classes and capacity pools to classify and separate resources. The web-based GUI includes the following wizards to guide you through the steps of managing your storage: • Use the Provision Storage wizard to assign storage to servers or hypervisors. • Use the Analyze Tiering wizard to manage storage resources efficiently and reduce storage costs. • Use the Analyze Tiering, Balance Pools, and Transform Storage wizards to optimize and transform storage volumes.

Navigation pane	Description
 Reporting	Create reports and view predefined reports to analyze multiple storage systems, switches, servers, and hypervisors. For example, you can compare performance metrics across multiple storage systems. You can also view a report about the aggregate volume load that is on a server or hypervisor. Use the reporting tool to schedule reports to run on different systems and with different options. You can specify schedules and output formats for the reports, and different ways to share the reports. Use the web-based GUI when you want to individually view details about resources, for example when you want to troubleshoot a resource. When you click a reporting option, a separate browser tab is opened. For information about how to get started with reporting tool, see Getting started with reporting.
 Settings	Modify the authentication repository and assign Tivoli Storage Productivity Center roles to user groups. Roles determine the product functions that are available to users.

Navigating in the stand-alone GUI: The stand-alone GUI provides a different of navigation tools for accessing its functions. For more information, see Overview of the stand-alone GUI.

Keyboard navigation in the web-based GUI

Most of the features of the web-based GUI are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI), except as noted in the product release notes.

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. The following sections describe the keys or key combinations for different parts of the web-based GUI:

For navigating the web-based GUI and the context-sensitive help system:

- To navigate to the next link, button, or topic within a panel, press Tab.
- To move to the previous link, button, or topic within a panel, press Shift+Tab.
- To select an object, when the object is in focus, press Enter.

For actions menus:

- To navigate to the grid header, press Tab.
- To reach the drop-down field, press the Left Arrow or Right Arrow key.
- To open the drop-down menu, press Enter.
- To select the menu items, press the Up Arrow or Down Arrow key.
- To start the action, press Enter.

For filters:

To specify a filter option and text:

1. Press Tab to navigate to the magnifying glass icon.
2. Press the Up Arrow or Down Arrow key to navigate the filtering list.
3. Press Enter to select a filtering option.
4. When a filtering option is selected, the cursor moves to the filter text box. Type the filter text and press Enter. To reset a filter, press Enter.

For text fields:

- To navigate to text fields, press Tab.
- To navigate to the fields that are available for editing, press Tab.
- To navigate to the next field or to the **Submit** button, press Tab.

For tables or lists:

- To navigate between column headers, focus on a column header and use the Left Arrow and Right Arrow keys to move to other column headers.
- To navigate between data cells, focus on a data cell and use the Left, Right, Up, Down, Pageup, and Pagedown Arrow keys.
- To sort a column, focus on a column header and press Enter. The focus remains on the column header after the sort occurs.
- To change the size of a column, focus on the column header, hold Shift+Control, and press the Left or Right Arrow keys.
- To follow a link in a data cell, focus on a data cell and press Shift+F9.
- To open a menu for a table row, focus on the row and press Shift+F10.
- To select consecutive rows, select the first row and hold Shift, press the Up or Down Arrow keys to go to the last row in the range, and press the Space bar to add the new rows to the selection.
- To select non-consecutive rows, select a row and hold Control, press the Up or Down Arrow keys, and press the Space bar to add the new row to the selection.

Restriction: For Chinese languages, the keyboard combination Control+Space bar is not enabled for selecting multiple rows at the same time.

Chapter 2. Product overview

IBM Tivoli Storage Productivity Center provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance, and resources. It can reduce the complexity of managing a storage environment by centralizing, simplifying, and optimizing storage tasks that are associated with storage systems, storage networks, replication services, and capacity management.

Because it is policy-based, Tivoli Storage Productivity Center can detect potential problems and automatically adjust storage and resources that are based on the policies and actions that you define. For example, it can notify you when a server or storage system is running out of disk space or warn you of impending storage hardware failure. By alerting you to these and other issues that are related to your stored data, you can use Tivoli Storage Productivity Center to prevent unnecessary system and application downtime.

Tivoli Storage Productivity Center provides the following benefits and functions:

- Simplifies the management of storage infrastructures
- Manages, administers, and provisions SAN-attached storage
- Monitors and tracks performance of SAN-attached devices
- Monitors, manages, and controls (through zones) SAN fabric components
- Manages the capacity utilization and availability of file systems and databases

Architecture

The IBM Tivoli Storage Productivity Center consists of several components that form the infrastructure of its storage-management functions.

Data server
: This component is the control point for product scheduling functions, configuration, event information, reporting, and graphical user interface (GUI) support. It coordinates communication with and data collection from agents that scan file systems and databases to gather storage demographics and populate the database with results. Automated actions can be defined to perform file system extension, data deletion, and Tivoli Storage Productivity Center backup or archiving, or event reporting when defined thresholds are encountered. The Data server is the primary contact point for GUI user interface functions. It also includes functions that schedule data collection and discovery for the Device server.

Device server
: This component discovers, gathers information from, analyzes performance of, and controls storage subsystems and SAN fabrics. It coordinates communication with and data collection from agents that scan SAN fabrics and storage devices.

Replication server
: This component coordinates communication and processes tasks that are related to replication and Tivoli Storage Productivity Center for Replication. It also provides the infrastructure for hosting the Tivoli Storage Productivity Center for Replication GUI.

© Copyright IBM Corp. 2001, 2013

Web server
Tivoli Storage Productivity Center uses the WebSphere Application Server as the web application server to host its web-based GUI, storage management API for cloud, and VMware vSphere web client extension.

Jazz for Service Management and Tivoli Common Reporting (optional)
These components enable Tivoli Storage Productivity Center reports in the web-based GUI. You can use this reporting tool to create reports and view predefined reports to analyze multiple storage systems, switches, servers, and hypervisors.

For large storage environments where one server is not sufficient to manage all the Tivoli Storage Productivity Center components, you can install the product on multiple servers. In this type of environment, Jazz for Service Management and Tivoli Common Reporting can run on a separate server.

Database
A single database instance serves as the repository for all Tivoli Storage Productivity Center components.

Agents
Storage Resource agents, CIM agents, and Out of Band fabric agents gather host, application, storage system, and SAN fabric information and send that information to the Data Server or Device server.

GUI Tivoli Storage Productivity Center provides two graphical user interfaces for managing the storage infrastructure in an enterprise environment. a stand-alone GUI and the web-based GUI. Each GUI provides different functions for working with monitored resources. To view a comparison of the functions that are available in each GUI, see "Available functions in the interfaces" on page 34.

CLI Use the command-line interface (CLI) to issue commands for key Tivoli Storage Productivity Center functions.

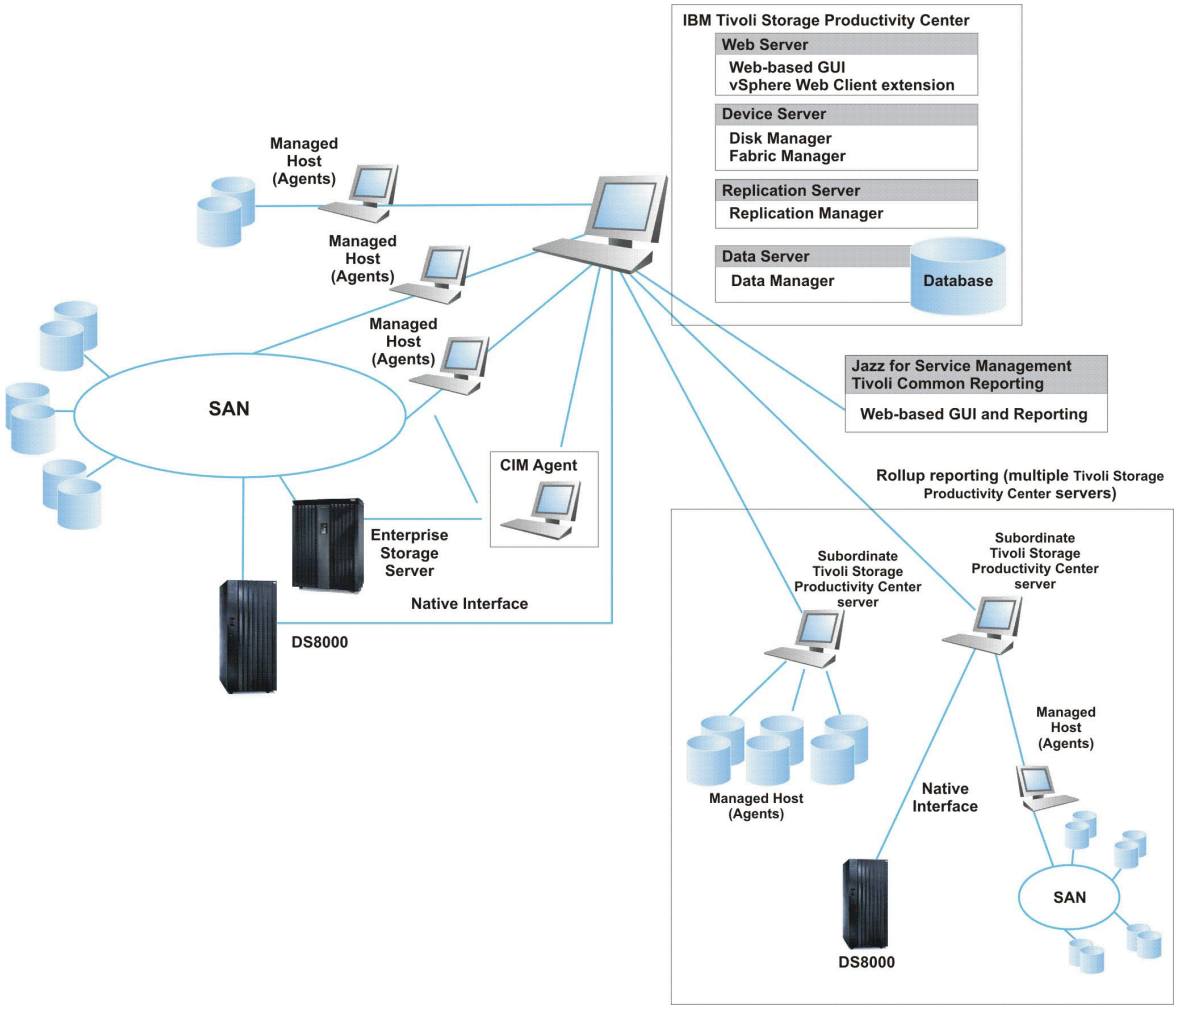

Figure 1. Tivoli Storage Productivity Center components

Product licenses

Product licenses are available for IBM Tivoli Storage Productivity Center. Each license determines the functions that are accessible in the stand-alone GUI, web-based GUI, command-line interface, and the API. You can review the functions that are included with each license and determine which license meets your storage management needs.

Restriction: When you upgrade your license, all of the servers, except the web server are recycled. After you upgrade your license, to access the functionality in the license, you must stop and restart the web server.

IBM Tivoli Storage Productivity Center

This license contains the following functions:

Data management

A file and capacity management solution for heterogeneous storage environments. Data management includes enterprise-wide reporting and monitoring, and automated capacity provisioning for

Direct Attached Storage (DAS), network-attached storage (NAS), and Storage Area Network (SAN) environments.

Disk management
Disk management helps you manage SANs and heterogeneous storage from one console.

Fabric management
Fabric management helps you manage the SAN fabric that connects the host systems and applications to the storage devices. This feature is a comprehensive management solution for multi-vendor SANs and includes automatic resource and topology discovery, monitoring, and alerts.

Tivoli Storage Productivity Center for Replication
This license includes all the Tivoli Storage Productivity Center for Replication functions.

Performance management
Disk and Fabric performance management functions.

Device support
You can see a complete list of devices that can be used with Tivoli Storage Productivity Center at https://www.ibm.com/support/docview.wss?uid=swg21386446. Click the appropriate release under **Storage**.

IBM Tivoli Storage Productivity Center Select Edition
This license contains all the functions of the IBM Tivoli Storage Productivity Center license. This license is priced by enclosure rather than by capacity.

IBM SmartCloud Virtual Storage Center Storage Analytics Engine
This license contains everything that is in the IBM Tivoli Storage Productivity Center license, and also includes advanced analytical functions. This license was previously called IBM Tivoli Storage Productivity Center Advanced. The user interfaces and messages still refer to this license as IBM Tivoli Storage Productivity Center Advanced.

In the web-based GUI, the following advanced analytical functions are available:

- Optimize storage tiering by using the Analyze Tiering wizard
- Distribute the workload of volumes across pools on the same tier by using the Balance Pools wizard
- Provision block storage by using the Provision Storage wizard.
- Enable automatic zoning to create zones during block storage provisioning to connect a server to a storage system

In the stand-alone GUI, the following advanced analytical functions are available:

- Configuration analysis
- Configuration history
- Scans
- Policy management
- Server probes
- Scans
- Profiles

- Data Manager for databases
- Data Manager for Chargeback

This license is only available as part of the IBM SmartCloud Virtual Storage Center (VSC). The complete VSC solution bundles Tivoli Storage Productivity Center with IBM System Storage SAN Volume Controller and Tivoli Storage FlashCopy Manager. In the VSC solution, System Storage SAN Volume Controller provides a virtualization platform and remote replication functions, and the Tivoli Storage FlashCopy Manager provides data backup and restore capabilities. For more information about IBM SmartCloud Virtual Storage Center, see http://www.ibm.com/software/products/us/en/vsc/.

If you purchased a license for IBM Tivoli Storage Productivity Center Version 4.2, when you upgrade to Version 5.2, you can upgrade to the appropriate license. Table 5 shows the parallel functions in Version 5.2.

Table 5. Tivoli Storage Productivity Center licenses when you upgrade from Version 4.2 to Version 5.2

If you have this Tivoli Storage Productivity Center Version 4.2 license	You are entitled to this Tivoli Storage Productivity Center Version 5.2 license
Basic	When you upgrade to Version 5.2, you get the same functionality that you had with the Version 4 Basic Edition license.
Disk Select	IBM Tivoli Storage Productivity Center Select Edition
Data	Not applicable.
Disk	IBM Tivoli Storage Productivity Center
Standard Select	Not applicable.
Standard	Not applicable.

Functions available in the web-based GUI based on role and license

When you use the web-based GUI, your Tivoli Storage Productivity Center role and product license determine the functions that are available.

Users who are assigned the Administrator role or the Monitor role can use the web-based GUI. The functions that are available depend on the role that is assigned to the user:

Administrator role

Users who are assigned the Administrator role have access to all monitoring and administrative functions and are limited only by license restrictions.

Monitor role

Users who are assigned the Monitor role can view information about monitored resources and other objects such as tasks, alerts, and service classes. They can acknowledge alerts and resource statuses, open logs, and open management GUIs. Users who are assigned the Monitor role do not have access to administrative functions, with the following exceptions:

- They can assign storage pools to tiers.

- They can be granted permission in a service class to provision storage by using the service class. If so, they can use the Provision Storage wizard to create a provisioning task. Users can delete provisioning tasks that they create. If the service class specifies that administrator approval is not required, the users can run or schedule the provisioning tasks that they create.

To use some functions of the web-based GUI, you must have the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license. If you have the Tivoli Storage Productivity Center license or the Tivoli Storage Productivity Center Select license, the following functions are not available:

- Storage tier optimization. The Analyze Tiers wizard is not available.
- Pool balancing. The Balance Pools wizard is not available.
- Block storage provisioning. You are not able to provision volumes by using either the Provision Storage wizard or the vSphere Web Client extension. However, there are no license restrictions for file storage provisioning, therefore, Network Attached Storage (NAS) file shares can be provisioned.

The following table outlines the functions available by role. Some functions are restricted to the Administrator role. Functions that are restricted to the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license are also shown. In addition to the restrictions listed in this table, users who are assigned the Monitor role do not have access to user management functions.

Table 6. Role and license restrictions in the web-based GUI

Resource or object	**Functions available by role**	**Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license**
Performance monitors	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about the switch or storage system that is being monitored • Open performance monitor logs • Export performance monitor information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Start or stop a performance monitor • Schedule a performance monitor	All functions are available, depending on role.
Alerts	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about alerts • Acknowledge alerts • Export alert information to a file Users who are assigned the Administrator role can also remove alerts.	All functions are available, depending on role.

Table 6. Role and license restrictions in the web-based GUI (continued)

Resource or object	Functions available by role	Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license
Tasks	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about tasks • Open task logs • Export task information to a file If a user assigned the Monitor role has permission to provision storage by using a service class, the user can delete provisioning tasks that the user creates. If the service class does not require administrator approval, the user can also schedule and run the provision tasks that the user creates. Users who are assigned the Administrator role can also complete the following administrative actions: • Run tasks • Schedule tasks • Delete tasks	All functions are available, depending on role.
Storage systems	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about storage systems • Acknowledge storage system status • Open the management GUI or element manager for a storage system • Export storage system information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Add and remove storage systems • Schedule data collection jobs • Optimize storage tiering • Add storage systems to capacity pools • Administer connections	Optimize storage tiering
Volumes	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about volumes • Acknowledge volume status • Export volume information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Transform storage volumes • Optimize storage tiering	Optimize storage tiering

Table 6. Role and license restrictions in the web-based GUI (continued)

Resource or object	Functions available by role	Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license
Pools	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about pools • Acknowledge pool status • Export pool information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Add pools to capacity pools • Balance pools • Optimize storage tiering	• Balance pools • Optimize storage tiering
Shares	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about shares • Export pool information to a file	All functions are available, depending on role.
Servers	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about servers • Acknowledge server status • View and collect logs • Export server information to a file If a user assigned the Monitor role has permission to provision storage by using a service class, the user can provision storage to servers. Users who are assigned the Administrator role can also complete the following administrative actions: • Add and remove servers • Schedule data collection jobs • Provision storage to servers • Modify Storage Resource agents	• Provision block storage. There are no license restrictions for provisioning file storage. • Optimize storage tiering.

Table 6. Role and license restrictions in the web-based GUI (continued)

Resource or object	Functions available by role	Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license
Hypervisors	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about hypervisors • Acknowledge hypervisor status • Export hypervisor information to a file If a user who is assigned the Monitor role has permission to provision storage by using a service class, the user can provision storage to hypervisors. Users who are assigned the Administrator role can also complete the following administrative actions: • Add and remove hypervisors • Schedule data collection jobs • Provision storage to hypervisors • Administer connections	Provision block storage. There are no license restrictions for provisioning file storage.
Switches	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about switches • Acknowledge switch status • Open the management GUI or element manager for a switch • Export switch information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Add and remove switches • Schedule data collection jobs • Administer connections • Enable automatic zoning	Enable automatic zoning
Fabrics	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about fabrics • Acknowledge fabric status • Export fabric information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Add and remove fabrics • Schedule data collection jobs • Administer connections • Enable automatic zoning	Enable automatic zoning

Table 6. Role and license restrictions in the web-based GUI (continued)

Resource or object	Functions available by role	Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license
Service classes	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about service classes • Export service class information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Create service classes • Modify service classes • Delete service classes	Create block-storage service classes. There are no license restrictions for working with file-storage service classes.
Capacity pools	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about capacity pools • Export capacity pool information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Create capacity pools • Modify capacity pools • Delete capacity pools	All functions are available, depending on role.

Functions available in the stand-alone GUI based on role and license

When you use the stand-alone GUI, your Tivoli Storage Productivity Center role and product license determine the functions that are available.

Users who are assigned the Administrator role or the Monitor role can use the stand-alone GUI. The functions that are available depend on the role that is assigned to the user:

Administrator role
Users who are assigned the Administrator role have access to all monitoring and administrative functions and are limited only by license restrictions.

Monitor role
Users who are assigned the Monitor role can view data that is collected by Tivoli Storage Productivity Center, and can create, generate, and save reports.

To use some functions of the stand-alone GUI, you must have the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license. The following table outlines the functions that are restricted to the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license. The table mirrors the navigation tree of the stand-alone GUI. The table shows nodes of the navigation tree and indicates

whether that function is available with all licenses, or whether the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license is required.

Table 7. Functions with each Tivoli Storage Productivity Center license

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
Administrative Services		
Services	Yes	Yes
Data Sources[1] • CIMOM agents • Storage Resource agents • Storage Subsystems • TPC Servers • VMWare VI Data Source	 • Yes • Yes • Yes • No • Yes	 • Yes • Yes • Yes • Yes • Yes
Discovery	Yes	Yes
Configuration	Yes[2]	Yes
IBM Tivoli Storage Productivity Center		
Job Management	Yes	Yes
Reporting • My Reports (User ID reports) • System Reports • Rollup Reports • Data Source Reports	Yes See Table 8 on page 31 for a list of available reports.	Yes See Table 8 on page 31 for a list of available reports.
Topology Viewer	Yes	Yes
Monitoring (data collection) • Probe Alert Configurations • TPC Server Probes	 • Yes • No	 • Yes • Yes
Storage Resource Group Management	Yes[5]	Yes
Analytics • Configuration History • Configuration Analysis	 • No • No	 • Yes • Yes
Alerting: • Authentication Configuration Alerts • Alert Log	 • Yes • Yes[3]	 • Yes • Yes
Data Manager		
Monitoring: (data collection) • Groups • Pings • Scans • Profiles	 • Yes • Yes • No • No	 • Yes • Yes • Yes • Yes
Alerting	Yes[4]	Yes
Policy Management	No	Yes

Table 7. Functions with each Tivoli Storage Productivity Center license (continued)

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
Reporting	Yes Not all reports are available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license. See Table 8 on page 31 for a list of available reports.	Yes
Data Manager for Databases	No	Yes
Data Manager for Chargeback	No	Yes
Disk Manager		
Storage Subsystem	Yes	Yes
Monitoring (data collection) • Groups • Jobs • Subsystem Performance Alert Configurations	 • Yes • Yes • Yes	 • Yes • Yes • Yes
Alerting	Yes	Yes
Reporting • Groups • Storage Subsystems • Storage Subsystem Performance	 • Yes • Yes • Yes	 • Yes • Yes • Yes
Fabric Manager		
Monitoring (data collection) • Groups • Switch Performance Alert Configurations	 • Yes • Yes	 • Yes • Yes
Fabrics	Yes	Yes
Alerting	Yes	Yes
• Switch Performance (See “Creating switch performance reports” on page 1503)	Yes	Yes
Replication Manager	Yes	Yes

Table 7. Functions with each Tivoli Storage Productivity Center license (continued)

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
Note: 1. You cannot associate TPC Servers with Tivoli Storage Productivity Center in the Tivoli Storage Productivity Center or Storage Productivity Center Select license. VMware Data Sources are associated with the Tivoli Storage Productivity Center or Storage Productivity Center Select license. TPC Servers are associated with the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license. 2. The following nodes under **Configuration** are not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: • License Keys • Quota and Constraint e-mail Address Rules • Resource History Retention for Databases • Removed Resource Retention for Databases • Configuration History Settings 3. The following nodes under **Alert Log** are not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: • Directory • User • OS User Group • Configuration Analysis 4. The following node under **Data Manager** > **Alerting** is not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: Directory Alerts. Hypervisor Alerts and Other NAS Alerts are available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license. 5. You cannot create a storage resource group for a type of capacity pool.		

Table 8. Available reports with each IBM Tivoli Storage Productivity Center license

Reports	Tivoli Storage Productivity Center or Storage Productivity Center Select	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
Batch reports	Yes[1], [2]	Yes[1]
System Reports: • Data • Fabric • Disk	 • Yes[4] • Yes • Yes	 • Yes • Yes • Yes
Rollup Reports	No	Yes
Data Source Reports	Yes	Yes
Asset reports:		

Table 8. Available reports with each IBM Tivoli Storage Productivity Center license (continued)

Reports	Tivoli Storage Productivity Center or Storage Productivity Center Select	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
• By Cluster • By Computer • By Hypervisor • By IBM Storwize V7000 Unified/IBM SONAS • By OS Type • By Storage Subsystem • System-wide	• Yes • Yes • Yes • Yes • Yes • Yes • Yes[3]	• Yes • Yes • Yes • Yes • Yes • Yes • Yes
Availability reports	Yes	Yes
TPC-wide Storage Space reports: • Disk Space • File System Space • Consumed File System Space • Available File System Space	 • Yes • Yes • Yes • Yes	 • Yes • Yes • Yes • Yes
Usage reports	No	Yes (these reports are a result of a scan)
Usage Violation reports	No	Yes (these reports are a result of a scan)
Backup reports	No	Yes (these reports are a result of a scan)
Monitored Computer Storage Space reports	Yes	Yes
Storage Subsystem reports	Yes	Yes
Storage Subsystem Performance reports	Yes	Yes
Switch Performance reports	Yes	Yes

Table 8. Available reports with each IBM Tivoli Storage Productivity Center license (continued)

Reports	Tivoli Storage Productivity Center or Storage Productivity Center Select	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
Note: 1. A Storage Resource agent must be deployed on the computer where you want to run batch reports. 2. The following batch reports are not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: • Rollup • Usage • Usage Violations • Backup 3. The following **Asset** > **System-wide** reports are available with the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license: • Monitored Directories • Unmanaged Computers • Users • OS User Groups 4. The following Data system reports are available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: • Disk Space Summary • Disk Defects • Computer Storage Availability • Computer Disk Space • Available File System Space		

Interfaces for Tivoli Storage Productivity Center

Tivoli Storage Productivity Center provides multiple user interfaces for managing the storage infrastructure in an enterprise environment.

Web-based GUI

This interface runs in a web browser and includes the ability to monitor, manage, and troubleshoot storage resources. You can also report on the condition, capacity, and performance resources in the Tivoli Common Reporting interface. In 5.2 and later, the web-based GUI includes many of the functions for managing storage, disks, and fabrics. You can access this interface from anywhere that you have a web browser and connectivity to a network. For more information about the functions that are available in this GUI, see "Available functions in the interfaces" on page 34.

Tivoli Common Reporting interface

This interface runs in a web browser. Use this interface to view predefined reports and create custom reports about Tivoli Storage Productivity Center. You access reports from the web-based GUI, and work with the reports in Tivoli Common Reporting. Tivoli Common Reporting includes Cognos® reporting software. For more information about reporting in the web-based GUI, see Chapter 7, "Reporting," on page 789.

Stand-alone GUI

The stand-alone GUI is the same, Java based interface that was available in Tivoli Storage Productivity Center v4.2.2.x and earlier. In version 5.2 and later, many of the functions for monitoring and managing storage in this interface were moved to the web-based GUI. For more information about the functions that were moved from this GUI to the web-based GUI, see "Functions that were moved to the web-based GUI in V5.2" on page 38.

Command-line interface

Use this interface to run Tivoli Storage Productivity Center commands from a command prompt. For more information about CLI commands, see the *IBM Tivoli Storage Productivity Center Command-Line Interface Reference*.

Tivoli Storage Productivity Center for Replication interface

Tivoli Storage Productivity Center for Replication provides a GUI that runs in a web browser and a command-line interface. Use these interfaces to manage data replication and disaster recovery. For more information about these interfaces, see "Interfaces for Tivoli Storage Productivity Center for Replication" on page 48. For more information about the features that are available in Tivoli Storage Productivity Center for Replication, see Chapter 6, "Managing replication," on page 629.

Related concepts:

Chapter 5, "Managing resources," on page 257
Use Tivoli Storage Productivity Center to monitor and manage the resources and storage infrastructure in an enterprise environment.

Available functions in the interfaces

Tivoli Storage Productivity Center provides the web-based GUI and stand-alone GUI for managing storage resources. In V5.2, many of the functions for managing storage were moved to the web-based GUI. Determine which GUI to use based on the tasks that you must complete.

Configuring and administering

Use the following list to determine the configuring and administering tasks that you can complete in each GUI.

Table 9. Available configuring and administering tasks in each GUI

Tasks	Available in web-based GUI	Available in stand-alone GUI
Discovering resources for monitoring	Yes[1]	Yes[1]
Adding resources for monitoring	Yes[2]	Yes[2]
Removing resources (data sources)	Yes[3]	Yes[3]
Administering resources (data sources)	Yes: • Viewing and modifying connections • Testing connections • Updating credentials	Yes, but only for the following actions and data sources: • Showing managed resources: CIM agents, VMware vSphere • Testing connections: storage systems, TPC servers • Launching GUI: TPC servers

Table 9. Available configuring and administering tasks in each GUI (continued)

Tasks	Available in web-based GUI	Available in stand-alone GUI
Administering Tivoli Storage Productivity Center services	No	Yes
Administering servers and Storage Resource agents	Yes	Yes, but only to view logs
Configuring Tivoli Storage Productivity Center	Yes: • Assigning roles to user groups • Modifying the authentication repository • Upgrading Storage Resource agents	All other configuration tasks
Viewing and managing tasks	Yes, for the following types of tasks: • Provisioning • Optimizing	All other tasks
Viewing and managing data collection jobs	Yes, for the following types data collection jobs: • Probes • Performance monitors • Agent upgrades	All other jobs
Opening the management GUIs for resources	You can open the start page of the management GUI for a resource. The GUI must be web-based or use Java Web Start.	No
Notes for configuring and administering tasks: 1. In the web-based GUI, the discovery process is completed automatically when you add resources. In the stand-alone GUI, you must run a discovery job to find resources before you can add them. 2. In the web-based GUI, you can add the following resources for monitoring: storage systems (other than IBM SONAS storage systems), servers (without a Storage Resource agent), hypervisors, fabrics, and switches. In the stand-alone GUI, you can add the following resources: servers (with a Storage Resource agent), IBM SONAS storage systems, and Tivoli Storage Productivity Center servers as subordinate servers. 3. In the web-based GUI, you can remove the following resources: servers, Hypervisors, Storage systems, Switches, Fabrics, NAS filers, CIM agents for storage systems, switches, fabrics, and NAS filers. In the stand-alone GUI, you can remove the following resources: CIM agents, VMware vSphere servers, and Tivoli Storage Productivity Center servers as subordinate servers.		

Monitoring

Use the following list to determine the monitoring tasks that you can complete in each GUI.

Table 10. Available monitoring tasks in each GUI

Tasks	Available in web-based GUI	Available in stand-alone GUI
Scheduling data collection	Yes: • Probes • Performance monitors	Yes: • Pings • Scans
Creating groups of resources	No	Yes
Creating storage resource groups	No	Yes
Viewing information about resources	Yes, except databases	Yes
Viewing the following information about resources: • Status • Capacity (including used and available space) • Relationships between resources • Data paths	 • Yes • Yes • Yes • Yes	 • Yes • Yes • Yes • Yes
Exporting information about resources to a file	Yes	Yes
Viewing a visual topology of resources	Yes, for storage systems, servers, and hypervisors	Yes
Viewing a visual data path of resources	Yes	Yes
Troubleshooting performance issues	Yes	Yes, but only for the following actions: • Defining alerts to notify you when performance conditions are detected or thresholds are violated. • Viewing the performance reports that were available in previous versions.
Running analytics and configuring storage: • Configuration history • Configuration analysis • Optimizing and balancing storage • Provisioning storage • Cloud configuration	 • No • No • Yes • Yes • Yes	 • Yes • Yes • No • No • No
Managing volumes for storage systems[1]: • Viewing information about volumes • Creating volumes • Deleting volumes • Assigning volumes to host ports	 • Yes • Yes, when provisioning storage • No • Yes, when provisioning storage	 • Yes • No • No • No

Table 10. Available monitoring tasks in each GUI (continued)

Tasks	Available in web-based GUI	Available in stand-alone GUI
Managing virtual disks for storage systems[1]: • Viewing information about virtual disks • Creating virtual disks • Deleting virtual disks • Assigning volumes to host ports	• Yes • Yes, when provisioning storage • No • Yes, when provisioning storage	• Yes • No • No • No
Managing fabrics	Yes, but only for setting a zoning policy to enable automatic zoning. When automatic zoning is enabled, Tivoli Storage Productivity Center can create zones during storage provisioning.	Yes, but only for viewing zone configuration information
Troubleshooting performance issues	Yes	
Managing zones and zone sets	No	No
Creating storage policies	No	Yes
Charging back for storage usage	No	Yes
Profile management	Yes, managing service classes	No
Restriction: [1] Certain actions for provisioning storage and managing virtual disks are no longer available in Tivoli Storage Productivity Center. To complete actions that were not added to the web-based GUI, use the element manager or management GUI for the related resource.		

Alerting

Use the following list to determine the alerting tasks that you can complete in each GUI.

Table 11. Available alerting tasks in each GUI

Tasks	Available in web-based GUI	Available in stand-alone GUI
Creating alerts	No	Yes
Viewing alert conditions that are detected on monitored resources	Yes	Yes
Viewing replication alerts	Yes	Yes
Acknowledging the status of alerts	Yes	Yes
Removing alerts	Yes	Yes

Reporting

Use the following list to determine the reporting tasks that you can complete and information that you can view in each GUI.

Reporting tasks

Table 12. Available reporting tasks in each GUI

Tasks	Available in web-based GUI	Available in stand-alone GUI
Accessing the web-based reporting GUI	Yes	No
Saving reports	Yes	Yes
Scheduling reports	Yes	Yes

Reporting information

Table 13. Available reporting information in each GUI

Information	Available in web-based GUI	Available in stand-alone GUI
Resource properties	Yes	Yes
Performance	Yes	Yes
Capacity (including used and available space)	Yes	Yes
Historical capacity	Yes	Yes
Storage tiering	Yes	No
Policy violations	No	Yes
Rollup servers	No	Yes
File statistics	No	Yes
Data sources	No	Yes
Databases	No	Yes

Functions that were moved to the web-based GUI in V5.2

In Tivoli Storage Productivity Center V5.2, a number of product functions were moved from the stand-alone GUI to the web-based GUI. Use the following list to determine which function were moved.

Table 14. Functions that were moved from the stand-alone GUI to the web-based GUI

Tasks in V5.2	Moved to the web-based GUI
Configuring[1]	• Assigning roles to user groups • Upgrading Storage Resource agents
Discovering resources	• Hypervisors (VMware virtual infrastructure) • Storage systems • Switches • Fabrics In the web-based GUI, discovering resources is done automatically when you add resources in the Add resource wizard.

Table 14. Functions that were moved from the stand-alone GUI to the web-based GUI (continued)

Tasks in V5.2	Moved to the web-based GUI
Adding resources (data sources) for monitoring	• Servers[2] (agentless) • Hypervisors (VMware virtual infrastructure) • Storage systems • Switches • Fabrics • CIM agents for storage systems, switches, fabrics, and NAS filers
Removing resources (data sources)	• Servers • Hypervisors • Storage systems • Switches • Fabrics • NAS filers
Administering resources (data sources)	Yes: • Viewing and modifying connections • Testing connections • Updating credentials
Administering servers and Storage Resource agents	Yes: • Disabling agents • Enabling agents • Upgrading agents • Updating user credentials • Viewing logs • Collecting agent logs • Testing connections • Enabling and disabling fabric functions • Enabling and disabling running scripts on agent • Removing servers
Starting element managers (management GUIs) in context	You can open the start page of an element that is web-based or uses Java Web Start.
Starting Tivoli Storage Productivity Center for Replication in context[3]	Yes, from the Alerts page. If a replication alert includes a link to a storage system, you can click the link to startTivoli Storage Productivity Center for Replication.
Scheduling data collection	Yes, for the following types of jobs: • Probes • Performance monitors

Table 14. Functions that were moved from the stand-alone GUI to the web-based GUI (continued)

Tasks in V5.2	Moved to the web-based GUI
Viewing and managing jobs and tasks	Yes, for the following types of jobs and tasks: • Probes • Performance monitors • Agent upgrades • Provisioning (task) • Optimizing (task)
Troubleshooting performance	Yes
Viewing a visual topology of resources	Yes, but only for viewing data paths
Analytics • Provisioning storage • Optimizing storage • Cloud Configuration – Assigning storage to tiers – Creating and modifying service classes – Creating and modifying capacity pools	Yes
Viewing information about managed disk groups (pools)	Yes
Managing volumes and virtual disks for storage systems[4]: • Viewing information about volumes and virtual disks • Creating volumes and virtual disks • Assigning volumes to host ports	 • Yes • Yes, when provisioning storage • Yes, when provisioning storage
Managing fabrics	Setting a zoning policy to enable automatic zoning. When automatic zoning is enabled, Tivoli Storage Productivity Center can create zones during storage provisioning.

Restrictions:

- [1] The ability to migrate Data agents and Fabric agents is no longer available in either GUI. To monitor a resource where a Data agent or Fabric agent is installed, you must uninstall that agent and manually deploy a Storage Resource agent.
- [2] There are two types of servers that you can add in Tivoli Storage Productivity Center: agentless servers and servers with Storage Resource agents. Add agentless servers in the web-based GUI. Add servers with Storage Resource agents in the stand-alone GUI.
- [3] The ability to open pages in the Tivoli Storage Productivity Center for Replication GUI from the **Replication Management** node in the stand-alone GUI is no longer available. In 5.2 and later, you must start the Tivoli Storage Productivity Center for Replication GUI as a stand-alone application in a web browser. The exception is starting the GUI from the Alerts page as described in the table.
- [4] Certain actions for managing volumes and virtual disks are no longer available in Tivoli Storage Productivity Center (for example, deleting). To complete actions that were not added to the web-based GUI, use the storage management GUI for the related resource.

Starting Tivoli Storage Productivity Center

Tivoli Storage Productivity Center provides two graphical user interface (GUI) applications for managing and monitoring the resources in a storage environment: a web-based GUI and a stand-alone GUI. You can start these GUIs on the Tivoli Storage Productivity Center server or on a remote system.

Each GUI provides different functions for working with monitored resources.

Starting the Tivoli Storage Productivity Center stand-alone GUI

The stand-alone GUI contains functions for monitoring the condition of storage resources, and all the tools for managing data, disk, fabric, and replication.

How you start the stand-alone GUI depends on whether the stand-alone GUI component is installed on your system.

- On the Tivoli Storage Productivity Center server, or on a remote system where the stand-alone GUI component is installed, you start the stand-alone GUI by running a batch file or a shell script. For Windows operating systems, you run the batch file `tpc.bat`, which you can run from the **Start** menu. For AIX or Linux operating systems, you run the shell script `TPCD.sh`.
- If the stand-alone GUI component is not installed on a remote system, you use a Java Web Start application to download and start the stand-alone GUI.

Starting the Tivoli Storage Productivity Center stand-alone GUI by using Java Web Start

Start the stand-alone Tivoli Storage Productivity Center GUI by using Java Web Start to interact remotely with Tivoli Storage Productivity Center running on the server.

Before you start Tivoli Storage Productivity Center, ensure that you are using a supported web browser. For a list of web browsers that you can use with Tivoli Storage Productivity Center, see the support matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446. In the **Agents, Servers and Browser** column, click the version of Tivoli Storage Productivity Center that is installed on your system. On the next page, click **Web Browsers** to find the web browsers that you can use.

On remote systems where the stand-alone GUI component is not installed, you can start the stand-alone Tivoli Storage Productivity Center GUI using Java Web Start. To start the stand-alone GUI by using Java Web Start, you click a web link that downloads a Java Network Launching Protocol (JNLP) file from the server. The JNLP file specifies all the files needed by the application. The Java Web Start Launcher on the remote system opens the JNLP file, downloads and caches all the required files, and starts the stand-alone GUI in a new window.

1. Start a web browser, and type the following address for the Java Web Start page.

   `http://`*`hostname`*`:`*`port`*`/ITSRM/app/welcome.html`

   In the preceding address, specify the following values:

   `hostname`
   The Tivoli Storage Productivity Center server. You can specify the *hostname* as an IP address or a Domain Name System (DNS) name. To verify your host name, ask your Tivoli Storage Productivity Center administrator.

port
The port number for the Device server. The default port number for connecting to the Device server using the HTTP protocol is 9550. However, this port number might be different for your site. For example, the port number might be different if the default port range was not accepted during installation. If the default port number does not work, ask your Tivoli Storage Productivity Center administrator for the correct port number.

2. The Tivoli Storage Productivity Center GUI requires an IBM Java Runtime Environment 1.6.0. Links are provided on the Java Web start page for you to download the IBM Java Runtime Environment 1.6.0 for Windows, Linux, or AIX. If an IBM Java Runtime Environment 1.6.0 is not already installed on the system, click the link to download the IBM Java Runtime Environment 1.6.0 for your operating system. On Windows, download a self-extracting executable. On Linux, download an RPM file. On AIX, download an executable JAR file. Install the JRE from the file that you downloaded before you start the stand-alone GUI. If you are prompted to install the JRE as the system JRE or to overwrite the current system JRE, click **No**.
3. From the Java Web Start page, click **TPC GUI (Java Web Start)**

   The JNLP file is downloaded.

   Important: Depending on your browser and operating system and their default settings for your browser and operating system) you might need to:
   - Verify to your browser that the tpcgui.jnlp file is safe to download
   - Indicate that JNLP files are to be opened by the Java Web Start Launcher

   For more information, go to the Tivoli Storage Productivity Center information center at http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp and search for *Cannot start the Tivoli Storage Productivity Center stand-alone GUI remotely*.
4. The stand-alone Tivoli Storage Productivity Center GUI opens in a separate window. To log on to Tivoli Storage Productivity Center, enter the following information and click **OK**.
   a. Enter your user ID and password.

      Important: If you are logging on immediately after installation and, see "Overview of required user names for initial logon to the Tivoli Storage Productivity Center family of products" on page 125 for information on the user ID you must use.
   b. If the server field does not contain the address and port number of for the server, specify the following values for the *hostname* and *port*.

      hostname
      The Tivoli Storage Productivity Center server. You can specify the *hostname* as an IP address or a Domain Name System (DNS) name. To verify your host name, ask your Tivoli Storage Productivity Center administrator.

      port
      The port number for the Data server. The default port number for connecting to the Data server is 9549. However, this port number might be different for your site. For example, the port number might be different if the default port range was not accepted during installation.

Starting the Tivoli Storage Productivity Center stand-alone GUI on a system where the GUI component is installed

On the Tivoli Storage Productivity Center server, or on a remote system where the stand-alone GUI component is installed, you start the Tivoli Storage Productivity Center stand-alone GUI using the batch file `tpc.bat` on a Windows operating system or the shell script `TPCD.sh` on an AIX and Linux operating system.

A typical installation of Tivoli Storage Productivity Center on the server includes the stand-alone GUI component. A remote system can also have the stand-alone GUI component installed. On a system where the stand-alone GUI component is installed, use the following instructions for starting the stand-alone GUI. If you are on a remote system where the stand-alone GUI component is not installed, start the stand-alone GUI by using Java Web Start instead.

1. Start the Tivoli Storage Productivity Center GUI on Windows operating systems.
   - Open the Tivoli Storage Productivity Center stand-alone GUI. . If you are on a remote system, the **Start** menu or Start page for the stand-alone GUI might not be available. In this case, run the `tpc.bat` file directly. The default location for the batch file is `C:\Program Files\IBM\TPC\gui\tpc.bat`.
   - On AIX or Linux, type the following path and command at the command line.

     ```
     /opt/IBM/TPC/gui/TPCD.sh
     ```
2. From the Tivoli Storage Productivity Center GUI logon window, log on to Tivoli Storage Productivity Center.
   a. Enter your user ID and password.
   b. If the server field does not contain the address of the server, enter the information in the format *`hostname:port`*. Specify the following values for the *hostname* and *port*:

      `hostname`
      The Tivoli Storage Productivity Center server. You can specify the *hostname* as an IP address or a Domain Name System (DNS) name.

      `port`
      The port number for the Data server. The default port number for connecting to the Data server is 9549. However, this port number might be different for your site. For example, the port number might be different if the default port range was not accepted during installation.

   c. Click **OK**.

Starting the Tivoli Storage Productivity Center web-based GUI

You can start the web-based Tivoli Storage Productivity Center GUI by opening a web browser and entering a web address for the Tivoli Storage Productivity Center logon page. For example, you might enter `http://storage.example.com:9568/srm`.

Before you start IBM Tivoli Storage Productivity Center, ensure that you are using a supported web browser. For a list of web browsers that you can use with Tivoli Storage Productivity Center, see the support matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446. In the **Agents, Servers and Browser** column, click the version of Tivoli Storage Productivity Center that is installed on your system. On the next page, click **Web Browsers** to find the web browsers that you can use.

The web-based GUI provides quick access to pages that you can use to monitor the condition, capacity, and relationships of the resources within your storage environment. Start the web-based GUI if you are interested in monitoring your storage environment, but do not need the full set of administrative tools available in the stand-alone GUI.

1. On a server running the Windows operating system, start **TPC Web-based GUI**. If you are not on a server running the Windows operating system, start a web browser and enter the following address in the address field:

   `http://hostname:port/srm`

   In the preceding address, specify the following values:

   hostname
   The Tivoli Storage Productivity Center server. You can specify the host name as an IP address or a Domain Name System (DNS) name.

   port
   The port number for Tivoli Storage Productivity Center. The default port number for connecting to Tivoli Storage Productivity Center by using the HTTP protocol is 9568. However, this port number might be different for your site. For example, the port number might be different if the default port range was not accepted during installation. If the default port number does not work, ask your Tivoli Storage Productivity Center administrator for the correct port number.

   Tip: If you have a non-default port, check the value of the `WC_defaulthost` property in *TPC_installation_directory*`/web/conf/portdef.props` file.

   You might be redirected from the address that you enter to another address and port that provides secure access using the HTTPS protocol. This page is the Tivoli Storage Productivity Center logon page.
2. From the Tivoli Storage Productivity Center logon page, type your user name and password and click **Log in**. The web-based Tivoli Storage Productivity Center GUI opens in the browser.

Starting Tivoli Storage Productivity Center for Replication

You can start the IBM Tivoli Storage Productivity Center for Replication GUI by opening a web browser and entering a web address for the Tivoli Storage Productivity Center for Replication logon page.

Before you start Tivoli Storage Productivity Center for Replication, ensure that you are using a supported web browser. For a list of web browsers that you can use with Tivoli Storage Productivity Center, see the support matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446. In the **Agents, Servers and Browser** column, click the version of Tivoli Storage Productivity Center that is installed on your system. On the next page, click **Web Browsers** to find the web browsers that you can use.

The Tivoli Storage Productivity Center for Replication GUI provides a single point of control to configure, manage, and monitor copy services. Whether you start it on the server or on a remote system, the GUI is web-based and is displayed in a browser.

1. On a server that is running on the Windows operating system, click **Start** > **Programs** > **IBM Tivoli Storage Productivity Center** > **TPC Replication Manager GUI** to open a browser to the correct address. If you are not using a

server that is running on the Windows operating system, start a web browser and enter the following address in the address field. The address is case-sensitive.

`https://hostname:port/CSM`

In the preceding address, specify the following values:

hostname
The Tivoli Storage Productivity Center for Replication server. You can specify the host name as an IP address or a Domain Name System (DNS) name.

port
The port number for Tivoli Storage Productivity Center for Replication. The default port number for connecting to Tivoli Storage Productivity Center for Replication using the HTTPS protocol is 9559. However, this port number might be different for your site. For example:

- If you upgraded from Tivoli Storage Productivity Center for Replication version 3.3 or earlier, the default HTTPS port is 9443.
- If you upgraded from Tivoli Storage Productivity Center for Replication version 3.4 or later, the default HTTPS port is 3443.
- If you are using IBM WebSphere Application Server for z/OS, the default HTTPS port is 9443.
- If you are using IBM WebSphere Application Server OEM Edition for z/OS, the default HTTPS port is 32209.

You can find the port number for the Replication server in the *install_root*`/wlp/usr/servers/replicationServer/properties/portdef.props` file. The default port for the HTTP protocol is defined by the **`WC_defaulthost_secure property`** property in the file.

You can find the port number for the web-based GUI in the *install_root*`/eWAS/profiles/webServerProfile/properties/portdef.props` file. The default port for the HTTP protocol is defined by the `WC_defaulthost` property in the file.

2. From the Tivoli Storage Productivity Center for Replication logon page, type your user name and password and click **Login**. The web-based Tivoli Storage Productivity Center for Replication GUI opens in the browser.

Tivoli Storage Productivity Center for Replication

Tivoli Storage Productivity Center includes Tivoli Storage Productivity Center for Replication. Tivoli Storage Productivity Center for Replication manages copy services in storage environments. *Copy services* are features that are used by storage systems such as System Storage DS8000 to configure, manage, and monitor data-copy functions. Copy services include IBM FlashCopy, Metro Mirror, Global Mirror, and Metro Global Mirror.

Tivoli Storage Productivity Center for Replication provides copy services management for the storage systems and session types that are listed in the following table:

Storage system	Session type
System Storage DS8000	• Basic HyperSwap • FlashCopy • Metro Mirror Single Direction • Metro Mirror Failover/Failback • Metro Mirror Failover/Failback with Practice • Global Mirror Single Direction • Global Mirror Either Direction with Two-Site Practice • Global Mirror Failover/Failback • Global Mirror Failover/Failback with Practice • Metro Global Mirror • Metro Global Mirror with Practice
System Storage DS6000	• Basic HyperSwap • FlashCopy • Metro Mirror Single Direction • Metro Mirror Failover/Failback • Metro Mirror Failover/Failback with Practice • Global Mirror Single Direction • Global Mirror Either Direction with Two-Site Practice • Global Mirror Failover/Failback • Global Mirror Failover/Failback with Practice
TotalStorage Enterprise Storage Server® Model 800	• Basic HyperSwap • FlashCopy • Metro Mirror Single Direction • Metro Mirror Failover/Failback • Metro Mirror Failover/Failback with Practice • Global Mirror Single Direction • Global Mirror Either Direction with Two-Site Practice • Global Mirror Failover/Failback • Global Mirror Failover/Failback with Practice • Metro Global Mirror (only H1 site) • Metro Global Mirror with Practice (only H1 site)
SAN Volume Controller	• FlashCopy • Metro Mirror Single Direction • Metro Mirror Failover/Failback • Metro Mirror Failover/Failback with Practice • Global Mirror Single Direction • Global Mirror Failover/Failback • Global Mirror Failover/Failback with Practice • Global Mirror Failover/Failback with Change Volumes
Storwize V3500	• FlashCopy

Storage system	Session type
Storwize V3700	• FlashCopy • Metro Mirror Single Direction • Metro Mirror Failover/Failback • Metro Mirror Failover/Failback with Practice • Global Mirror Single Direction • Global Mirror Failover/Failback • Global Mirror Failover/Failback with Practice • Global Mirror Failover/Failback with Change Volumes
Storwize V7000	• FlashCopy • Metro Mirror Single Direction • Metro Mirror Failover/Failback • Metro Mirror Failover/Failback with Practice • Global Mirror Single Direction • Global Mirror Failover/Failback • Global Mirror Failover/Failback with Practice • Global Mirror Failover/Failback with Change Volumes
Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)	• FlashCopy • Metro Mirror Single Direction • Metro Mirror Failover/Failback • Metro Mirror Failover/Failback with Practice • Global Mirror Single Direction • Global Mirror Failover/Failback • Global Mirror Failover/Failback with Practice • Global Mirror Failover/Failback with Change Volumes
The XIV system	• Snapshot • Metro Mirror Failover/Failback • Global Mirror Failover/Failback

Key components of Tivoli Storage Productivity Center for Replication for System z

Tivoli Storage Productivity Center for Replication consists of several key components.

Tivoli Storage Productivity Center for Replication server

The Tivoli Storage Productivity Center for Replication server is the control point for the sessions, storage systems, hosts, and other components that are managed by Tivoli Storage Productivity Center for Replication.

Database

A single database instance serves as the repository for all Tivoli Storage Productivity Center for Replication data.

An embedded database is configured automatically with Tivoli Storage Productivity Center for Replication.

GUI The graphical user interface (GUI) is used to manage all Tivoli Storage Productivity Center for Replication functions.

CLI The command-line interface (CLI) is used to issue commands for key Tivoli Storage Productivity Center for Replication functions.

Interfaces for Tivoli Storage Productivity Center for Replication

Tivoli Storage Productivity Center for Replication provides a graphical user interface (GUI) and a command-line interface (CLI) for managing data replication and disaster recovery.

Graphical user interface

Tivoli Storage Productivity Center for Replication contains the following GUI components:

Navigation tree
The navigation tree provides categories of tasks that you can complete in Tivoli Storage Productivity Center for Replication. Clicking a task opens a main page in the content window.

Health Overview
The **Health Overview** area in the navigation tree shows a status summary for all sessions, storage systems, host systems, and management servers that Tivoli Storage Productivity Center for Replication is managing.

Content area
The content window displays content based on the item that you select in the navigation tree.

You can view help for a window by clicking the **?** icon.

You can view the information center by clicking the **i** icon. You must have Internet access to view the information center.

When you log on to the GUI, the **Health Overview** window is displayed in the content area.

Command-line interface

You can use the Tivoli Storage Productivity Center for Replication CLI by using the **csmcli** utilities. You can use the CLI directly or as a script to automate functions.

For security, the CLI runs only on the Tivoli Storage Productivity Center for Replication management server, which is the central point of control for managing data replication. You can run the CLI remotely by using a remote-access utility, for example, Secure Shell (SSH) or Telnet.

For Tivoli Storage Productivity Center for Replication on Windows systems, you can specify remote access to Linux or AIX terminals if you enabled Telnet on your Windows server.

Tivoli Storage Productivity Center for Replication integration with Tivoli System Automation Application Manager

IBM Tivoli System Automation Application Manager includes a Distributed Disaster Recovery (DDR) feature that uses IBM Tivoli Storage Productivity Center for Replication to provide a data replication and disaster recovery solution for

distributed systems. The DDR feature is available in Tivoli System Automation Application Manager version 3.2 or later and requires Tivoli Storage Productivity Center for Replication version 4.1.1.1 or later.

Tivoli System Automation Application Manager automates the availability of information technology resources. The DDR feature and Tivoli Storage Productivity Center for Replication provide a full end-to-end solution for site switching on distributed systems at the application level.

For more information about the integration of Tivoli Storage Productivity Center for Replication with Tivoli System Automation Application Manager, including the steps required to perform the integration, see the *Tivoli System Automation Application Manager Administrator's and User's Guide*. This guide is provided by version at http://www.ibm.com/developerworks/servicemanagement/dca/saam/resources.html.

Management server

The *management server* is a server that has Tivoli Storage Productivity Center for Replication installed. The management server provides a central point of control for managing data replication.

For more information about the management server, see "Management servers" on page 634.

Storage systems

A *storage system* is a hardware device that contains data storage. Tivoli Storage Productivity Center for Replication can control data replication within and between various storage systems.

To replicate data among storage systems using Tivoli Storage Productivity Center for Replication, you must manually add a connection to each storage system in the Tivoli Storage Productivity Center for Replication configuration.

For more information about the storage systems that you can use with Tivoli Storage Productivity Center for Replication, see "Storage systems" on page 643.

Host systems

A *host system* is an IBM z/OS or AIX system that connects to storage systems to enable certain replication features for those systems.

A connection to a host system is required in the following scenarios:

- A connection to a z/OS host system is required if you want to enable z/OS features such as HyperSwap and hardened freeze in Tivoli Storage Productivity Center for Replication sessions.
- A connection to an AIX host system is required if you want to use the Open HyperSwap feature. This feature enables the automatic swap of input/output (I/O) to the volumes on a secondary site when a failure occurs when I/O is written to the primary site. For the swap to occur, the storage systems that contain the volumes must be attached to an AIX host and the AIX host must be connected to Tivoli Storage Productivity Center for Replication. The Open HyperSwap feature is optional for the Metro Mirror with Failover/Failback session type.

You can add both AIX and z/OS host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed. For example, you can add a connection to a z/OS host in a Tivoli Storage Productivity Center for Replication instance that is running on Windows, Linux, or AIX.

For more information about the host systems and how to manage them in Tivoli Storage Productivity Center for Replication, see "Managing host systems" on page 659.

Ports

Tivoli Storage Productivity Center for Replication uses ports for communication with the management servers that are in a high-availability relationship; the graphical user interface (GUI); the command-line interface (CLI); and storage systems.

Tivoli Storage Productivity Center for Replication

Port numbers are defined during the installation of Tivoli Storage Productivity Center for Replication. To determine the port number for the ports that are used for Tivoli Storage Productivity Center for Replication, refer to the following table.

Table 15. Where to find Tivoli Storage Productivity Center for Replication port numbers

Tivoli Storage Productivity Center for Replication component	**Location of port information**
Web browser	To view the port number for the web browser, open the `portdef.props` file in the *`TPC_install_directory`*`\wlp\usr\servers\replicationServer\server.xml` directory.
Standby management server	To view the port number for each management server, click **Management Servers** in the navigation panel of the Tivoli Storage Productivity Center for Replication GUI. You can also run the **lshaservers** command from the command-line interface to view the ports for the management servers. **Important:** The port number for the standby management server must be the same on both the management server and the standby management server in a high-availability relationship. If you change the standby management server port number on either the management server or the standby management server, you must also change the port number on the other server.

Table 15. Where to find Tivoli Storage Productivity Center for Replication port numbers (continued)

Tivoli Storage Productivity Center for Replication component	Location of port information
Client and CLI port	To view the port number for the client, click **About** in the navigation panel of the Tivoli Storage Productivity Center for Replication GUI. You can also run the **whoami** command from the command-line interface to view the port for the client. **Important:** The client port number must be the same on both the management server and the standby management server in a high-availability relationship. If you change the client port number on either the management server or the standby management server, you must also change the port number on the other server.

Storage system ports

The following table lists the default ports for each storage system type.

Table 16. Storage system default ports

Storage System	Connection Type	Port
• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000	Direct Connection	2433
• System Storage DS8000	Hardware Management Console Connection	1750
• SAN Volume Controller • Storwize V3500 • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified	Direct Connection	443 and 22
• The XIV system	Direct Connection	7778

Ensure that your network configuration is set up so that Tivoli Storage Productivity Center for Replication can send outgoing TCP/IP packets to the storage controllers.

Because there are typically multiple applications that are running on the management server, it is possible that port conflicts might arise if other applications attempt to use the same ports that Tivoli Storage Productivity Center for Replication uses. Use the **netstat** command to verify which ports the various applications on the management server are using.

When you add a storage system to the Tivoli Storage Productivity Center for Replication configuration, the port field is automatically populated with the appropriate value. You can accept the default port or enter another port number.

If firewalls are being used in your configuration, ensure that none of these ports are being blocked. Also, ensure that not only is the Tivoli Storage Productivity Center for Replication server granted access to reach the other components, but that the other components are granted access to reach the Tivoli Storage Productivity Center for Replication server.

If the management server is running on Windows Server: Configure the inbound and outbound rules for Tivoli Storage Productivity Center for Replication. To create a new outbound rule, start the New Outbound Rule Wizard from the Windows Firewall with Advanced Security menu.

Sessions

A *session* manages a specific type of data replication against a specific set of volumes. The type of data replication that is associated with the session determines the replication actions that are available for the session. For example, the options for FlashCopy sessions are different from the options for Metro Mirror sessions.

For more information about sessions, see "Sessions" on page 665.

Copy sets

During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a *copy set*.

For more information about copy sets, see "Copy sets" on page 665.

Volume roles

A *volume role* is given to every volume in the copy set. The role defines how the volume is used in the copy set and, for multi-site sessions, the site location of the volume. For example, the H1 role is made up of host-attached volumes that are at the primary site.

For more information about volume roles, see "Volume roles" on page 669.

Role pairs

A *role pair* is the association of two volume roles in a session that take part in a copy relationship. For example, in a Metro Mirror session, the role pair can be the association between host volumes at the primary site and host volumes at the secondary site (H1-H2).

The flow of data in the role pair is indicated by an arrow. For example, H1>H2 denotes that H1 is the source and H2 is the target.

For more information about role pairs, see "Role pairs" on page 670.

Practice volumes

You can use a *practice volume* to practice for disaster recovery without interrupting current data replication. Practice volumes are available in Metro Mirror, Global Mirror, and Metro Global Mirror sessions.

You can practice disaster recovery actions without using practice volumes. However, without practice volumes, you cannot continue to copy data changes between volumes while you are practicing.

For more information about practice volumes, see "Practice volumes" on page 671.

Consistency groups

For Global Mirror and Metro Global Mirror sessions, Tivoli Storage Productivity Center for Replication manages the consistency of dependant writes by creating a consistent point-in-time copy across multiple volumes or storage systems. A *consistency group* is a set of target volumes in a session that were updated to preserve write order and are recoverable.

For more information about consistency groups, see "Consistency groups" on page 671.

Failover and failback operations

If a session has *failover* and *failback* capabilities, you can move the operation of Tivoli Storage Productivity Center for Replication from one site to another, and change the direction of the copy without completing a full copy.

Failover is the process of temporarily switching production to a backup facility by following a scheduled maintenance period or a disaster at a production site. A failover operation can be followed by a failback operation, which is the process of returning production to its original location and resuming the copy between the two sites.

For more information about failover and failback operations, see "Failover and failback operations" on page 671.

Session commands

The commands that are available for a session depend on the session type.

For descriptions of the session commands for each session type, see "Session commands" on page 706.

Session types

Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

The session types that are available depend on the storage system type and the edition of Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication that you are using as shown in this topic.

Basic HyperSwap

Basic HyperSwap is a copy services solution for IBM z/OS version 1.9 and later. Basic HyperSwap provides high availability of data if a storage system failure occurs.

The Basic HyperSwap session type enables Tivoli Storage Productivity Center for Replication to manage Basic HyperSwap by using Metro Mirror or Metro Global Mirror replication.

Basic HyperSwap sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 17. Basic HyperSwap

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication Basic Edition for System z • Tivoli Storage Productivity Center for Replication for System z	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000

FlashCopy

FlashCopy replication creates a point-in-time copy of the data on a source volume to a target volume. Data that existed on the target volume is replaced by the copied data.

FlashCopy sessions are available for the following Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 18. FlashCopy

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication Basic Edition for System z • Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3500 • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Snapshot

Snapshot replication creates a point-in-time copy of a volume or set of volumes without having to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

Snapshot sessions are available for the following Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 19. Snapshot

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication Basic Edition for System z • Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	The XIV system

Metro Mirror Single Direction

Metro Mirror is synchronous replication that operates between two sites that are up to 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system.

Metro Mirror replication maintains identical data in both the source and target volumes. When a write operation is issued to a source volume, the changes are propagated to the target volume before the write operation finishes processing.

With the Metro Mirror Single Direction session type, Metro Mirror replication is available only from the primary site.

Metro Mirror Single Direction sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 20. Metro Mirror Single Direction

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Metro Mirror Failover/Failback

Metro Mirror Failover/Failback replication enables you to switch the direction of the data flow so that you can use your secondary site as your production site. You can then copy changes that are made at the secondary site back to the primary site.

Metro Mirror Failover/Failback sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 21. Metro Mirror Failover/Failback

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage) • The XIV system

Metro Mirror Failover/Failback with Practice

Metro Mirror Failover/Failback with Practice replication combines Metro Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Metro Mirror Failover/Failback with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 22. Metro Mirror Failover/Failback with Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Global Mirror Single Direction

Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system.

Global Mirror replication maintains identical data in both the source and target volumes. When a write operation is issued to a source volume, the changes are typically propagated to the target volume a few seconds after the data is written to the source volume.

With the Global Mirror Single Direction session type, Global Mirror replication is available only from the primary site.

Global Mirror Single Direction sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 23. Global Mirror Single Direction

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Global Mirror Either Direction with Two-Site Practice

Global Mirror Either Direction with Two-Site Practice replication enables you to run Global Mirror replication from either the primary or secondary site. This session type provides volumes on the primary and secondary site that you can use to practice for disaster recovery without losing your disaster recovery capability.

Global Mirror Either Direction with Two-Site Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 24. Global Mirror Either Direction with Two-Site Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000

Global Mirror Failover/Failback

Global Mirror Failover/Failback replication enables you to switch the direction of the data flow so that you can use your secondary site as your production site. You can then copy changes that are made at the secondary site back to the primary site.

Global Mirror Failover/Failback sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 25. Global Mirror Failover/Failback

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage) • The XIV system

Global Mirror Failover/Failback with Change Volumes

Global Mirror Failover/Failback with Change Volumes replication provides the same capabilities as the Global Mirror Failover/Failback session type. However, Global Mirror Failover/Failback with Change Volumes sessions also provide the option of enabling or disabling the use of change volumes. Change volumes are available in SAN Volume Controller and Storwize storage systems.

Global Mirror Failover/Failback with Change Volumes sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 26. Global Mirror Failover/Failback with Change Volumes

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• SAN Volume Controller Version 6.3 and later • Storwize V3700 Version 6.4.1 and later • Storwize V7000 Version 6.3 and later • Storwize V7000 Unified Version 1.3 and later (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Global Mirror Failover/Failback with Practice

Global Mirror Failover/Failback with Practice replication combines Global Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Global Mirror Failover/Failback with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 27. Global Mirror Failover/Failback with Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Metro Global Mirror

Metro Global Mirror replication provides continuous, remote data replication that operates between three sites that are varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source. Metro Global Mirror replication enables you to switch the direction of the data flow so that you can use your secondary or tertiary site as your production site.

Metro Global Mirror sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 28. Metro Global Mirror

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 (only H1 site) • System Storage DS8000

Metro Global Mirror with Practice

Metro Global Mirror with Practice replication combines Metro Mirror, Global Mirror, and FlashCopy capabilities to provide a point-in-time copy of the data on the tertiary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Metro Global Mirror with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 29. Metro Global Mirror with Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 (only H1 site) • System Storage DS8000

Related concepts:

"Basic HyperSwap" on page 680
Basic HyperSwap is a copy services solution for z/OS version 1.9 and later. It provides high availability of data if a disk storage system failure occurs. Basic HyperSwap does not replace a disaster recovery solution. If a session is suspended but the suspend operation was not caused by a HyperSwap trigger, the session is not frozen to ensure that it is consistent.

"FlashCopy" on page 681
FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established.

"Snapshot" on page 682
Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes. You do not have to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

"Metro Mirror" on page 683
Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is in one storage system and the target is in another storage system.

"Global Mirror" on page 689
Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system. Global Mirror replication maintains identical data in both the source and target volumes.

"Metro Global Mirror" on page 694
Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that are varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source.

Metro Mirror heartbeat

The Metro Mirror *heartbeat* guarantees data consistency across multiple storage systems when the IBM Tivoli Storage Productivity Center for Replication management server cannot communicate with one or more storage systems.

The Metro Mirror heartbeat is disabled by default. When determining whether to use the Metro Mirror heartbeat, analyze your business needs. If you disable the Metro Mirror heartbeat, data might become inconsistent. If you enable the Metro Mirror heartbeat and a freeze occurs on the storage system, your applications cannot perform write operations during the freeze.

For more information about the Metro Mirror heartbeat, see "Metro Mirror heartbeat" on page 736.

Site awareness

You can associate a location with each storage system and each site in a session. This *site awareness* ensures that only the volumes whose location matches the location of the site are allowed for selection when you add copy sets to the session. Site awareness prevents a session relationship from being established in the wrong direction.

For more information about site awareness, see "Site awareness" on page 649.

Users and groups

For authentication and authorization, Tivoli Storage Productivity Center for Replication uses users and groups.

For information about users and groups for Tivoli Storage Productivity Center for Replication for System z, see the IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide.

For information about users and groups for Tivoli Storage Productivity Center for Replication, see "Users and groups" on page 782.

User roles

A *user role* is a set of privileges that is assigned to a user or user group to allow the user or user group to perform certain tasks and manage certain sessions.

To be assigned to a role, each user or group of users must have a valid user ID or group ID in the user registry on the management server.

Both individual users and a group of users can be assigned to a role. All users in a group are assigned the role of the group. If a user is assigned to one role as an individual and a different role as a member of a group, the user has access to the permissions of the role that has greater access.

IBM Tivoli Storage Productivity Center for Replication provides the following predefined user roles: monitor, session operator, and administrator.

For more information users roles, see "User roles" on page 783.

Key concepts

This section contains a technical overview that will help you understand how IBM Tivoli Storage Productivity Center works. An understanding of the concepts in this section will help you use Tivoli Storage Productivity Center effectively.

Data collection

To help you implement a storage management strategy, it is critical that you determine the information that you want to gather about the resources within your environment. You can schedule different data collection jobs depending on the types of information that you want to gather and the resources that you want to monitor.

Use probes, performance monitors, scans, and pings to collect detailed information about resources. You must schedule and run these data collection jobs before you can complete other tasks within Tivoli Storage Productivity Center, such as viewing reports, enforcing storage policies, setting alerts, viewing data paths, and managing storage systems.

You schedule probe and performance monitor jobs for a resource when you add the resource for monitoring in the web-based GUI. After you add the resource, you can modify the probe and performance monitor job schedules on the list and details pages for the resource. For example, to modify the probe job schedule for a storage system, go to the Storage Systems page or the Storage System details page in the web-based GUI.

Exception: For servers with Storage Resource agents, and IBM SONAS storage systems, you must first add the resource for monitoring by using the stand-alone GUI. After you add the resource, you use the web-based GUI to schedule probe and performance monitor jobs.

Use the web-based GUI to schedule data collection jobs to gather the following types of information:

Asset and status information
Use probes to collect asset, storage statistics, and status information about resources. You can run probes on all the resources that are monitored by Tivoli Storage Productivity Center.

Performance monitoring
Use performance monitors to collect metrics that measure the performance of switches and storage systems.

Use the stand-alone GUI to schedule data collection jobs to gather the following types of information:

Availability monitoring
Use pings to collect information about the availability of resources. You can run pings against servers and computers only. Pings are unavailable for storage systems and other resources.

Usage and trending information
Use scans to collect statistics about the usage and trends of your actual storage consumption. Scans provide all the data for usage reporting, quota analysis, and constraint analysis. They gather information at the file and directory level. You can run scans against servers and computers only. Scans are unavailable for storage systems, switches, fabrics, and other devices.

Tivoli Storage Productivity Center server monitoring
Use TPC Server probes to collect the storage asset and capacity information that was gathered by other Tivoli Storage Productivity Center servers in your environment. Create master and subordinate relationships among servers to use a single interface to generate reports that are based on data and metrics that are collected by multiple servers.

Performance monitoring and troubleshooting

Tivoli Storage Productivity Center can collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations that can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage.

To monitor the performance of resources and check for threshold violations, complete the following tasks:

- Add resources for monitoring and schedule data collection
- Define performance thresholds
- View and troubleshoot performance issues

Collect performance data

Before you can troubleshoot and view reports about performance, you must collect data about monitored resources. Performance monitors are data collection jobs that gather performance information about resources. This information includes metrics

that measure the performance of the components within a resource. Metrics measure the performance characteristics of volumes, ports, and disks on storage systems and switches. Tivoli Storage Productivity Center provides many different metrics for measuring performance. For example, some key metrics for storage systems are I/O rate in I/O operations per second, data rate in MiB per second, and response time in milliseconds.

You can use metrics in Tivoli Storage Productivity Center to track growth or change in I/O rates, data rates, and response times. In many environments, I/O and data rates grow over time, and response times increase as those rates increase. This relationship can help with "capacity planning" for your storage. As rates and response times increase, you can use these trends to project when more storage performance and capacity is required.

Define performance thresholds

Performance thresholds alert you when the performance of a monitored resource falls outside of a specified range and might represent a potential problem. When you define a performance threshold for an internal component of a resource, select a specific metric that you want to measure and its boundary values. When the performance of a resource falls outside the boundary values, a threshold violation occurs. For example, you can define a performance threshold that alerts you when the overall back-end response time for a managed disk on a SAN Volume Controller exceeds a certain value. The overall back-end response time is a metric that measures the average number of milliseconds that it takes to service each I/O operation on a managed disk.

Restriction: You only set performance thresholds for certain metrics. For a complete list of performance thresholds, see “Performance thresholds” on page 1690.

View and troubleshoot performance issues

After data collection and performance thresholds are configured, you can use the web-based GUI to complete the following tasks:

- Measure, compare, and troubleshoot the performance of switches, storage systems, and their internal resources.
- Review the threshold violations and alerts that were triggered when the performance of a resource fell outside of a specific range.
- View performance information in a chart or table format to help you quickly identify where and when performance issues are occurring. The chart is a visual representation of how the performance of resources trend over time.
- Customize views of performance so that you can analyze specific resources and metrics during time ranges that you specify.
- Drill down into resources to view detailed information about the performance of internal and related resources. For example, if a SAN Volume Controller storage system is shown in the chart, you can quickly view and compare the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.
- Implement server-centric monitoring of SAN resources without requiring a Storage Resource agent. When you add an agentless server for monitoring, Tivoli Storage Productivity Center automatically correlates that server with the ports on known host connections. If matches are found between the server and host connections on monitored storage systems, you can view the performance of the

internal resources that are directly associated with the SAN storage that is assigned to the server. For example, if a SAN Volume Controller maps two volumes to the server, you can view the performance of those volumes and the related managed disks.

- Export performance information to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application.
- In the reporting function of the web-based GUI, you can also view and create performance reports about multiple resources.

Related concepts:

"Monitoring the performance of resources" on page 383
Tivoli Storage Productivity Center can collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations. Metrics can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage. Threshold violations alert you when the performance of a monitored resource falls outside of a specified range.

Chapter 7, "Reporting," on page 789
Tivoli Storage Productivity Center provides multiple user interfaces for viewing reports about the storage infrastructure in an enterprise environment.

Related tasks:

"Adding resources" on page 257
You must add resources to Tivoli Storage Productivity Center before you can collect data, generate reports, and manage storage that is related to those resources.

Related reference:

"Performance metrics" on page 1701
Tivoli Storage Productivity Center can collect information about the performance of monitored storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks.

"Performance thresholds" on page 1690
Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

Storage optimization

Storage optimization tasks optimize the resources in your storage environment. You can use the web-based GUI to create tasks to balance pools, re-tier volumes, or transform volumes that are in storage virtualizer pools.

Balancing the workload of volumes across storage pools

Use the Balance Pools Analysis wizard to balance the workload of volumes across pools on the same tier. The pools are analyzed and recommendations are generated to move volumes from pools with high utilization values to pools with low utilization values.

Re-tiering volumes

Use the Analyze Tiering wizard to tier volumes automatically or based on the criteria that you set in tiering policies. For example, you can tier volumes that are based on the volume workload or on file usage, or both. Depending on the conditions that are set in the tiering policy, recommendations are generated. For example, you can reduce storage costs by moving volumes with low workloads to

lower or less expensive tiers. You can also improve performance and use storage more efficiently by moving volumes with heavy workloads to the tiers that best meet their workload requirements.

Transforming volumes

Use the Transform Storage wizard to complete the following tasks for one or more volumes in storage virtualizer pools:

- Move volumes from one storage virtualizer pool to another pool in the same storage virtualizer pool.
- Move volumes in a storage virtualizer pool to a pool that is enabled for Easy Tier.
- Convert fully allocated volumes to thin-provisioned volumes and convert thin-provisioned volumes to fully allocated volumes.
- Convert fully allocated volumes to compressed volumes and convert compressed volumes to fully allocated volumes.

The pools are analyzed and recommendations are generated. For example, analysis occurs to ensure that there is sufficient space in the pool to convert volumes or add volumes.

Related tasks:

"Optimizing storage tiering" on page 428
To optimize the placement of volumes on storage tiers, use the Analyze Tiering wizard.

"Balancing pools" on page 436
To distribute the workload of volumes across pools on the same tier, use the Balance Pools wizard. To improve the performance of the pools, recommendations are generated to move volumes to other pools on the same tier.

"Transforming and migrating volumes" on page 437
You can transform fully allocated volumes to compressed or thin-provisioned volumes, or transform compressed or thin-provisioned volumes to fully allocated volumes. You can move volumes to other pools or to pools that are enabled for Easy Tier.

"Managing tasks for tiering storage, balancing pools, and transforming storage" on page 452
To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the Tasks page in the web-based GUI to manage all the tasks that are used by Tivoli Storage Productivity Center to optimize storage. Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.

Cloud configuration

To take advantage of the simplified provisioning and optimization capabilities that are available by using the web-based GUI, some configuration is required.

To provision storage, you are required to specify only the storage capacity and storage quality that is required. After volumes are created, Tivoli Storage Productivity Center can analyze and optimize volume performance. To take advantage of these capabilities, create service classes and, optionally, capacity pools.

A tutorial of cloud configuration concepts is available in the web-based GUI. From the navigation pane, select **Advanced Analytics** > **Cloud Configuration** > **Learn the Concepts**.

Service classes

A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.

Service classes simplify provisioning requests by representing a level or type of storage quality. When you are requesting storage, only the required capacity and service class must be specified. Before you can provision storage, you must create service classes that describe the capabilities and characteristics of the storage you want to be able to provision. Service classes can later be modified or deleted as the needs of your installation change.

Block-storage service classes

A block-storage service class describes attributes and capabilities of block storage resources. When you provision volumes, you specify the requirements by using a block-storage service class. Based on the requirements of the service class, IBM Tivoli Storage Productivity Center identifies a storage pool for the volume.

File-storage service classes

A file-storage service class describes attributes of file storage resources. The file storage resources include Network Attached Storage (NAS) filers, Network Shared Disks (NSDs), and file systems. When you provision shares, you specify the requirements by using a file-storage service class. Based on the requirements of the service class, Tivoli Storage Productivity Center identifies a file system or IBM SONAS NSD for the share. Space on an IBM SONAS NSD is allocated if a new file system is required or an existing file system requires more space.

Service class properties

A service class represents storage resources with common traits. These common traits are the service class properties, and collectively describe a particular quality of service. Before you create service classes, consider the levels of storage quality that you want to have available for provisioning requests. For block storage, for example, you can organize service classes by storage tier or RAID level. Some service class properties are used during provisioning to identify the best location for storage placement. Other service class properties determine how the storage is configured, and reflect the performance, reliability, or security considerations for the level of service.

In addition to the standard properties of a service class, you can create your own properties by using custom tags. Custom tags can represent any common trait that storage systems must have. For more information about creating your own properties by using custom tags, see “Custom tags” on page 68.

Block-storage service classes can specify the following properties:

- Storage tier or range of storage tiers
- RAID level

- Whether volume virtualization is required or not allowed. This property is available only when at least one IBM storage virtualizer is managed by Tivoli Storage Productivity Center. The IBM storage virtualizers are SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified.
- Whether volumes must be thin provisioned or must not be thin provisioned.

  When a thin-provisioned volume is created in a DS8000, an extent space efficient (ESE) volume is created.

  If thin provisioning is required, you can specify thin-provisioning configuration properties for SAN Volume Controller, IBM Storwize V7000, IBM Storwize V7000 Unified, and IBM XIV Storage System.

  If the volume is created on a SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified, the following configuration properties affect the volume that is created:

  – Allocated space
  – Auto-expand
  – Warning level
  – Grain size

  If the volume is created on an XIV system, the locking behavior configuration property affects the volume that is created.
- A multipathing policy for the IBM System Storage Multipath Subsystem Device Driver (SDD) on the client host. If the client host has a storage resource agent (SRA) and uses the SDD, the multipathing policy property is used to configure the driver on the client host. When a multipathing policy is set, and fabrics are managed by Tivoli Storage Productivity Center, the following properties affect how the storage system is connected to the client host:
  – Whether the storage system and the client host are connected through fully redundant paths
  – Number of paths between the storage system and the client host
- If at least one DS8000 Encryption Group is discovered by Tivoli Storage Productivity Center, the following properties are available and affect the storage placement and configuration of volumes:
  – Whether volumes must be encrypted
  – If volumes must be encrypted, the DS8000 encryption group to use

File-storage service classes have the following properties:

- Whether dedicated storage is required, or whether shared storage is allowed. In other words, whether the file system from which the NAS share is provisioned can contain other NAS shares.
- File set type. The file set type is either an independent file set with its own allocated nodes, or a dependent file set allocated on the file system.
- Access path host name replacement. For more information, see “Special consideration for provisioning by using the VMware vSphere Web Client extension” on page 69.

Predefined service classes

To illustrate service classes, a set of predefined service classes are provided. From the Service Classes page in the web-based GUI, you can view any of these service classes for more information. Select a service class and then select **Actions** > **View / Modify**. To display help information about a service class property, complete the following steps:

1. Move the mouse pointer over the field or control.
   A question mark icon is displayed next to the field.
2. Move the mouse pointer over the question mark icon to display an explanation of the property.

The following sample block-storage service classes are provided for illustration.

Gold The Gold service class represents the highest-performing storage resources for mission-critical applications.

Silver The Silver service class represents high-performing storage resources for applications in production.

Bronze
The Bronze service class represents standard storage resources for non-mission-critical applications.

The following sample file-storage service classes are provided for illustration.

NormalIsolation
The NormalIsolation service class describes normal isolation file storage by specifying that shared storage is allowed. The file system from which the NAS share is provisioned can contain other NAS shares.

EnhancedIsolation
The EnhancedIsolation service class describes enhanced isolation file storage by specifying that dedicated storage is required. The file system from which the NAS share is provisioned cannot contain other NAS shares.

Custom tags

During provisioning, Tivoli Storage Productivity Center determines candidates for storage placement by comparing the requirements of a service class with the known capabilities, configuration, and performance of the available storage systems. You can think of the storage placement determination as a filtering process. Tivoli Storage Productivity Center filters the set of available storage resources against properties of the service class until only the storage resources that can satisfy all of the requirements remain.

In addition to the standard properties of a service class, you can create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.

When you specify tags on a block-storage service class, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

When you specify tags on a file-storage service class, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

During provisioning, Tivoli Storage Productivity Center filters candidates for provisioning against the custom tags.

Before you create service classes, consider whether there are any special requirements for provisioning that are not addressed by the standard properties of

the service class. If so, tag the appropriate storage resources to satisfy custom requirements of a service class. When you are defining the service class, specify the custom tags that are required.

Adding users to a service class

By default, only administrators can provision storage. However, an administrator can add users to a service class to grant them permission to provision by using the service class. In order for an administrator to grant permission to provision to a user, the user must be assigned to the Monitor or External Application role.

When you are granting users permission to provision by using a particular service class, you can also specify whether execution of provisioning plans created by these users requires administrator approval. The requirement for administrator approval applies to all non-administrative users who are added to the service class, and also applies to external applications such as the vSphere Web Client extension. Approval cannot be required for some non-administrative users and not for others.

Before you create service classes, consider whether you want to allow users to request their own storage. Consider whether you want user provisioning requests to require administrator approval.

Associating capacity pools with a service class

Before you create service classes, consider whether certain storage requests must always be satisfied from a particular set of resources. If so, you can associate one or more capacity pools with a service class. For example, suppose that a set of storage resources are allocated to a particular department of your business. All storage requests for that department must be satisfied by those storage resources. In this case, you can add all the storage resources to a capacity pool and associate that capacity pool with a service class. You can also add the members of the department to the service class. The department members can then request their own storage as needed. However, the requests can be satisfied only by the storage resources that are allocated to the department.

When you are configuring capacity pools and service classes, keep in mind that a storage resource can be a member of only one capacity pool. Adding storage resources to a capacity pool is not a viable approach if the resources are already members of another capacity pool. Consider specifying custom tags to achieve the same result. In the preceding example, suppose that some or all of the storage resources that are allocated to the department are already members of capacity pools. In this case, you can instead tag the storage systems and the service class with matching tags to achieve the same result. With either approach, only the storage systems that are allocated to the department are candidates for provisioning when storage is requested by using the service class.

Special consideration for provisioning by using the VMware vSphere Web Client extension

When a file share is provisioned by the VMware vSphere Web Client extension, an access path to the share is provided by Tivoli Storage Productivity Center. By default, the path includes the cluster name of the file storage system in which the share is created. A cluster name must be a valid host name that is registered with a Domain Name System (DNS) server. If the cluster name is not valid, the share is not accessible under the access path that is returned by Tivoli Storage Productivity Center.

To address this situation, you can set one of its custom tags for a file storage system to specify an access path host name replacement. A file-storage service class can specify which of the file storage system custom tags contains the access path host name replacement.

If file shares are provisioned by using the VMware vSphere Web Client extension, determine whether you must set access path host name replacements.

Capacity pools

Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. Configure capacity pools to track the used and available space for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

The following types of storage resources can be grouped into capacity pools:

- Storage systems
- Storage pools
- File systems of file storage systems
- IBM SONAS Network Shared Disks (NSDs)

You can organize storage resources into capacity pools in any way that serves your business needs. For example, you can separate the storage resources that are allocated for separate divisions of your business into separate capacity pools. You can then track the storage use for each division separately, and restrict provisioning requests to the appropriate set of storage resources.

Tracking used and available space in a capacity pool

You can track the used and available space for any set of storage resources by adding the storage resources to a capacity pool. From the Capacity Pools pane of the web-based GUI, you can view the total capacity of the resources, and monitor the following capacity measurements:

- Used file space
- Available file space
- Used block space
- Available block space

Restricting a provisioning request to a capacity pool

You can use capacity pools to define a set of storage resources from which provisioning requests must be satisfied. When you are requesting storage by using the Provision Storage wizard or the vSphere Web Client extension, you can specify a capacity pool. If you do, the provisioning request is restricted to resources in the capacity pool. Only those resources are candidates for provisioning.

You can also associate a service class with capacity pools. If a service class is associated with capacity pools, provisioning requests for the service class must specify, and be constrained to, one of the associated capacity pools. Because you can also grant non-administrative users permission to provision storage by using the service class, associating capacity pools with the service class restricts the users to a specific set of resources.

Provisioning storage

The Provision Storage wizard simplifies the steps for submitting provisioning requests. When you use the Provision Storage wizard, you concern yourself only with which service class is needed for the new volumes or shares, and how much capacity you require.

Technology Preview: The provisioning feature is part of a Technology Preview program. Use this preview to evaluate the functions and usefulness of the provisioning feature as it is being developed and enhanced for future releases. To report issues, send feedback, and submit enhancement requests, join the conversation on Service Management Connect.

Provisioning volumes

Restriction: To provision volumes, you must have the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license.

To provision volumes, you select one or more servers or hypervisors and start the Provision Storage wizard. You specify the capacity that is required for one or more volumes, and a service class. IBM Tivoli Storage Productivity Center identifies the storage pools that can provide the capacity and the service class. From the set of pools that can provide the capacity and service class, Tivoli Storage Productivity Center identifies the best location for the storage. The best location for the storage is based on the unallocated volume space in the pool and performance data. Preference is given to systems that have available performance data. When you request multiple volumes, they are provisioned from pools in the same storage system.

Provisioning shares

To provision a share, you select one or more servers or hypervisors and start the Provision Storage wizard. You specify the capacity that is required for the share, a service class, and information about how to export the share. Tivoli Storage Productivity Center identifies the file systems and Network Shared Disks (NSDs) that can provide the capacity and the service class. From the set of file systems and NSDs that can provide the capacity and service class, Tivoli Storage Productivity Center identifies the best location for the storage. The best location for the storage is based on the available space on the file system or NSD.

The provisioning task

When you complete the steps of the Provision Storage wizard, Tivoli Storage Productivity Center displays a provisioning task. If you have Administrator privileges, you can save, run, or schedule the task.

In the service class, an administrator can grant users permission to provision by using the service class. In the service class, an administrator can also specify whether scheduling or running provisioning tasks that are created by using the service class requires administrator approval. If you are a user with permission to provision by using the service class, and administrator approval is not required, you can save, run, or schedule the task. If administrator approval is required, you can only save the task. If you save the task, an administrator can later schedule or run it.

Although Tivoli Storage Productivity Center identifies the best location for storage when it creates the provisioning task, the implementation of the task might fail. For example, the implementation of the task might fail because of changes in the environment that occur after the task is created and before it is run. The implementation of the task might also fail because of storage system restrictions that are not considered when the task is created.

Viewing task details and logs

When a provisioning task is running, you can view its status in a details page. If you run the provisioning task immediately after you complete the steps of the Provision Storage wizard, the details page is already displayed. You can also display details of a provisioning task from the Tasks page. From the details page, you can open logs for the provisioning task. The logs show the steps that are taken by the task during processing and include detailed information about any warnings or errors.

Tracking

In the Provision Storage wizard, you can specify a ticket identifier for tracking purposes. The ticker identifier is associated with the provisioning task, and with any volume or share that is created by the provisioning task. The ticket identifier can be viewed by showing the Ticket column in the Volumes page or the Shares page. The ticket identifier can also be viewed in the properties notebook for a volume or share.

Storage requirements

In the Provision Storage wizard, you specify your storage requirements by specifying a service class. Certain properties of the service class describe capabilities that storage resources must have in order to provide the service class. For example, when you provision volumes, a block-storage service class specifies properties such as a storage tier, a RAID level, and whether the pool must be able to encrypt or thin provision volumes.

A service class typically represents a quality or service. For example, Tivoli Storage Productivity Center provides predefined block-storage service classes named Gold, Silver, and Bronze. Gold represents the highest-performing storage resources, and Silver and Bronze represent lesser levels of storage quality.

The different levels of storage quality that are provided by each service class are defined by the service class properties. For example, each service class specifies a different required storage tier that depends on the level of service that is required. Because the Gold service class is for mission critical applications, it sets thin provisioning off.

When you use the Provision Storage wizard, you concern yourself only with which service class is required. If you are unfamiliar with the service classes, however, you can open a separate window to view service class information. If you have Administrator privileges, you can modify or create service classes.

Storage constraints

Tivoli Storage Productivity Center identifies the storage resources that can provide the capacity and the service class from a set of storage resources. This set might be all the storage resources that are known to Tivoli Storage Productivity Center, or it

might be constrained to a subset of those resources. In particular, your site can create capacity pools. *Capacity pools* are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. For example, a capacity pool might contain the storage resources that are allocated to a particular department or division of your business.

A service class might restrict storage placement to one or more capacity pools. If the service class you are provisioning from restricts placement, you must select one of the allowed capacity pools. If the service class does not restrict storage placement, you can specify that candidates for provisioning can be selected from all storage resources that are known to Tivoli Storage Productivity Center, all storage resources in any capacity pool, or storage resources in a particular capacity pool. If you are unfamiliar with the capacity pools that you can select when you are using the Provision Storage wizard, you can open a separate window to view capacity pool information. If you have Administrator privileges, you can modify or create capacity pools.

If the fabrics are managed by Tivoli Storage Productivity Center, the set of storage resources might be further constrained by your zoning policy. If the zoning policy is not configured to automatic zoning, only storage systems with connectivity to the server or hypervisor are candidates for storage placement.

Storage configuration

Whether you are provisioning volumes or shares, Tivoli Storage Productivity Center configures the storage and resources as specified by certain properties of the service class, and, in the case of volumes, according to your zoning policy.

When you provision volumes, Tivoli Storage Productivity Center configures volumes and resources according to properties of the service class:

- If the service class requires thin provisioning, the volume is created in a pool that can thin provision volumes. During provisioning, other properties of the service class are used to configure the thin provisioned volume for the particular type of storage system that contains the pool. Because you configure thin provisioned volumes differently for the different storage system types, the service class has a unique set of thin provisioning properties for each type.
- If the service class specifies a multipathing policy, the IBM System Storage Multipath Subsystem Device Driver (SDD) on the client host is configured to use that policy. For the SDD to be configured, a Storage Resource agent (SRA) must be running on the client host.

When you are provisioning volumes, if automatic zoning is enabled, new zones might be created to connect a server to the storage system. Existing zones are used if the server already has connectivity to the storage system. Otherwise, one or more new zones are created between a host initiator port and a controller, node, or module port.

When you provision shares, Tivoli Storage Productivity Center configures at least one file access protocol for the file share.

Tivoli Storage Productivity Center VASA provider

Tivoli Storage Productivity Center VASA provider enables the Tivoli Storage Productivity Center administrator to bring visibility of Tivoli Storage Productivity Center managed storage resources to a vCenter server administrator.

About vSphere API for Storage Awareness APIs Providers (VASA)

VASA enables coordination between VMware vCenter and storage arrays. It is designed to provide built-in storage insight to support VM storage provisioning, assist with storage troubleshooting and enable new DRS-related use cases for storage.

VASA benefits

For VMware vSphere users, the Tivoli Storage Productivity Center VASA provider improves the ability to monitor and automate storage-related operations in VMware environments.

A Tivoli Storage Productivity Center storage administrator can share certain Tivoli Storage Productivity Center alerts and storage capabilities for resources that are mapped to a vCenter server.

The Tivoli Storage Productivity Center VASA provider provides the underlying connectivity for alerts and information to be shared.

Tivoli Storage Productivity Center VASA provider architecture

This figure shows the Tivoli Storage Productivity Center VASA provider running on a Tivoli Storage Productivity Center connected to both the VMware vCenter server and the managed storage resources.

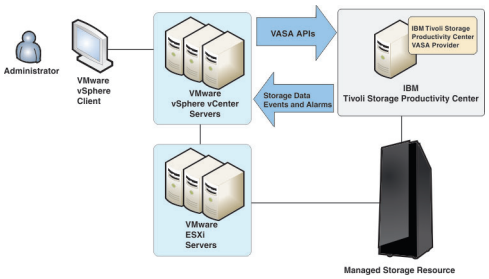

A Tivoli Storage Productivity Center storage administrator can share certain Tivoli Storage Productivity Center alerts for resources that are mapped to a vCenter server.

The Tivoli Storage Productivity Center VASA provider provides the underlying connectivity, and is automatically deployed and running after a Tivoli Storage Productivity Center installation. You must register a Tivoli Storage Productivity Center server to view its information in vCenter storage reports and views.

vSphere Web Client extension for Tivoli Storage Productivity Center

With the vSphere Web Client extension for Tivoli Storage Productivity Center, you can manage virtual storage in a vSphere environment. You can use the vSphere Web Client extension to complete provisioning tasks and to view information about storage resources in the vSphere Web Client.

Provisioning storage in the vSphere Web Client extension

You can use the vSphere Web Client extension to provision storage capacity on storage systems that are monitored by Tivoli Storage Productivity Center. You can provision file and block storage and use the service classes and capacity pools that you create in Tivoli Storage Productivity Center. Additionally, you can choose to create data stores on your storage volumes.

Viewing storage reports in the vSphere Web Client extension

You can use the vSphere Web Client extension to view information about the storage systems that are monitored by Tivoli Storage Productivity Center. You can view the mapping of your virtual storage resources to storage systems that are monitored by Tivoli Storage Productivity Center. You can view metrics for your storage systems. You can also view information about the fabric that is connected to a storage adapter. From the storage reports, you can click links to open the resources in the Tivoli Storage Productivity Center web-based GUI.

vSphere supported locales

The locales that are supported by the vSphere Web Client extension for Tivoli Storage Productivity Center depend on the vSphere support. In vSphere 5.x, the following locales are supported:

- English: en_US
- German: de_DE
- Japanese: ja_JP
- Simplified Chinese: zh_CN
- French: fr_FR
- Korean: ko_KR

When a locale does not have a translated version, the vSphere Web Client extension is displayed in English, which is the default language.

Related concepts:

"Managing resources in a vSphere environment" on page 458
You can use the vSphere Web Client extension for Tivoli Storage Productivity Center to manage and monitor your virtual environment and storage devices.

"Reporting in a vSphere environment" on page 1643
You can use the vSphere Web Client extension for Tivoli Storage Productivity Center to view reports on your virtual environment and storage devices.

Alerts

Determining when and how you are alerted to conditions and violations of performance thresholds within a storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines the data about your resources for the conditions and performance thresholds that you specify.

Alerts are triggered by conditions and violations of performance thresholds that are detected during data collection and event processing. For storage systems such as SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, or the XIV system, events are polled every minute from the resource. For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to Tivoli Storage Productivity Center when conditions change on the resource.

Triggering conditions for alerts

The conditions that trigger alert notifications depend on the type of resource that you are monitoring. Some triggering conditions, like performance thresholds, require you to enter values for triggering alerts. In general, the following types of conditions can trigger alerts:

- A performance threshold was violated
- A data collection job did not complete
- A change occurred in the storage infrastructure

For example, you can use performance thresholds to be notified when the total I/O rate for storage systems falls outside a specified range. This information can help you identify areas in your storage infrastructure that are over utilized or under utilized. Tivoli Storage Productivity Center provides many metrics for measuring performance and determining violations of the thresholds that you specify.

Alert notifications and triggered actions

When an event occurs and triggers an alert, the alert is written to a log. You can also select one or more other ways to be notified of the event. These alert notifications include SNMP traps, Tivoli Enterprise Console events, login notifications, entries in Windows event log or UNIX syslog, and emails. Additionally, if a Storage Resource agent is deployed on a monitored server, you can run a script or start an IBM Tivoli Storage Manager job in response to the alert.

Acknowledging alerts: Some alerts are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify other alerts that must be reviewed and resolved.

Prerequisites for using alerts

The following conditions must be met to successfully use alerts:

- Data collection schedules are configured and scheduled to run regularly. For example, to detect violations of performance thresholds, you must run performance monitors to collect performance data about resources. Running performance monitors regularly also helps to establish a history of performance for trending analysis.
- If you want to be notified about an alert in some way other than an entry in the log file, such as using SNMP traps, Tivoli Enterprise Console events, or email, you must configure those alert destinations before you use the alert.
- If an alert is triggered based on an SNMP trap from the monitored resource, you must properly configure the SNMP server of the monitor resource to enable Tivoli Storage Productivity Center to listen to SNMP traps. The default port number is 162, and the default community is public.

Related concepts:

"Alerting" on page 341
Determining when and how you are alerted to conditions and violations of performance thresholds within your storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines the data about your resources for the conditions and performance thresholds that you specify.

Related tasks:

"Configuring alert destinations" on page 355
In the stand-alone GUI, configure the control information that is used to generate alerts when a corresponding event is discovered. You can specify default configuration settings for SNMP, TEC/OMNIbus, and email-directed alerts.

"Acknowledging alerts" on page 359
Some alerts in Tivoli Storage Productivity Center are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify the other alerts in lists that must still be reviewed and resolved.

Reporting

Use reporting functions to view overview and detailed information about your storage.

From the web-based GUI, you can view predefined reports and create custom reports about the resources managed by Tivoli Storage Productivity Center. You access reports from the web-based GUI, and work with the reports in Tivoli Common Reporting. Tivoli Common Reporting provides the reporting service for Tivoli Storage Productivity Center reports. Tivoli Common Reporting includes Cognos reporting software.

You can view over 70 predefined reports about the capacity and performance of your resources in Tivoli Common Reporting. Charts are automatically generated for most of the predefined reports. Depending on the type of resource, the charts show statistics for space usage, workload activity, bandwidth percentage, and other statistics.

You can schedule reports and specify to create the report output in HTML, PDF, and other formats. You can also configure reports to save the report output to your local file system, and to send reports as email attachments.

In the stand-alone GUI, use the Reporting facilities in the different managers of IBM Tivoli Storage Productivity Center to view both overview and detailed information about your storage. The over 500 available reports are constructed from the statistical information accumulated in the repository and enable you to select exactly the computers, storage systems, switches, clusters, disks, file systems, hypervisors, users and even other Tivoli Storage Productivity Center servers upon which you want to report. You can view reports as tables of data or as graphical charts, which provides you with the flexibility to get a very detailed look at your information or a very high level overview.

You must collect information about your environment before you can use reports to view details about the storage resources in it. You can use monitoring jobs in Tivoli Storage Productivity Center such as probes, scans, pings, and Tivoli Storage Productivity Center server probes, and discovery jobs to gather comprehensive information and statistics about your storage resources.

Role-based authorization

Roles determine the functions that are available to users of Tivoli Storage Productivity Center. When a user ID is authenticated to Tivoli Storage Productivity Center through the GUI, CLI, or APIs, membership in an operating system or LDAP group determines the authorization level of the user.

In Version 5.2, the roles that were previously defined in Tivoli Storage Productivity Center were consolidated into a smaller set of roles. The following table shows how the roles in versions earlier than 5.2 are mapped to the current set of roles:

Table 30. How roles in previous versions of Tivoli Storage Productivity Center are mapped to the roles in version 5.2 and later

Roles in previous versions	Roles in version 5.2 and later	Authorization level
Superuser Productivity Center administrator Disk administrator Fabric administrator Data administrator Tape administrator	Administrator	This role has full access to all monitoring and administrative functions. At least one group must have the Administrator role. **Note:** When Tivoli Storage Productivity Center is first installed, the following operating system groups are assigned the Administrator role: • Windows: Administrators • UNIX and Linux: root • AIX: system
Disk operator Fabric operator Data operator Tape operator	Monitor	In the web-based GUI, this role has access to the following read-only functions: • Viewing and exporting information about monitored resources • Viewing and acknowledging alerts • Viewing tasks and data collection jobs • Viewing data paths • Opening management GUIs • Opening logs • Tiering storage In the stand-alone GUI, this role has access to the following functions: • Viewing data that is collected by Tivoli Storage Productivity Center • Creating, generating, and saving reports
This role did not exist in versions 5.1 or earlier.	External Application	This role enables users of other applications to use the provisioning capability of Tivoli Storage Productivity Center to provision storage. For example, a VMware user with this role can provision storage in the vSphere GUI by using the vSphere Web Client extension for Tivoli Storage Productivity Center. If you assign the External Application role to the user, you must also assign one or more service classes to the user. This role does not enable users to log in to the Tivoli Storage Productivity Center GUIs.

Tips:

- If a user belongs to multiple groups and the groups have different roles, the role with the highest level of authorization is granted to the user. For example, if a user belongs to a group that is assigned the Administrator role and also belongs to a group that is assigned a Monitor role, the user is granted the authorization of the Administrator role.
- If a user is not a member of a group that is assigned a Tivoli Storage Productivity Center role, no access is granted to that user.
- For rollup reports, you need Administrator authority to do the following actions:
  - Add, remove, or modify the Tivoli Storage Productivity Center subordinate server that the master server is monitoring.
  - Create or run probe jobs that include Tivoli Storage Productivity Center subordinate servers.

  Any Tivoli Storage Productivity Center role can generate rollup reports.
- The Tivoli Storage Productivity Center installation program adds the administrator, external application, and monitor roles to the Tivoli Storage Productivity Center installation user.

Fabrics and zones

Use the Tivoli Storage Productivity Center to learn more about fabrics and zones.

Fabrics and zones in the web-based GUI

You can monitor switches and fabrics in the web-based GUI. You can also enable automatic zoning so that new zones are created during storage provisioning, if necessary, to connect a server to a storage system.

In the web-based GUI, you can add switches and fabrics for monitoring. You can add the following types of switches:

- Brocade
- Cisco
- QLogic

After switches and fabrics are probed, you can view detailed status and performance information about the switches and fabrics. You can also test the connection to the switch or fabric.

Information that you can view about a fabric includes the following properties:

- The fabric type
- The number of physical and virtual switches in the fabric
- The number of ports that are on all the switches in the fabric
- The number of switch ports that are online and connected to other ports in a fabric
- The name of the principal switch of the fabric

Information that you can view about a switch includes the following properties:

- The number of ports on a switch
- The number of ports that are connected to a storage resource
- The name of the fabric where a switch is a member
- The name of the vendor or manufacturer for a switch

From the Details page of a fabric, you can open the Zone Sets page to view information about the zone sets on the fabric, and to identify which is the active zone set.

You can enable automatic zoning in the web-based GUI. When automatic zoning is enabled, Tivoli Storage Productivity Center creates new zones during provisioning if new zones are needed to connect the storage system with the server. When you set the zoning policy, you can specify whether zone changes from automatic zoning are made to the active zone set, or to a new inactive active zone set.

Zones, zone aliases, and zone sets

Zones, zone aliases, and zone sets allow logical grouping of ports and storage devices within a storage area network. This section describes zoning concepts and elements.

In a storage area network a *zone* is a logical grouping of ports to form a virtual private storage network. Zones that belong to a single SAN can be grouped into a *zone set*, which can be activated or deactivated as a single entity across all switches in the fabric. A zone set can contain one or more zones, and a zone can be a member of more than one zone set. Using zoning, you can automatically or dynamically arrange fabric-connected devices into logical groups across a physical fabric.

Ports and devices in a zone are called *zone members*. A zone can contain one or more zone members. Ports that are members of a zone can communicate with each other, but they are isolated from ports in other zones. Devices, however, can belong to more than one zone. A *zone alias* is a collection of zone members. A zone alias can be added to one or more zones.

Note: In some cases, inactive zones might not have any zone members. Activating an inactive zone set that contains empty zones will fail if a switch does not support empty zones in active zone definitions.

Zone membership can be specified by:

- The N_Port_Name of the N_Port connected to the switch (also known as WWN zoning or port name zoning)
- The N_Port address identifier assigned during fabric login (also known as PortId or FCID zoning)
- The Node_Name associated with the N_Port
- The Domain identification (Domain_ID) and physical port ID of the Switch Port to which the N_Port is attached (also known as domain port zoning)
- An alias name

Zoning supports the use of aliases, which are meaningful names assigned to devices. An alias can also be a group of devices that are managed together to make zoning easier.

There are two types of zoning:

Hardware zoning (port zoning)
In hardware zoning (also called *port zoning*), the members of a zone are the physical ports on a fabric switch.

Software zoning (WWN zoning)
Software zoning uses the Simple Name Server (SNS) that runs inside a

fabric switch. It is based on the node WWN or port WWN of the zone members to be included. Software zoning lets you create symbolic names for the zones and zone members.

A *default zone* is a group of devices that are not members of the active zone set. These can communicate with each other but not with members of any other zone. Default zoning is enabled by default. You can use a switch element manager to configure the Default Zone option to enable or disable the default zone independently of the active zone set.

Note:

1. If the default zone is disabled, devices that are not members of the active zone set cannot communicate.
2. If the default zone is disabled and no zone set is active, no devices can communicate.
3. If default zoning is enabled, deactivating the active zone set makes all devices members of the default zone. If default zoning is disabled, all communication stops.
4. If you activate one zone set while another zone set is active, the currently active zone set is deactivated.
5. If your EFC Manager manages multiple fabrics, ensure that you have the correct zone set for the fabric you are currently updating.

You can view zones in SANs using the topology viewer console. Fabric Manager lets you view zone membership.

Note: Certain types of switches and the types of agents configured can affect the zoning operations that you can perform.

Switch zoning capabilities

When you select a fabric in which to perform zoning operations, Fabric Manager determines the capabilities of the switches in the SAN and limits the zoning operations based on that information.

To see a list of supported switches and their capabilities, default values, ranges, and possible effects, go to http://www.ibm.com/support/docview.wss?uid=swg27027300.

Zoning for Cisco MDS 9000 devices:

Cisco switches support virtual SANs (VSANs), which is the logical partitioning of a fabric into multiple fabrics. The overall network is referred to as the physical infrastructure, and the logical fabrics are the VSANs. Fabric Manager provides basic zone discovery for the Cisco MDS 9000 series of devices using the FC-GS-3 protocol.

The following tasks are supported:

- Creating, deleting, and updating zones
- Creating, deleting, updating, activating, and deactivating zone sets

Fabric Manager supports these zone member types:

- N_Port WWN
- FC ID (the fibre channel ID of an N_port attached to the switch)

Unsupported zone and zone member types are displayed as non-standard on the topology display. Fabric Manager does not allow zone management of zones and zone sets that contain unsupported members.

Brocade switches in zones:

Brocade switches support some non-standard zones such as quick loop zones, fabric assist zones, and protocol zones. If the switch configurations have these zones already defined, Fabric Manager preserves them and does not modify them in any way. If they are part of a zone set that is active at some time, the devices that are part of such zones that are also online are displayed in the topology Zone View.

You can create, change, and delete non-standard zones by using the Brocade switch management application.

Note: Tivoli Storage Productivity Center allows you to add empty zones to an inactive zone definition. Activation of a zone set containing empty zones may still fail if the switch does not support empty zones in active zone definitions.

Agents

The IBM Tivoli Storage Productivity Center uses agents to gather data: Common Information Model (CIM) agents, Storage Resource agents, Fabric agents, and out-of-band Fabric agents.

Common Information Model agents

Common Information Model (CIM) agents enable communication between storage devices, such as storage systems and fabric switches, and IBM Tivoli Storage Productivity Center. Each storage device that you want to manage must have a CIM agent either installed or embedded unless it is a device that uses Native API connection.

CIM agents are provided by the vendor of the storage subsystem or fabric switch. Each vendor provides unique agent code for their family of storage devices. This code implements a Common Information Model Object Manager (CIMOM) that conforms to the Storage Management Initiative Specification (SMI-S) of the Storage Networking Industry Association (SNIA).

The CIM agent enables communication between the storage device and Tivoli Storage Productivity Center. Commands and responses are transmitted between Tivoli Storage Productivity Center and the CIM agent using an XML transport layer. The CIM agent to storage device layer uses a vendor-specific proprietary interface.

The CIM agent usually must be installed and configured, so that it can identify the storage devices with which it communicates. Some storage devices, such as Cisco fabric switches, contain embedded CIM agents and so do not require that CIM agents be installed, although Tivoli Storage Productivity Center must be configured to point directly to the storage devices that contain the embedded CIM agents.

CIM agents can be referred to as SMI-S providers, CIM proxy agents, and CIMOM agents. CIM agents can be embedded in the device or installed on a separate computer.

Note:

- Do not install multiple CIM agents on a single computer because of port conflicts.
- Do not install a CIM agent on the system where a Tivoli Storage Productivity Center server component is installed.

Storage Resource agents

Use Storage Resource agents to collect information about storage resources such as servers, virtual machines, workstations, HBAs, and fabrics.

You must deploy Storage Resource agents on resources where you want to gather the following information:

- Asset information
- File and file system attributes
- Database application information
- Network-attached storage (NAS) device information
- Topology information
- Information about zoning and the fabrics that are visible to the server

Tip: You can also monitor servers without deploying a Storage Resource agent. When you add an agentless server, Tivoli Storage Productivity Center correlates information about that server with the known host connections on monitored resources. If a match is made between the server and a monitored resource, you can view topology information and the capacity and performance of the storage that is assigned to that server. For more information about when to deploy a Storage Resource agent or add an agentless server, see "Adding servers" on page 262.

Deployment guidelines and limitations for Storage Resource agents:

There are a number of guidelines and limitations that you must consider when managing the Storage Resource agents in your environment.

Capacity guidelines for Storage Resource agents: For the capacity guidelines for Storage Resource agents by Tivoli Storage Productivity Center version, see http://www.ibm.com/support/docview.wss?uid=swg21424912.

Use the following information when you deploy Storage Resource agents:

- Multiple Storage Resource agents that are probing or scanning the same storage resources
- Platforms that support the deployment of Storage Resource agents
- Product functions that are not available for storage devices monitored by Storage Resource agents
- Required authority for deploying Storage Resource agents
- Orphan zones
- Firewalls and Storage Resource agent deployments
- Multiple computers in a deployment
- Communication between the Tivoli Storage Productivity Center server and a Storage Resource agent
- Daemon and non-daemon services
- Authentication between the Tivoli Storage Productivity Center server and a Storage Resource agent
- Storage Resource agents on the same computer

- Time zones for computers monitored by Storage Resource agents
- AIX systems using RSH protocol
- Deployments on Windows - NetBIOS setting
- Deployments on Windows - User Account Control (UAC) remote restrictions
- Additional information about Storage Resource agents

Multiple Storage Resource agents that are probing or scanning the same storage resources

If multiple Storage Resource agents are set up to probe or scan the same storage resources, the Storage Resource agent that was added to Tivoli Storage Productivity Center first is used for the probe or scan. Therefore, only data that is gathered by the first Storage Resource agent is shown.

Platforms that support the deployment of Storage Resource agents

For a list of platforms on which you can deploy Storage Resource agents, go to the Tivoli Storage Productivity Center support site at http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition, click the **Documentation** link, and enter `Platform Support: Agents, Servers and Browsers` in the **Search** box. Click the link to the document for the appropriate release.

Product functions that are unavailable for resources that are monitored by Storage Resource agents

Before you deploy a Storage Resource agent, make sure that the product functions you want to use on the monitored devices are available for those agents. The following functions are not available for resources that are monitored by Storage Resource agents:

- Certain relational database monitoring. For list of relational databases that can be monitored by Storage Resource agents, go to the Tivoli Storage Productivity Center Interoperability Matrix and click the link for *Agents, Servers and Browsers*.
- The reporting of HBA, fabric topology, or zoning information for fabrics that are connected to hosts that are running Linux on IBM System z hardware. These limitations also apply to Storage Resource agents on all guest operating systems for VMware configurations.

Required authorities for deploying and running Storage Resource agents

Before you can create deployment schedules and deploy Storage Resource agents on target computers, you must meet the following requirements:

- To create deployment schedules, you must be logged in to Tivoli Storage Productivity Center with a user ID that has the **Administrator** role. For more information about user roles, see "Role-based authorization" on page 77.
- To deploy Storage Resource agents on target computers, you must provide a user ID that has administrative rights on those computers. You enter this ID when you create a deployment schedule. Tivoli Storage Productivity Center uses this ID to log on to the target computers and install and configure the necessary runtime files for the agents.

The user under which a Storage Resource agent (daemon or non-daemon) runs must have the following authorities on the target computers:

- On the Linux or AIX operating systems, the user must have root authority. By default, an agent runs under the user 'root'.

- On the Windows operating systems, the user must have Administrator authority and be a member of the Administrators group. By default, a Storage Resource agent runs under the 'Local System' account.

Orphan zones

Storage Resource agents do not collect information about orphan zones. An orphan zone is a zone that does not belong to at least one zoneset.

Firewalls and Storage Resource agent deployments

Before you can deploy a Storage Resource agent on a computer, you must turn off the firewall on that computer. If you do not turn off the firewall, the deployment fails.

To turn off the firewall on a Windows 2008 computer, complete the following steps:

1. Open Administrative Tools.
2. Click **Windows Firewall with Advanced Security**.
3. Click **Windows Firewall Properties**.
4. Change the **Firewall state** field to **Off** on the following tabs:
   - **Domain Profile**
   - **Private Profile**
   - **Public Profile**
5. Click **OK** to accept the changes and exit.
6. Deploy a Storage Resource agent to the Windows 2008 computer.

Multiple computers in a deployment

You can define a deployment schedule to include multiple computers on which to install Storage Resource agents. The computers that you include in a deployment schedule must share administrative user ID and password. Tivoli Storage Productivity Center uses these user credentials to log on to the computers when you installing Storage Resource agents. If the computers in a deployment schedule do not share administrative user credentials, you must create separate deployment jobs for them.

Tip: When you deploy a Storage Resource agent to multiple computers, a globally unique identifier (GUID) is created for each computer (if one does not exist).

Communication between the Tivoli Storage Productivity Center server and a Storage Resource agent

The Tivoli Storage Productivity Center server connects to a monitored computer when a Storage Resource agent is deployed and whenever a data collection schedule runs against that agent.

During deployment, the server communicates with the target computer by using one of the following protocols:

- Windows server message block protocol (SMB)
- Secure Shell protocol (SSH)
- Remote execution protocol (REXEC)
- Remote shell protocol (RSH)

After deployment, the type of communication between the server and agent on that computer depends on whether you deployed the agent as daemon service or non-daemon service. See the *IBM Tivoli Storage*

Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide for more information about these protocols.

Daemon and non-daemon services

You can deploy a Storage Resource agent agent as a daemon or non-daemon service:

- A Storage Resource agent that is deployed as a daemon service runs in the background on the monitored computer and listens for requests from the Tivoli Storage Productivity Center server. Connectivity between the server and agent is established by using SSL. The server and agent have their respective certificates and no additional information is required besides those certificates and the security that is provided by the SSL protocol.
- A Storage Resource agent deployed as a service on demand (non-daemon service) runs as a stand-alone executable file on the monitored computer. Communication from the server to the agent uses the same protocol that was used during the deployment of the agent. Communication from the agent to the server uses SSL.

Authentication between the Tivoli Storage Productivity Center server and a Storage Resource agent

Tivoli Storage Productivity Center requires the correct authentication information (user name, password, port, certificate location, or passphrase) for monitored computers each time it communicates with Storage Resource agents on those computers. If the authentication information changes for a host computer on which a Storage Resource agent is deployed, the authentication information for that agent must be updated by using the **Modify Agents** > **Update Credentials** action on the Servers page in the web-based GUI.

Storage Resource agents on the same computer

You cannot deploy a Storage Resource agent on a computer where a Storage Resource agent is already installed *and* pointing to the same Data server. You can deploy a Storage Resource agent on the same computer as another Storage Resource agent if those agents communicate with different Data servers and use different ports when you listen for requests.

Time zones for computers that are monitored by Storage Resource agents

The time zones of computers that are monitored by Storage Resource agents are shown as Greenwich Mean Time (GMT) offsets in Tivoli Storage Productivity Center reports. For example, a computer in Los Angeles shows the following time zones in the By Computer report in Asset reporting:

```
(GMT-8:00) GMT-8:00
```

Linux and AIX operating systems by using Remote Shell protocol (RSH)

If RSH is configured to use a user ID and password, the connection fails. To successfully connect to a system by using RSH, you must set up the `.rhosts` file (in the home directory of the account). RSH must be configured to accept a login from the system that is running your application.

Deployments on Windows operating systems - NetBIOS setting

If you want to install a Storage Resource agent on Windows targets, the **Enable NetBIOS over TCP/IP** option must be selected in the Control Panel settings for the computer's network connections properties. To set this option, complete the following steps:

1. Open Windows Control Panel.
2. Select **Network and Dial-Up Connections** > *some_connection* > **Properties** > **Internet Protocol (TCP/IP)** > **Advanced** > **WINS** > **Enable NetBIOS over TCP/IP**.

See the documentation for your firewall to determine if these ports are not blocked for inbound requests.

To determine if security policies are blocking the connection ports, open Administrative Tools.

Depending on whether your policies are stored locally or in Active Directory, follow these directions:

Locally stored policies

For locally stored policies, complete the following steps:

1. Open Windows Administrative Services.
2. Click **Local Security Policy** > **IP Security Policies on Local Computer**.

Policies stored in Active Directory

For policies stored in Active Directory, examine the IP security policies and edit or remove filters that block the ports:

- Click **Administrative Tools** > **Default Domain Security Settings** > **IP Security Policies on Active Directory**.
- Click **Administrative Tools** > **Default Domain Controller Security Settings** > **IP Security Policies on Active Directory**.

For all Windows systems, the Server service must be running to connect to a Windows system by using the Windows protocol.

The following table lists the ports reserved for NetBIOS. Ensure that these ports are not blocked.

Port	Description
135	NetBIOS Remote procedure call. (Not currently used.)
137	NetBIOS name service.
138	NetBIOS datagram. (Not currently used.)
139	NetBIOS session (for file and print sharing).
445	CIFS (on Windows XP).

For Windows Server 2008, shares must be shared for the Guest or Everyone accounts, and password protected sharing must be disabled. To disable password protected sharing, follow these steps:

1. Click **Control Panel** > **Networking and Sharing Center**.
2. Click the down arrow next to **Password protected sharing**.
3. Click **Turn off password protected sharing**.
4. Click **Apply**.
5. Exit from the Control Panel.

Deployments on Windows 2008 - User Account Control (UAC) remote restrictions

If you are planning to install Storage Resource agents remotely on a Windows 2008 operating system, you must disable the User Account Control (UAC) remote restrictions on the Windows operating system. User Account Control is a security component on Windows operating systems.

Tip: This task contains steps that tell you how to modify the registry. Serious problems might occur if you modify the registry incorrectly. Therefore, make sure that you follow these steps carefully. For added protection, back up the registry before you modify it. Then you can restore the registry if problems occur. For information about how to back up and restore the registry, see http://support.microsoft.com/kb/322756/.

To disable UAC remote restrictions, follow these steps:

1. Open the Windows Run window.
2. Enter **regedit** and click **OK**.
3. Locate and click the following registry subkey:

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\
       Policies\System
   ```

4. If the `LocalAccountTokenFilterPolicy` registry entry does not exist, follow these steps:
   a. On the **Edit** menu, click **New** > **DWORD Value**.
   b. Enter `LocalAccountTokenFilterPolicy` as the name for the DWORD value in the content pane.
   c. Click **Enter**.
   d. Click **LocalAccountTokenFilterPolicy**, and click **Modify**.
   e. In the Edit DWORD Value window, enter **1**, then click **OK**.

      This value can be 0 or 1:

      0 This value builds a filtered token. This value is the default value. The administrator credentials are removed.

      1 This value builds an elevated token.

   f. Exit the registry editor.

Additional information about Storage Resource agents

See *Planning for Storage Resource agents* in the *IBM Tivoli Storage Productivity Center Installation and Configuration Guide* for more information about Storage Resource agents.

Out-of-band fabric agents

Out-of-band fabric agents issue SNMP queries to the fabric switches and collect topology information.

Be sure that your out-of-band fabric agents point to each switch in the Storage Area Network (SAN) that you want to manage.

Out-of-band fabric agents are necessary for Virtual storage area network (VSAN) information for Cisco switches. For Brocade fabrics, if no SMI-S agent is configured, Out-of-band fabric agents are necessary for Zoning information for Brocade switches.

IBM Tivoli Storage Productivity Center only supports SNMP V1.

See the Planning > Planning for Fabric Manager > *Configuring Cisco switches* topic in the *Installation and Configuration Guide* for more information about using Out-of-band fabric agents and Cisco switches.

Tivoli Storage Productivity Center Monitoring Agent

The IBM Tivoli Storage Productivity Center Monitoring Agent can be used by the IBM Tivoli Enterprise Monitoring Server to monitor systems in your enterprise. This agent is an optional program you can install and use in your enterprise.

IBM Tivoli Monitoring monitors and manages system and network applications on a variety of operating systems, tracks the availability and performance of your enterprise system, and provides reports to track trends and troubleshoot problems. The Tivoli Storage Productivity Center Monitoring Agent participates in that environment.

For performance reasons, install the Tivoli Storage Productivity Center server on a separate system from the IBM Tivoli Monitoring server. You can install the Tivoli Storage Productivity Center Monitoring agent on any system you want to monitor.

Before using the Tivoli Storage Productivity Center Monitoring Agent, you must have IBM Tivoli Enterprise Monitoring Server installed. For information about the IBM Tivoli Enterprise Monitoring Server, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.1/welcome.htm.

For information about installing IBM Tivoli Monitoring, see the Quick Start Guide at http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc/GI11-8058-00.htm.

To use the Tivoli Storage Productivity Center Monitoring Agent, install this agent on the system you want to monitor.

To see the type of reports displayed by the Tivoli Storage Productivity Center Monitoring Agent, see "Tivoli Storage Productivity Center Monitoring Agent reports" on page 1636.

Single sign-on

You can use the *Single sign-on* authentication process to enter one user ID and password to access multiple applications. *Single sign-on* integrates with the launch and launch-in-context features to enable the smooth movement from one application to another application.

For the current version of Tivoli Storage Productivity Center, *single sign-on* is available for the following applications:

- Tivoli Storage Productivity Center for Replication
- IBM System Storage DS8000
- IBM System Storage SAN Volume Controller
- Storwize V3500
- Storwize V3700
- IBM Storwize V7000

For Tivoli Storage Productivity Center for Replication, *single sign-on* is enabled by default.

For DS8000 Storage Manager, SAN Volume Controller, and Storwize V7000, *single sign-on* must be enabled.

To enable *single sign-on*, see the following topics in the Tivoli Storage Productivity Center information center:

- Configuring Tivoli Storage Productivity Center for DS8000 LDAP authentication
- Configuring Tivoli Storage Productivity Center for SAN Volume Controller and Storwize V7000 *single sign-on*

How single sign-on works

A *single sign-on* environment requires a centralized authentication repository that is accessed by all applications within the environment. When you log on to Tivoli Storage Productivity Center, a Lightweight Third-Party Authentication (LTPA) token is created and used for *single sign-on* authentication when you start other applications from within the Tivoli Storage Productivity Center GUI. LTPA is the security technology that is used in the IBM WebSphere Application Server for passing the user authentication information between applications. To use the LTPA tokens, each application must have the same set of public keys to encrypt the user's information. The authenticating service uses the corresponding private keys to decrypt the user's information for authentication. As an extra security mechanism, LTPA tokens expire after a determined amount of time. By default, the tokens expire after 8 hours. You can change the LTPA token expiration time that uses the IBM WebSphere Application Server administration console.

During the period between when, you log on to Tivoli Storage Productivity Center and when you start other applications, such as element managers, from the Tivoli Storage Productivity Center GUI, the following conditions might occur:

Condition	Explanation
The user password that was used to log on toTivoli Storage Productivity Center is changed in the user repository.	Under this condition, the original user credentials that were used to log in to Tivoli Storage Productivity Center are used to access other applications until the timeout period for the LTPA token that is used for *single sign-on* expires. When the LTPA token expires, you are prompted to reenter your user ID and password when you attempt to start another application that uses *single sign-on*.
The user repository is not accessible.	Under this condition, the *single sign-on* feature does not work. You are always prompted to reenter your user ID and password when you attempt to start another application through *single sign-on*.

Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication use their respective IBM WebSphere Application Server instances to authenticate LTPA tokens. However, other applications, such as the System Storage DS8000 element manager and other element managers that do not run within an IBM WebSphere Application Server instance must be configured for LDAP.

The *single sign-on* feature is not supported by the Tivoli Storage Productivity Center command-line interface (CLI).

Selecting the user authentication method to use with single sign-on

During the installation, Tivoli Storage Productivity Center uses the operating system (OS) users and groups on the server where Tivoli Storage Productivity

Center is installed for authentication and authorization. After the installation, you can configure Tivoli Storage Productivity Center to use LDAP as the authentication and authorization repository.

OS Authentication
This method authenticates the user against the users that are defined for the local operating system (OS).

LDAP/Active Directory
This method authenticates the user against a Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory repository.

If OS authentication is selected, the use of the *single sign-on* feature is limited. OS authentication does not support *single sign-on* for element managers, even when the element manager is installed on the same computer as Tivoli Storage Productivity Center. LDAP or Microsoft Active Directory authentication supports *single sign-on* for element managers regardless of where they are installed.

You can change the user authentication method you use for the WebSphere Integrated Solutions Console that is associated with the TPC 5.2 Web Server

Related concepts:

"Changing the user authentication configuration" on page 146
The Tivoli Storage Productivity Center installation program establishes a default authentication configuration by using the federated repositories feature of the IBM WebSphere Application Server. You can change this authentication configuration.

Related tasks:

"Changing the LTPA token expiration for single sign-on" on page 173
For single sign-on, the user authentication information is passed between applications by using Lightweight Third-Party Authentication (LTPA) tokens. By default, the tokens expire after eight hours. You can change the LTPA token expiration time by using the WebSphere Integrated Solutions Console.

Launch and launch in context

The launch and launch in context features are used to start the Tivoli Storage Productivity Center stand-alone GUI and to start other applications from Tivoli Storage Productivity Center. The terms inbound and outbound are used to differentiate starting Tivoli Storage Productivity Center (inbound) and starting another application from Tivoli Storage Productivity Center (outbound).

Inbound launch in context refers to starting Tivoli Storage Productivity Center at a specific point in the stand-alone GUI. You can use inbound launch in context in either of the following ways:

- Configure an application such as Tivoli Enterprise Portal to start Tivoli Storage Productivity Center at functionally related locations in the stand-alone GUI.
- Start Tivoli Storage Productivity Center from Java Web Start, the command line, or the **Start** menu or Start page in Windows. Use parameters to specify the starting location in the stand-alone GUI.

Outbound launch refers to starting an application at the starting point for that application. Use outbound launch to open the management GUIs of monitored storage systems and switches from the Tivoli Storage Productivity Center web-based GUI.

Inbound launch in context

Start Tivoli Storage Productivity Center at a specified point in the stand-alone GUI.

Prerequisite: You must have View and Modify permissions to use the launch in context feature.

You can use the launch in context feature to perform the following tasks in Tivoli Storage Productivity Center:

- Create a capacity report
- Create a wasted space report

You can run the launch in context feature in the following ways:

- **Use Java Web Start** from a Web browser. See "Launching with Java Web Start" on page 613.
- **Use the command line on Windows, AIX, Linux, or UNIX** to enter the command and arguments on the command line. See "Launching from the command line" on page 615.
- **Use the Windows Start menu or Start page**. See "Launching from Windows" on page 616.
- **Use the Tivoli Enterprise Portal GUI** to define, edit, and launch commands. See "Launch from Tivoli Enterprise Portal" on page 617.

The launch in context feature has a number of parameters. See "Launch in context parameters" on page 1729, "Launch-in-context task parameters" on page 1731, and "Available windows" on page 1731.

Outbound launch

Resources such as storage systems and switches can have their own management GUIs or element managers. In theTivoli Storage Productivity Center web-based GUI, you can open the start page for these management GUIs.

The action for opening a management GUI is available for storage systems and switches under the following conditions:

- The resource is monitored by Tivoli Storage Productivity Center.
- The resource has a management GUI that is web-based or can be accessed with Java Web Start.

Related tasks:

"Opening the management GUI for a resource" on page 440
Resources such as storage systems and switches can have their own management GUIs or element managers. In Tivoli Storage Productivity Center, you can open the start page in these GUIs from resource list pages, such as the Storage Systems page and Switches page.

Chapter 3. Scenarios

This section contains scenarios which illustrate how to perform specific tasks using IBM Tivoli Storage Productivity Center.

The provided scenarios set out situations that can be solved using Tivoli Storage Productivity Center. When using these scenarios you should remember the following:

- The scenarios are intended to be step-by-step instructions that you follow to complete a specific job. Each scenario contains multiple tasks that must be followed in the order given.
- Specific values are given in the scenarios, such as user name, IP Address, probe name, and so forth. These are for illustration purposes only and you must replace them with values appropriate for your system.
- The scenarios do not provide in-depth conceptual information about the tasks. See the appropriate topics in the IBM Tivoli Storage Productivity Center Information Center for more information about the concepts behind the tasks.
- You will complete a scenario successfully if you follow the instructions correctly and your system performs as expected.

Optimizing the performance of storage virtualizers

In these scenarios, you run the Balance Pools and the Analyze Tiering wizards to help resolve performance hot spots and manage storage resources efficiently.

You can optimize the placement of volumes on the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

To optimize the placement of volumes in storage virtualizer pools, you must complete the following tasks:

- Probe the storage virtualizers.
- Set the tier level of the storage pools that you want to analyze.
- Collect performance data to gauge the utilization percentage for pools and the average workload activity of the volumes.

With the Balance Pools wizard, you optimize storage performance by redistributing volume workloads across pools on the same tier.

With the Analyze Tiering wizard, you optimize storage performance by re-tiering volumes to balance pools and by setting minimum thresholds to regulate the tiering of volumes.

Related concepts:

"Tiering storage overview" on page 429
To manage storage efficiently and reduce storage costs, optimize the placement of volumes in storage virtualizer pools.

Related tasks:

© Copyright IBM Corp. 2001, 2013

"Collecting performance data" on page 299
Before you can monitor the performance of the resources in your environment, you must collect data about those resources. Tivoli Storage Productivity Center uses performance monitors to collect metrics for measuring the performance of storage systems and switches. Use the Performance Monitors page to view and manage performance monitors and to track threshold violations.

"Setting the tier level of storage pools" on page 433
To analyze tiering and to balance pools, you must set the tier level of storage pools.

"Balancing pools" on page 436
To distribute the workload of volumes across pools on the same tier, use the Balance Pools wizard. To improve the performance of the pools, recommendations are generated to move volumes to other pools on the same tier.

"Managing tasks for tiering storage, balancing pools, and transforming storage" on page 452
To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the Tasks page in the web-based GUI to manage all the tasks that are used by Tivoli Storage Productivity Center to optimize storage. Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.

Redistributing volume workloads to balance pools

In this scenario, you identify performance hot spots and then run the Balance Pools wizard to help resolve them.

The pools that you want to analyze are on a SAN Volume Controller with two DS8000 back-end storage systems. The name of the SAN Volume Controller is `svc_payment_records`, and the names of the DS8000 back-end storage systems are `ds_pr1` and `ds_pr2`.

The value for the maximum utilization of the pools is used as a threshold value to identify whether the pools on a storage tier require balancing. The maximum utilization threshold is also used to ensure that the performance capability of destination pools does not become degraded when volumes are added to the pool. For example, recommendations are not generated to move volumes to target pools if the threshold for the maximum utilization of the pool is violated.

In this scenario, if the utilization value of a pool exceeds 80%, the performance of the pool deteriorates and it causes delays in the response times for sales transactions. When you run the Balance Pools wizard, you set the maximum utilization value for the pools.

The information in Table 31 is provided to illustrate how the volumes are balanced in the pools on the storage virtualizer.

Table 31. Utilization of the pools on the storage virtualizer `svc_payment_records`

Back-end storage system	Tier	Pool	Current utilization (%)
ds_pr1	1	pool_ds81_1	88
ds_pr2	1	pool_ds82_101	10
ds_pr1	2	pool_ds81_2	90
ds_pr1	2	pool_ds81_3	65

Table 31. Utilization of the pools on the storage virtualizer `svc_payment_records` (continued)

Back-end storage system	Tier	Pool	Current utilization (%)
ds_pr2	2	pool_ds82_103	45
ds_pr1	3	pool_ds81_4	42
ds_pr1	3	pool_ds81_5	49
ds_pr2	3	pool_ds82_104	55

When you examine the current utilization of the pools on the Pools page of the web-based GUI, you make the following decisions:

- Redistribute volume workload across the pools on tier 1.
- Redistribute volume workload across the pools on tier 2.

To help resolve the performance hot spots, you open the Balance Pools wizard.

1. Log in to the Tivoli Storage Productivity Center web-based GUI.
2. On the Dashboard, click **Storage Resources** > **Pools**.
3. To group the pools by storage system, click the **Storage System** column heading. Examine the utilization values for the pools on the **svc_payment_records** storage virtualizer. In this scenario, The tier 1 pools and the tier 2 pools require balancing because one or more of the pools have high utilization values.
4. Select the pools on tier 1 and tier 2, right-click, and then select **Balance Pools**.
5. Type `80` as the maximum utilization percentage for the target pools.
6. To specify the number of days for estimating the utilization of the pools, type `5`.
7. Click **Analyze**.

The Balance Analysis task opens.

The **Analysis** tab shows the following information:

- The names of the pools on tier 1 and tier 2 that were analyzed
- The name of the storage virtualizer which is **svc_payment_records**
- The current utilization percentage of the pools
- The projected utilization percentage of the pools after the volumes are moved to the recommended destination pools
- The amount of pool space that is available
- The projected amount of space that is available after the volumes are moved to the recommended destination pools

The **Recommendations** tab shows the following information:

- The names of the pools that contain the volumes that should be relocated
- The names of the destination pools
- The tier level of the pools that contain the volumes that should be relocated
- The tier levels of the destination pools
- The name of the server to which the volumes are assigned

Click **Execute** to implement the recommendations. To create a schedule, click **Schedule** > **Analysis**.

Re-tiering volumes to balance pools

In this scenario, you must re-tier some of the volumes to help resolve the performance hot spots that you identified. To re-tier the volumes, you run the Analyze Tiering wizard and select the default or automatic tiering policy.

The pools that you want to analyze are on a SAN Volume Controller with two DS8000 back-end storage systems. The name of the SAN Volume Controller is `svc_sales_records`, and the names of the two DS8000 back-end storage systems are `ds81` and `ds82`.

The value for the maximum utilization of the pools is used as a threshold value to identify whether the pools on a storage tier require balancing. The maximum utilization threshold is also used to ensure that the performance capability of destination pools does not become degraded when volumes are added to the pool. For example, recommendations are not generated to move volumes to target pools if the threshold for the maximum utilization of the pool is violated.

In this scenario, if the utilization value of a pool exceeds 80%, the performance of the pool deteriorates and it causes delays in the response times for sales transactions. When you run the Analyze Tiering wizard, you set the maximum utilization value for the pools.

The information in Table 32 is provided to illustrate which pools require the re-tiering of volumes to balance the pools.

Table 32. Utilization of the pools on the `svc_sales_records` storage virtualizer

Back-end storage system	Tier	Pool	Utilization (%)
ds81	1	pool_ds81_1	86%
ds82	1	pool_ds82_101	10%
ds81	2	pool_ds81_2	90%
ds81	2	pool_ds81_3	86%
ds82	2	pool_ds82_102	95%
ds82	2	pool_ds82_103	87%
ds81	3	pool_ds81_4	60%
ds82	3	pool_ds81_5	15%
ds82	3	pool_ds82_104	50%
ds82	3	pool_ds82_105	65%

In the following scenario, the pools are balanced by moving volumes to pools on the same storage tier and by moving volumes to higher storage tiers.

- On tier 3, the utilization values of all of the pools are below the threshold value of 80%. The pools on storage tier 3 do not require balancing; therefore, no recommendations are generated.
- On tier 2, the utilization values of all of the pools exceed the threshold value of 80%. Because the volumes cannot be redistributed to pools on the same storage tier, recommendations are generated to move the most active volumes from the pools on storage tier 2 to `pool_ds82_101` on storage tier 1.
- On tier 1, the utilization value of one of the pools exceeds the threshold value of 80%. Because volumes from storage tier 2 were added to `pool_ds82_101`, the current utilization value is not accurate. However, because of the low utilization

value for pool_ds82_101, recommendations are generated to move volumes on storage tier 1 from pool_ds81_1 to pool_ds82_101.

To help resolve the performance hot spots, you open the Analyze Tiering wizard.

1. Log in to the Tivoli Storage Productivity Center web-based GUI.
2. On the Dashboard, click **Storage Resources** > **Storage Systems.**
3. Right-click **svc_sales_records**, and then select **Analyze Tiering**. Because you open the Analyze Tiering wizard at the storage system level, all of the pools on the storage virtualizer are selected for analysis.
4. Select all of the pools as target pools, and then click **Next**. The pools that you selected for analysis are now also selected as target pools.
5. On the Optimize the placement of volumes page, set the maximum utilization for the pools to 80 percent, and then click **Analyze**.

The Analyze Tiering task is shown.

The **Analysis** tab shows the following information:

- The names of the source and target pools that were analyzed
- The names of the storage virtualizer that contains the pools
- The current utilization percentage of the pools
- The projected utilization percentage of the pools after the volumes are moved to the recommended destination pools
- The amount of pool space that is available
- The projected amount of space that is available after the volumes are moved to the recommended destination pools

The **Recommendations** tab shows the following information:

- The names of the pools that contain the volumes that should be relocated
- The names of the destination pools that were selected to relocate the volumes
- The tier level of the pools that contain the volumes that should be relocated
- The tier levels of the destination pools
- The name of the server to which the volumes are assigned

Click **Execute** to implement the recommendations. To create a schedule, click **Schedule** > **Analysis**.

Related concepts:

"Tiering storage overview" on page 429
To manage storage efficiently and reduce storage costs, optimize the placement of volumes in storage virtualizer pools.

Tiering volumes by I/O density and file usage

In this scenario, you use the Analyze Tiering wizard to regulate the placement of volumes in storage pools on tier 1. To ensure that only volumes that require the performance capability of tier 1 storage, you add minimum thresholds for I/O density and file usage for each storage tier.

The pools that you want to analyze are on a SAN Volume Controller with two DS8000 back-end storage systems. The name of the SAN Volume Controller is svc_sales_records, and the names of the DS8000 back-end storage systems are ds81 and ds82. The pools on the storage virtualizer are assigned to tier 1, tier 2, and tier 3.

With the Analyze Tiering wizard, you want to analyze the placement of the volumes on tier 1. You set minimum threshold values for I/O density and file usage to re-tier the volumes that do not meet the minimum threshold criteria that you set for I/O density and file usage.

On the Optimize the placement of volumes page of the Analyze Tiering wizard, you want to set the following minimum thresholds for each tier.

Table 33. Minimum threshold values for tier 1, tier 2, and tier 3

Tier	I/O density minimum threshold level	File usage minimum threshold level
1	4 operations per second per gigabyte of storage	Eighty percent of files were last accessed within 1 day.
2	3 operations per second per gigabyte of storage	Eighty percent of files were last accessed within 1 week.
3	1 operation per second per gigabyte of storage	Eighty percent of files were last accessed within 1 month.

The information in Table 34 is provided to illustrate the tiering of volumes.

Table 34. Current® I/O density and file usage data for the volumes in tier 1 pools

Pool	Volumes	Current I/O density	Current file usage
pool_ds81_1	volume_1	4 operations per second per gigabyte of storage	Ninety percent of the files were last accessed within one day.
pool_ds81_1	volume_2	2 operations per second per gigabyte of storage	Eighty percent of files were last accessed within 3 weeks.
pool_ds82_101	volume_11	2 operations per second per gigabyte of storage	Eighty percent of files were last accessed within 2 weeks.
pool_ds82_101	volume_12	3 operation per second per gigabyte of storage	Eighty-five percent of files were last accessed within 1 week
pool_ds82_101	volume_13	5 operations per second per gigabyte of storage	Ninety-five percent of the files were last accessed within one day.

When you run the Analyze Tiering wizard, the I/O density and file usage data of the volumes on tier 1 is compared with the minimum threshold levels that you set for the storage tiers. The volumes in the tier 1 pools with values lower than the minimum threshold levels are selected as candidate volumes for re-tiering. The criteria that determines the destination pool for the candidate volumes are as follows:

- The minimum threshold levels that you set for tier 2 and tier 3 storage on the Optimize the placement of volumes page of the Analyze Tiering wizard.
- The maximum utilization percentage that you set for destination pools on the Optimize the placement of volumes page of the Analyze Tiering wizard.

In this scenario, the following recommendations are generated:

- Move volume_12 from tier 1 to a pool on tier 2.
- Move volume_2 and volume_11 from tier 1 to a pool on tier 3.

Because volume_1 and volume_13 meet the minimum threshold levels that are set for I/O density and file usage, no recommendations are generated to move the volumes.

1. Log in to the Tivoli Storage Productivity Center web-based GUI.
2. On the Dashboard, click **Storage Resources** > **Storage Systems.**
3. Double-click **svc_sales_records**.
4. In the Internal Resources section of the navigation pane, double-click **Pools**, and then click the **Tier** column heading to group the pools by tier.
5. Select all of the pools on tier 1, then right-click and select **Analyze Tiering**.
6. Select all of the pools on tier 2 and tier 3 as target pools, and then click **Next**. To group the pools by tier, click the **Tier** column heading.
7. On the Optimize the placement of volumes page, set the maximum utilization for the pools to 80 percent.
8. Select **Advanced tiering policy**.
9. To specify the minimum threshold levels for placing volumes on tier 1, complete the following steps:
   a. Select **I/O density per second per gigabyte**, and type 4.
   b. Select **Percentage of files**, type 80 for the percentage of files, and ensure that the period is set to 1 day.
10. To specify the minimum thresholds for placing volumes on tier 2, complete the following steps:
    a. Select **I/O density per second per gigabyte**, and type 3.
    b. Select **Percentage of files**, type 80 for the percentage of files, and ensure that the period is set to 1 week.
11. To specify the minimum thresholds for placing volumes on tier 3, complete the following steps:
    a. Select **I/O density per second per gigabyte**, and type 1.
    b. Select **Percentage of files**, type 80 for the percentage of files, and ensure that the period is set to 1 month.
12. Click **Analyze**.

The Analyze Tiering task is shown.

The **Analysis** tab shows the following information:

- The names of the source and target pools that were analyzed
- The names of the storage virtualizer that contains the pools
- The current utilization percentage of the pools
- The projected utilization percentage of the pools after the volumes are moved to the recommended destination pools
- The amount of pool space that is available
- The projected amount of space that is available after the volumes are moved to the recommended destination pools

The **Recommendations** tab shows the following information:

- The names of the pools that contain the volumes in tier 1 pools that should be relocated

- The names of the destination pools on tier 2 and tier 3 that were selected to relocate the volumes
- The tier level of the pools that contain the volumes that should be relocated
- The tier levels of the destination pools
- The name of the server to which the volumes are assigned

Click **Execute** to implement the recommendations. To create a schedule, click **Schedule** > **Analysis**.

Related concepts:

"Tiering storage overview" on page 429
To manage storage efficiently and reduce storage costs, optimize the placement of volumes in storage virtualizer pools.

Creating a performance threshold and viewing violations

You can monitor the performance of a storage environment by defining performance thresholds for storage systems and SAN switches. A threshold violation is detected when a performance monitor discovers that the measured value of a metric falls outside of the specified threshold range for a resource. By defining alerts that are triggered by performance thresholds, you can get timely information about performance issues in your environment.

To monitor the performance of a storage system, you must collect performance data and define alerts. To collect performance data, schedule a performance monitor. Performance monitors collect data about monitored resources and at each sample interval the data is compared to the performance thresholds that you specify.

To specify a performance threshold, define an alert that includes a boundary value or set of boundary values that represent the acceptable range of a metric. A threshold violation is detected when a performance monitor discovers that the measured value of a metric falls outside of the specified threshold range. For example, the metric Total I/O Rate Threshold requires two upper-threshold and two lower-threshold boundaries or values. If these boundaries are matched or exceeded for a resource, an alert is generated.

In this scenario, you are defining a performance threshold alert to monitor a SAN Volume Controller named **SVC-01-IBM**. The metric for the threshold is total I/O rate, which measures the average number of I/O operations per second for both sequential and nonsequential read and write operations. After data is collected and the threshold is checked, you are viewing the threshold violation and exporting the performance information to a CSV file.

Complete the following tasks for this scenario:

- In the web-based GUI, schedule data collection for a storage system.
- In the stand-alone GUI, define a performance threshold for a storage system.
- In the web-based GUI, view the violation of the performance threshold.

Scheduling data collection

Before you can check for performance issues on a resource, you must set up data collection. Typically, data collection is automatically scheduled when resources are added to Tivoli Storage Productivity Center. To ensure that data collection was run, complete the following steps:

1. Log in to the web-based GUI.
2. In the navigation pane, click **Storage Resources** > **Storage Systems**.
3. On the Storage Systems page, locate the row for the storage system.
4. Check the values in the Probe Status and Performance Monitor Status columns to ensure that data collection was successful.
5. If data collection was not run, complete the following steps:
   a. To run a probe, right-click the storage system and select **Data Collection** > **Start Probe**.
   b. To run a performance monitor, right-click the storage system and select **Data Collection** > **Start Performance Monitor**.

Defining a performance threshold

Specify performance boundaries for the Total I/O Rate threshold condition in an alert. To define an alert, complete the following steps:

1. Log in to the stand-alone GUI.
2. In the navigation pane, expand **Disk Manager** > **Alerting**.
3. Right-click **Storage Subsystem Alerts** and select **Create Storage Subsystem Alert**.
4. Click the **Alert** tab.
5. Select **Total I/O Rate Threshold** from the **Condition** list. Fields are displayed for Critical Stress, Warning Stress, Warning Idle, and Critical Idle.
6. Define the performance boundaries for the threshold:
   a. In the **Critical Stress** field, type **3000** as the upper-threshold boundary for critical stress.
   b. In the **Warning Stress** field, type **2200** as the upper-threshold boundary for warning stress.

   Tip: The threshold values given in this scenario are examples and are not recommended values. They are for illustration only. Performance behavior varies tremendously between different resources and even different applications for different resources. It is recommended that you monitor the performance of a resource for several weeks and use this historical data to determine reasonable values for each threshold.
7. In the **Alert Suppression** section, select **Trigger alerts for both warning and critical conditions**. When you select this option, an alert is automatically triggered when the threshold is violated.
8. In the **Triggered-Actions** area, select **Email** to send an email to a mail address you specify. By default, alerts are always recorded to the error log. In this scenario, you also want to receive an email when the alert is generated.
   a. Click **Add** and type an email address in the text box. You can type multiple email addresses by clicking **Add** after each address that you enter.
   b. Click **Edit e-mail** to customize the email message. Errors that are sent to the error log are generic, so you might want to be more descriptive in the text of the email.
9. Click the **Storage Subsystems** tab.
10. Select **Storage Subsystems** > **SVC-01-IBM** in the **Available** column and move it to the **Current Selections** column.
11. Save the threshold alert:
    a. In the menu bar, click **File** > **Save**.

b. In the Save As window, type **SS Alert** and click **OK**. The entry **TPCUser.SS Alert** for the threshold alert is added to the navigation tree under the **Disk Manager** > **Alerting** > **Storage Subsystem Alerts** node. You can edit or rename the alert at any time.

Viewing threshold violations

A threshold violation is detected when a performance monitor discovers that the measured value of a metric falls outside of the specified threshold range for a resource. In the web-based GUI, you can view a list of all the threshold violations that are detected during data collection.

In this scenario, the total I/O rate for an internal resource exceeded 3000 for SVC-01-IBM and a threshold alert was generated. You want to open the alert for the threshold violation and view the performance of the resource that caused the violation.

1. In the navigation pane of the web-based GUI, go to **Storage Resources** > **Storage Systems**.
2. To view the alerts that were generated for storage systems, click the **Alerts** tab.
3. In the list of alerts, locate the alert for SVC-01-IBM. The name of the resource where a violation occurred is shown in the Resource column.
4. Right-click the alert and select **View Performance**. A chart shows the historical performance of SVC-01-IBM. The chart includes lines that represent the critical boundary and warning boundaries that were defined for the performance threshold. The default time range of the chart spans two hours before and two hours after the violation occurred.

   In this case, the total I/O rate for several volumes on SVC-01-IBM contributed to the threshold violation. A list of the volumes that were estimated to be the top contributors to the workload are shown below the chart.
5. Optional: To access the full performance view for SVC-01-IBM, click the open performance view icon on the chart. In the full performance view, you can select different metrics for the storage system, change the time range of the chart, view internal resources, and other tasks.
6. Optional: To export the performance information to a CSV file, click the export icon on the chart.

Tip: You can use the preceding steps, with some minor changes, to create threshold alerts for certain types of monitored switches.

Related tasks:

"Creating performance monitors" on page 300
Performance monitors are data collection jobs that collect metrics for measuring the performance of storage systems and switches. You can create performance monitors in the web-based GUI when you add resources for monitoring or at a later time.

"Defining performance thresholds" on page 397
Define performance thresholds to detect if the performance of the internal resources within a storage system or switch falls outside a specified range and might represent a potential problem. You can specify to be automatically alerted when thresholds are violated.

"Viewing performance information" on page 384
View the performance of storage systems, switches, and their internal resources. You can also view the performance of resources that violated a performance threshold and generated an alert.

Related reference:

"How performance information is displayed" on page 388
The performance view is displayed when you view the performance of resources, which includes resources in threshold violations, alerts, and performance monitors. Information in the performance view is organized into two main sections: a chart or table and a legend.

Archiving and deleting files

In this scenario, you are shown how to increase storage space by archiving and deleting files that meet the criteria, such as file size and file usage, that you set.

To increase storage space, you decide to archive and then delete files that contain data that is not often accessed. For the purpose of this scenario, the name of the server where the files that you want to archive and delete are stored is server1.

You log in to Tivoli Storage Productivity Center as `TPCUser`, and then you set the following criteria to identify the files that you want to archive and delete:

- The files is must be located in the `C:\` directory of the server1 server.
- The size of the files must be larger than 5 MB.
- The files must not have been accessed for 365 days or more.

To archive and delete the files, complete the following tasks in the Tivoli Storage Productivity Center web-based GUI:

- "Checking the status, date, and time of the last probe"
- "Starting a probe"

To archive and delete the files, complete the following tasks in the Tivoli Storage Productivity Center stand-alone GUI:

- "Defining the criteria to identify the files to be archived and deleted" on page 104
- "Scanning the server to find the files that meet the criteria" on page 104
- "Creating a report that shows the files that meet the criteria" on page 105
- "Archiving and deleting the files that meet the criteria" on page 105

Before you use the IBM Tivoli Storage Manager archive or backup function, you must ensure that the prerequisites listed in "Requesting an archive or backup" on page 538 are met.

Checking the status, date, and time of the last probe

1. Log in to the Tivoli Storage Productivity Center web-based GUI.
2. On the Dashboard, click **Sever Resources** > **Servers**.
3. Right-click the **server1** server, and then click **Start Probe** > **Data Collection**.
4. View the value in the **Probe Status** column for the **server1** server to determine if a probe was run

Starting a probe

1. On the Dashboard, click **Sever Resources** > **Servers**.
2. Right-click the **server1** server, and then click **Start Probe** > **Data Collection**.

Defining the criteria to identify the files to be archived and deleted

1. Log in to theTivoli Storage Productivity Center stand-alone GUI.
2. Expand **Data Manager** > **Policy Management**.
3. Right-click **Constraints**, and then click **Create Constraint**.
4. In the **Description** field, type a description of the constraint.
5. In the Available window of the **Filesystems** tab, expand **Filesystems** > **server1** > **C:**.
6. Move the selected items from the Available window to the Current Selections window.
7. On the **File Types** tab, ensure that the **Forbid files matching these patterns** check box is selected. The criteria that you set is applied to all types of file.
8. On the **Users** tab, ensure that **Forbid files owned by selected users** is selected. You want to include all of the files regardless of ownership.
9. On the **Options** tab, ensure that the **Max number of file names to be saved per agent is set to**`200`.
10. In the Alerts on Files section, select **number of days since last access greater than**, and then type `365`.
11. Select **bigger than**, type `5`, and then select **MB** from the list. In the File Filter Text window, the criteria that you selected are shown. In this example, the following text is shown: `(LAST ACCESSED earlier than 365 days 00:00 OR SIZE > 5MB)`.
12. In the Triggering Conditions section of the **Alert** tab, ensure that the triggering condition is set to **Violating files consume more than**, type `5`, and then select **Megabytes** from the list. You do not need to select an action in the Triggered Actions section. By default, if the criteria that you set are met, a message is sent to the error log.
13. Click **Save**, and type `Archive File Constraint`.

The constraint is saved in **Data Manager** > **Policy Management** > **Constraints** as **TPCUser.Archive File Constraint**.

Scanning the server to find the files that meet the criteria

1. Expand **Data Manager** > **Monitoring**.
2. Right-click **Scans**, and then select **Create Scan**.
3. On the **Filesystems** tab, select **Filesystems** > **server1** > **C:/**.
4. Move the selected item to the Current Selections window.
5. Click the **Profiles** tab.
6. Select all of the profiles that are listed in Available Profiles, and then move the profiles to **Profiles to Apply to Filesystems** and **Profiles to Apply to Directories**.
7. On the **When to Run** tab, ensure that **Run Now** and **Local Time in each time zone** are selected.
8. On the **Alert** tab, ensure that **Scan Failed** is selected as the **Triggering Condition**. You do not need to select an action in the Triggered Actions section. By default, if the criteria that you set are met, a message is sent to the error log.
9. Click the **Save** icon, and then type `Archive Scan`.

The scan is saved in **Data Manager** > **Monitoring** > **Scans** as **TPCUser.Archive Scan**.

The scan starts when it is saved. The length of time that it takes to complete the scan depends on the number of resources that are being scanned and the profiles that are used for the scan.

Creating a report that shows the files that meet the criteria

1. Expand **Data Manager** > **Reporting** > **Usage Violations** > **Constraint Violations** and click **By Constraint**. The **Constraint Violations: By Constraint** window opens in the **Selection** tab.
2. In the **Report Filter Specifications** area, type `100` in the **Return maximum of** field.
3. In the **Report Filter Specifications** area, click **Selection**. The **Select Resources** window opens.
4. Select **Archive File Constraint**. Ensure that **Archive File Constraint** is the only constraint that is selected.
5. Click **OK**.
6. Select the columns that you want to include in the report. By default, all of the columns are listed in the **Included Columns** window.
7. To remove columns from the report, select the columns, and then move them to the **Available Columns** window.
8. Click **Generate Report**. The Constraint Violation Report is shown.
9. On the **Constraints** tab, open the **Constraint Violations: By Constraint** report.
10. Click the magnifying glass to open the **Constraint Violations: By Filesystem** report.
11. Click the magnifying glass again to open the **Largest Files: By Filesystem** report. This report, the `Constraint Violation Report` is used to create a job to archive and back up the files.

Archiving and deleting the files that meet the criteria

1. Select the files in the **Constraint Violation Report** that you want to archive or back up.
2. Right-click on the selected files, and then select **Create a new archive\backup job**. The **Create New Archive/Backup Job** window opens.
3. On the **Options** tab, click **Archive**, and then click **Delete After Successful Archival**.

   Attention:
   The files that you selected to archive and delete are shown on the **Files** tab.
4. On the **When to Run** tab, click **Run Now**.
5. In the **How to handle time zones** section, click **Local time in each time zone**.
6. in the **Triggering-Condition** section of the **Alert** tab, ensure that **File Archive/Backup Failed** is selected. You do not need to select an action in the Triggered Actions section. If the criteria that you set are met, a message is sent to the error log.
7. Click the **Save** icon, type `Archive Log`, and then click **OK**. The job log is saved in **Data Manager** > **Monitoring** > **Policy Management** > **Archive/Backup** as **TPCUser.Archive Log**.

8. To close the notification that confirms that the job was submitted, click **OK**. To show the status of the job, click the job name. When the job finishes, click the date and time to view the job log file.

Related concepts:

"Requesting an archive or backup" on page 538
Request anIBM Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint. Tivoli Storage Manager protects your data from hardware failures and other errors by storing backup and archive copies of data on auxiliary storage.

Related tasks:

"Creating probes" on page 297
You can create and schedule probes to gather information about different storage resources at regular intervals. *Probes* are data collection jobs that provide detailed asset and status information about your storage resources.

"Creating a file constraint" on page 533
Learn how to create a constraint that defines the acceptable and unacceptable file types, file sizes, and file owners for a computer or a set of computers in your environment.

"Defining an archive or backup in a constraint" on page 539
Learn how to request an IBM Tivoli Storage Manager archive and backup of the largest violating files that are identified by a constraint. Request a Tivoli Storage Manager archive and backup on the alert page of a constraint definition.

"Creating scans" on page 307
Create and schedule a scan to gather usage information about your storage resources. To create a scan, use the stand-alone GUI. You can specify which file systems, computers, clusters, directories, databases, and other storage resources that you want to scan. You can define any number of scans to gather information about different storage resources at different times.

Limiting a set of storage to provision in the vSphere Web Client extension

In this scenario, a storage administrator limits the storage that a vSphere user can provision.

You plan to dedicate two storage pools to be used with vSphere. You want to allow a vSphere user to provision storage by using the vSphere Web Client extension for Tivoli Storage Productivity Center. You want to define the two storage pools as the resources that the user can use.

1. Create a capacity pool that the vSphere user can provision from. Place the two storage pools that you designated as provisioning targets for vSphere in the capacity pool. To add storage pools to the capacity pool, follow the instructions in "Adding resources to a capacity pool" on page 415.
2. Create a service class and specify that the service class can be used to provision storage only from the capacity pool. To create the service class, follow the instructions in "Creating a service class" on page 411.
3. Create a user ID for the vSphere user and assign the External Application role to that ID. To assign the role to the user, follow the instructions that are provided in the Tivoli Storage Productivity Center Information Center (http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp). Search for *assigning a role*.
4. Associate the user with the service class. To associate the user with the service class, follow the instructions in "Modifying a service class" on page 413.

When you complete these tasks, a vSphere user can provision storage by using the vSphere Web Client extension, but they are limited to the service class that was created. Because the service class is associated with the capacity pool that contains your designated vSphere storage pools, you limited where the vSphere user can provision data.

Related concepts:

"Cloud configuration" on page 65
To take advantage of the simplified provisioning and optimization capabilities that are available by using the web-based GUI, some configuration is required.

Related tasks:

"Provisioning block storage by using the vSphere Web Client extension for Tivoli Storage Productivity Center" on page 458
Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision block storage on Tivoli Storage Productivity Center. You can also select to have a data store created on the storage volume by the Block Storage Provisioning wizard.

"Provisioning file storage by using the vSphere Web Client extension for Tivoli Storage Productivity Center" on page 461
Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision file storage on Tivoli Storage Productivity Center.

Creating a profile with a filter

This scenario scans storage on a selected drive for the 100 largest media files, and summarizes space that is used by file system or directory and by owner.

There are a number of media files on computer1 that you suspect are taking up space that can be used for other purposes. You would like to search the C drive file systems and directories for the 100 largest files and create a report. You can use this report to decide what action to take (contact the owner, archive, and delete the files, and so on). You can also publish this report to the web, where other people with permission levels greater to or equal to yours (TPCUser) can access it.

With this scenario, you perform the following tasks.

1. Create a profile to look for the largest media files
2. Create a filter within the profile
3. Create a scan schedule that uses the profile
4. Create and run a batch report. Use the information that is gathered from the scan schedule
5. Save this report as an HTML file

For this scenario, complete the following steps.

1. Create a profile
   a. Expand **Data Manager** > **Monitoring**, right-click **Profiles** and select **Create Profile**.
   b. The **Create Profiles** window opens. Enter a description for the profile.
   c. In the **Statistics** tab, do the following steps.
      1) Under **Summarize space usage by**, select **filesystem or directory** and **owner**.
      2) Under **Accumulate History**, select **Per Scan** and type **10** as the number of days to keep.
      3) Under **Gather statistics by length of time since**, select **creation**.

4) Under **Gather the information on the**, select **largest files** and type `100` in the input field.

d. In the **File Filters** tab, create the filter by doing the following steps.

1) Highlight **(All Files Selected)** and click **New Condition**. The **Create Condition** window opens.

2) Select **NAME** and **Matches any of** from the pull-down lists.

3) Select **In any Directory**.

4) In the **match** field, enter the types of files you are looking for in either a comma-separated list or by pressing **Return** after each entry. As you press **Return**, the file type is listed in the box underneath the text field. For this scenario, the file types are `*.avi`, `*.au`, `*.mpeg3`, and `*.gif`. Click **OK**. The **Create Profile** window opens again, reflecting your choices.

e. Click the **Save** icon. Name the profile **Media Files Profile** and click **OK**. The profile is saved as **Data Manager** > **Monitoring** > **Profiles** > **TPCUser.Media Files Profile**.

2. Create the scan.

a. Expand **Data Manager** > **Monitoring**, right-click **Scans** and select **Create Scan**.

b. The **Create Scans** window opens. Type a description of the scan.

c. In the **Filesystems** tab, select **C:/** as the file system to scan.

d. Do not select anything in the **Directory Groups** tab.

e. In the **Profiles** tab, select **TPCUser.Media Files Profile** from Available Profiles and move it to **Profiles to apply to Filesystems**.

f. In the **When to run** tab, choose **Run Now** and select **Local time in each time zone**.

g. In the **Alerts** tab, the only **Triggering Condition** is **Scan Fails**. In the **Triggered Actions** area, select nothing to accept the default. If the job fails, a notification is sent to the error log automatically.

h. Click the **Save** icon. Name the scan **Media Files Scan** and click **OK**. The scan is saved as **Data Manager** > **Monitoring** > **Scans** > **TPCUser.Media Files Scan** and the job begins running immediately. This scan can take approximately 20 minutes to complete.

3. Create the batch schedule.

When you run a scan with the new profile, the data is in the database. You must go get it with a custom-designed report or batch schedule.

a. Expand **IBM Tivoli Storage Productivity Center** > **My Reports**, right-click **Batch Reports** and select **Create Batch Report**.

b. The **Create Batch Report** window opens. Enter a description for the report.

c. In the **Available** column in the **Report** tab, select **Usage** > **Files** > **Largest Files** > **by Computer** and move that to **Current Selections**.

d. In the **Selection** tab, select **TPCUser.Media Files Profile** from the available profiles. By default, all columns are included in **Included Columns**. If you do not want a column in your report, select it and move it to **Available Columns**.

e. In the **Options** tab, make the following selections.

- Under **Agent Computer Specification**, select **computer1**. The report is run on this computer and the data is stored. In most cases, this computer is the Agent running on the Tivoli Storage Productivity Center server. If **computer1** is a V4.2 Storage Resource Agent or the Tivoli Storage

Productivity Center server, you can accept the default path or you can specify a different path in the **Report destination** field. The default destination path varies depending on whether the computer is a Storage Resource Agent or the server.

- Under **Report Type Specification**, select **HTML File**.
- Select nothing under **Script**.
- Under **Output File Specification**, enter the mask that determines the output file name. It is recommended that you name the report something other than the default so that it can be easily identified by you and others. **{Report name}.{Report run date}** is recommended.

f. In the **When to run** tab, choose **Run Now** and select **Local time in each time zone**.

g. In the **Alerts** tab, the only **Triggering Condition** is **Report Failed**. In the **Triggered Actions** area, select nothing to accept the default. If the job fails, a notification is sent to the error log automatically.

h. Click the **Save** icon. Name the batch report **Largest Media Files** and click **OK**. The batch report criteria is saved under **IBM Tivoli Storage Productivity Center** > **Reports** > **Batch Reports** > **TPCUser.Largest Media Files**. The job begins running immediately. The job number, date, and time are listed under **TPCUser.Largest Media Files**.

4. Batch reports can be saved to any computer that runs Tivoli Storage Productivity Center. If the computer is a V4.2 Storage Resource Agent or the Tivoli Storage Productivity Center server, you can accept the default path. You can also specify a different path when you select batch report options.

Related tasks:

"Creating a profile" on page 317
Create a profile for use within a scan.

"Creating scans" on page 307
Create and schedule a scan to gather usage information about your storage resources. To create a scan, use the stand-alone GUI. You can specify which file systems, computers, clusters, directories, databases, and other storage resources that you want to scan. You can define any number of scans to gather information about different storage resources at different times.

"Creating batch reports" on page 1562
Create a batch report.

"Selecting batch report options" on page 1563
Define the agent computer that you want to run the batch report on and specify the format of the file to be saved.

Replication scenarios

These scenarios illustrate how to perform specific replication tasks using IBM Tivoli Storage Productivity Center for Replication.

The provided scenarios set out situations that can be solved using IBM Tivoli Storage Productivity Center. When using these scenarios you should remember the following:

- The scenarios are intended to be step-by-step instructions that you follow to complete a specific task. Each scenario contains multiple steps that must be followed in the order given.

- Specific values such as user names, passwords, and IP Address are given in the scenarios only for illustration purposes. You must replace them with values appropriate for your environment.
- The scenarios do not provide in-depth conceptual information about the tasks.

Performing planned maintenance scenarios

Scenarios for a planned outage consist of the removal of a unit from service to perform scheduled work on specific components for a predetermined duration.

Planned outage of local (H1), moving production to site 2 (H2) and optional return

In this scenario, your goal is to shut down the H1 box for maintenance, or for a planned power outage. As a result, you must move production to the H2 site, with the option of switching back to the H1 site after it is restored.

Follow these steps to carry out this scenario.

1. **Start H1->H2->H3**

   Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. The session will reach a Prepared state so there is a consistent point on Sites 2 and 3.

   Note: Since the outage is a planned outage, stop production I/O to H1.
2. **Suspend**

   This command creates the planned outage by suspending the Metro Mirror relationship between H1 and H2. When this step is complete, the session state should be suspended and recoverable.
3. **Recover H2**

   This command makes H2 target-available. When the session state is Target Available, the production I/O can then be moved to write to the H2 volumes.
4. **Start H2->H3**

   This command starts a Global Mirror session during the outage of H1. Production I/O can continue on H2. After the session reaches the Prepared state, a consistent copy of the data resides on Site 3, and can be used in the event of a disaster.
5. **Start H2->H1->H3**

   Issue this command when H1 is available and ready to be brought back into the configuration. This command starts a Metro Global Mirror session with H1 and H2 roles reversed from the original configuration. Production I/O can continue on H2. After the session reaches the Prepared state, both Site 1 and Site 3 will have a consistent copy of the data.

   Note: The following steps are the steps you must follow after you decide to switch production I/O to write on Site 1 again. Stop production I/O on H2 before continuing to the next step.
6. **Suspend**

   This command suspends the Metro Mirror relationship between H2 and H1. When this step is complete, the session state should be suspended and recoverable.
7. **Recover H1**

   This command makes H1 target-available. When the session state is target-available at H1, run production I/O to H1.
8. **Start H1->H2->H3**

This command returns the session to its original configuration, copying data from Site 1, to Site 2, to Site 3.

Planned outage on site 2 (H2) and return to three-site configuration

In this scenario, there is a planned outage expected on site 2. Disaster recovery capability is still needed so the copy will be setup between site 1 and site 3. After site 2 is brought back up, the session can then be brought back up to its original configuration.

Follow these steps to carry out this scenario:

1. **Start H1->H2->H3**

   Run this command to begin Metro Global Mirror on your system when the production I/O is running on H1. After the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

   Note: Production I/O does not stop running to site 1 to perform the following steps.

2. **Start H1->H3**

   This command starts a Global Mirror session between site 1 and site 3, during the outage of H2. Production I/O can continue on H1. After the session has reached a Prepared state, a consistent copy of the data will reside on Site 3.

   Note:

   - If running DS8000 devices on H1 with the correct level of microcode, and enough time has passed since the initial start of the session, the creation of the H1 to H3 pairs will *not* require a full copy.
   - Data on Site 2 is no longer consistent at this point.
   - After Site 2 is brought back up and is ready to hold a consistent copy of the data, go to the next step to return to the full three-site configuration. Production I/O never has to stop writing to H1 in this scenario.

3. **Start H1->H2->H3**

   After the session reaches a Prepared state, site 2 and 3 will have a consistent copy of the data again.

Planned outage of local (H1) with production move to remote (H3)

In this scenario, there is a planned outage for Site 1 (H1). Site 2, however, does not have the resources to run production, so production will be moved to Site 3 during the outage at Site 1. When the outage is complete, production can then be moved back to Site 1 and resume the full three-site configuration.

Complete the following steps to carry out this scenario.

1. **Start H1->H2->H3**

   Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. Once the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

   Note: Because this is a planned outage at Site 1, stop production I/O to Site 1 (H1) before continuing to the next step.

2. **Suspend**

This command creates the planned outage by suspending the Metro Mirror relationship between H1->H2. When this step is complete, the session state should be suspended and recoverable.

3. **Recover H3**

   This command makes H3 Target Available. When the session state is Target Available at H3, production I/O can be moved to H3.

   Note: When Site 1 is ready to resume production, continue with the following steps. Production I/O can continue to be written to H3 through the following step.

4. **Start H3->H1->H2**

   Issue this command when H1 and H2 are available and ready to be brought back into the configuration. Changes from production at the remote site (H3) flow from the local site (H1) to the intermediate site (H2). The H3-H1 and H1-H2 pairs are Global Copy. The session will remain in a preparing state.

   Note: While the session is Preparing, you can decrease the amount of time that production I/O has to be stopped by waiting for progress on the H3-H1 role pair to be at or near 100%. The next step will drive H3-H1 to sync, which depending on the percentage of the copy that is complete, can take a significant amount of time. When you are ready to switch back to Site 1, stop production I/O on H3 before continuing to the next step.

5. **Suspend**

   This command suspends Global Copy relationship between H3->H1 and H1->H2. When this step is complete, the session state should be suspended and recoverable.

6. **Recover H1**

   This command makes H1 Target Available. When the session state is Target Available at H1, production can be switched back to H1.

7. **Start H1->H2->H3**

   This command will resume copy from Site 1 back to sites 2 and 3. Once the session reaches a Prepared state, there will be a consistent copy on Site 2 and Site 3.

Planned outage of remote (H3) and returning

In this scenario, the remote site H3 has a planned outage, or consistency groups from the Global Mirror session need to be temporarily paused.

Complete the following steps to carry out this scenario.

1. **Start H1->H2->H3**

   Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. After the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

   Note: Production I/O to H1 can continue during the following steps.

2. **Suspend H2->H3**

   Issue this command to suspend only the Global Mirror portion of the session. The Metro Mirror continues between H1 and H2. This causes the session state to be put into a SuspendedH2H3 state.

3. **Start H1->H2->H3**

   This command restores the original configuration by restarting the Global Mirror portion of the session.

Recovering from a disaster scenarios

Scenarios for disaster recovery (or unplanned outages) describe the steps to recover from the loss of a complete site and restart operations after a disaster or disruption.

Unplanned outage of local (H1), production move to intermediate (H2)

In this scenario, there was an unplanned outage of the local site (H1), and production must be moved to Site 2 (H2) until Site 1 can be brought back up.

Complete the following steps to carry out this scenario.

Note: To ensure consistency in the event of a rolling disaster where the IBM Tivoli Storage Productivity Center for Replication server is lost on Site 1 before freezing all the copies in the session, enable the Metro Mirror heartbeat function.

1. **Start H1->H2->H3**

   Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. Once the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

2. Unplanned outage of H1

   This event causes an unplanned outage of H1, which causes the session to go suspended. It is recoverable if the session was in a Prepared state when the outage occurred. If the state was Preparing when the outage occurred, the session might be recoverable at Site 2 but will not be recoverable at Site 3.

3. **RecoverH2**

   This command makes H2 Target Available. When the session state is Target Available at H2, production can be switched to H2.

4. **Start H2->H3**

   This command starts a Global Mirror session from Site 2 to Site 3 during the outage of H1. Once the session reaches a Prepared state, consistent data should reside on Site 3.

   Note: Do not perform the following step until Site 1 is back up and ready to receive a copy of the data that has been written to Site 2 during the outage.

5. **Start H2->H1->H3**

   This command starts an MGM session with the H1 and H2 roles reversed from the original configuration. Production I/O can continue on H2. Wait for the session to go to the Prepared state, at which point there will be a consistent copy of the data on Site 1 and Site 3.

   Note: The following steps can be used to return to the original configuration so that production I/O can be switched back to Site 1. Stop production I/O on H2 before continuing to the next step.

6. **Suspend**

   This command suspends the Metro Mirror relationship between H2 and H1. When this step is complete, the session state should be suspended and recoverable.

7. **RecoverH1**

   This command makes H1 Target Available. When the session state is Target Available at H1, production I/O can be switched back to H1.

8. **Start H1->H2->H3**

This command restores the original configuration. Once the session reaches a Prepared state, a consistent copy of the data will reside on Site 2 and Site 3.

Unplanned outage of Site 2 (H2)

In this scenario, an unplanned outage of Site 2 (H2) occurs. Disaster Recovery capability is still required, so a Global Mirror session must be set up between Site 1 (H1) and Site 3 (H3) while Site 2 is down.

Complete the following steps to carry out this scenario.

1. **Start H1->H2->H3**

   Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. Once the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data

   Note: Production I/O to H1 can continue during the following steps.
2. Unplanned outage of H2

   This event causes an unplanned outage of H2, which causes the session to go suspended. It is recoverable if the session was in a Prepared state when the outage occurred. If the session state was Preparing, the session might be recoverable at Site 2 but would not be recoverable at Site 3.
3. **Start H1->H3**

   This command starts a Global Mirror session between Site 1 and Site 3 during the outage of H2. Production I/O can continue on H1.

   Note: If you are running DS8000 devices on H1 with the proper level of microcode, and enough time has passed since the initial start of the session, the creation of the H1 to H3 pairs will NOT require a full copy.

   Once the session has reached a Prepared state, a consistent copy of the data will reside on Site 3.

   Note: Data on Site 2 is no longer consistent at this point.

   Note: After Site 2 has been successfully brought back up, perform the next step to return to the three-site configuration.
4. **Start H1->H2->H3**

   Issue this command when H2 is ready to be brought back into the configuration. This command restores the original configuration. Production I/O can continue on H1. Once the session reaches a Prepared state, a consistent copy will reside on Site 2 and Site 3.

Unplanned outage of remote (H3)

In this scenario, an unplanned outage occurs on Site 3 (H3).

Complete the follow these steps to carry out this scenario.

1. **Start H1->H2->H3**

   Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. Once the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

   A failure of the remote site (H3) causes the session state to go from Prepared to SuspendedH2H3, with an error message that the connection to the H3 box has been lost.

Note: When the remote site (H3) is repaired and ready to return to the configuration, the session will remain SuspendedH2H3 and recoverable until step 2 is completed.

2. **Start H1->H2->H3**

   Issue this command only when the remote site (H3) is repaired and ready to return to the configuration. With this command, your session returns to the original configuration by starting the Global Mirror portion back up.

Unplanned outage of local (H1) with production move to remote (H3)

In this scenario, there is an unplanned outage for Site 1 (H1). Site 2, however, does not have the resources to run production, so production will be moved to Site 3 during the outage at Site 1. When the outage is complete, production can then be moved back to Site 1 and resume the full three-site configuration.

Complete the following steps to carry out this scenario.

Note: To ensure consistency in the event of a rolling disaster where the IBM Tivoli Storage Productivity Center for Replication server might be lost on Site 1 before freezing all the copies in the session, enable the Metro Mirror heartbeat function.

1. **Start H1->H2->H3**

   Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. When the session is in a Prepared state, both sites 2 and 3 have a consistent copy of the data.

2. **Unplanned outage of H1 and H2**

   This event causes an unplanned outage of H1 and H2 which causes the session to go suspended. It is recoverable if the session was in a Prepared state when the outage occurred. If the state was Preparing, the session might be recoverable at Site 2, but will not be recoverable at Site 3.

3. **Recover H3**

   This command makes H3 Target Available. The session will be severe/Target Available. When the session state is Target Available at H3, production I/O can be moved to H3.

   Note: Once Site 1 is brought back up, perform the following step to begin the process of returning production to Site 1.

4. **Start H3->H1->H2**

   Issue this command when H1 and H2 are available and ready to be brought back into the configuration. Changes from production at the remote site (H3) flow from the local site (H1) to the intermediate site (H2). The H3-H1 and H1-H2 pairs are Global Copy. The session will in the preparing state.

   Note: While the session is Preparing, you can decrease the amount of time that production I/O has to be stopped by waiting for the progress on the H3-H1 role pair to be at or near 100%. The next step will drive H3-H1 to sync, which depending on the percentage of the copy that is complete, can take a significant amount of time. When you are ready to switch back to Site 1, stop production I/O on H3 before continuing to next step.

5. **Suspend**

   This command suspends the Global Copy relationship between H3>H1 and H1>H2. When this step is complete, the session state should be suspended and recoverable.

6. **Recover H1**

This command makes H1 Target Available. When the session state is Target Available at H1, run production I/O to H1.

7. **Start H1->H2->H3**

   This command restores the original three-site configuration. Once the session reaches a Prepared state, a consistent copy of the data will reside on Site 2 and Site 3.

Practicing disaster recovery scenarios

This section provides scenarios for how to practice disaster recovery using a regular Metro Global Mirror session. You can follow these scenarios to practice disaster recovery while continuing to be disaster recovery-capable using only the Metro Mirror portion of the session. The Metro Global Mirror with Practice session is also available, which allows for practicing disaster recovery while maintaining full three-site disaster recovery capability.

Practice a recovery to H2 with production that is continuing on H1 in a normal configuration

In this scenario, your goal is to practice a recovery to H2 with production that is continuing on H1 in a normal configuration.

Follow these steps to carry out this scenario.

1. **Start H1->H2->H3**

   Issue this command to begin a Metro Global Mirror session on your system when the production I/O is running on H1. Wait for the session to go to the prepared state.

   You can continue to run production I/O to H1 during the following steps.

2. **Suspend**

   This command creates the planned outage by suspending the Metro Mirror relationship between H1 and H2. When this step is complete, the session state should be suspended and recoverable.

   Note: Stop running production I/O to H2 prior to issuing the Suspend command.

3. **Recover H2**

   This command makes H2 target-available. When the session is target-available at H2, run practice I/O to H2.

4. **Start H1->H3**

   Issue this command to ensure that any I/O that is written to site 1 is copied to site 3.

   You can also use the Start H1->H3 command to practice for a disaster at the H2 site while still maintaining disaster recovery capability. When the H1->H3 role pair is in Global Mirror mode, you can recover the session at site 3 if a disaster occurs.

   When you complete testing on H2, you can issue the Start H1->H2->H3 command to return to a three-site configuration without completing a full copy.

5. **Start H1->H2->H3**

   Issue this command to restart the Metro Global Mirror session in the H1-H2-H3 configuration. Any changes from production at the local site (H1) flow from the remote site (H3) to the intermediate site (H2), which overwrites the test data on H2.

Table 35. Role-pair states at the completion of each step (participating and nonparticipating role pairs)

Steps	Role-pair states					
	H1-H2	**H2-H3**	**H2-J3**	**H3-J3**	**H1-H3**	**H1-J3**
1	Prepared	Preparing	Prepared	Target available	Non-participating	Non-participating
2	Suspended					
3						
4		Non-participating	Non-participating		Participating	Prepared
5	Prepared	Preparing	Prepared		Non-participating	Non-participating

Practice recovery to H1 with production that is continuing on H2

In this scenario, your goal is to practice a recovery to H1 with production that is continuing on H2.

Follow these steps to carry out this scenario.

1. **Start H1->H2->H3**

   Issue this command to begin a Metro Global Mirror session on your system when the production I/O is running on H1. Wait for the session to go to the prepared state.

2. **Suspend**

   This command creates the planned outage by suspending the Metro Mirror relationship between H1 and H2. When this step is complete, the session state should be suspended and recoverable.

   Note: Stop running production I/O to H1 prior to issuing the Suspend command.

3. **Recover H2**

   This command makes H2 target-available.

   Note: Start running production I/O to H2.

4. **Start H2->H1->H3**

   This command starts a Metro Global Mirror session with H1 and H2 roles reversed from the original configuration.

   You can continue to run production I/O to H2 for the next two steps.

5. **Suspend**

   This command creates the planned outage by suspending the Metro Mirror relationship between H1 and H2. When this step is complete, the session should be suspended and recoverable.

6. **Recover H1**

   This command makes H1 target-available. When the session state is target-available at H1, run practice I/O to H1.

7. **Start H2->H3**

   Issue this command to ensure that any I/O that is written to site 2 is copied to site 3.

   You can also use the Start H2->H3 command to practice for a disaster at the H1 site while still maintaining disaster recovery capability. When the H2->H3 role pair is in Global Mirror mode, you can recover the session at site 3 if a disaster occurs.

When you complete testing on H1, you can issue the Start H2->H1->H3 command to return to a three-site configuration without completing a full copy.

8. **Start H2->H1->H3**

   Issue this command to restart the Metro Global Mirror session in the H2-H1-H3 configuration. Any changes from production at the local site (H2) flow from the remote site (H3) to the intermediate site (H1), which overwrites the test data on H1.

Table 36. Role-pair states at the completion of each step (participating and nonparticipating role pairs)

Steps	Role-pair states					
	H1-H2	**H2-H3**	**H2-J3**	**H3-J3**	**H1-H3**	**H1-J3**
1	Prepared	Preparing	Prepared	Target available	Non-participating	Non-participating
2	Suspended					
3						
4	Prepared	Non-participating	Non-participating		Preparing	Prepared
5	Suspended					
6						
7		Preparing	Prepared		Non-participating	Non-participating
8	Prepared	Non-participating	Non-participating		Preparing	Prepared

Practice at H3 for Metro Global Mirror

This topic describes how to practice at the remote site (H3) while running a Metro Global Mirror session.

To practice on the remote site, you first need to suspend the back-end. For example, if running in an H1->H2->H3 configuration, issue a **Suspend H2H3** command to pause the Global Mirror portion of the session while leaving the Metro Mirror portion of the session active. When the session is in a suspended state (that is, H2H3 is suspended while in an H1->H2->H3 configuration), you have the option to issue a **Recover H3** command.

The **Recover H3** command acts differently than other recover commands. In this case, the command acts only to make the remote site (H3) target-available and does not switch the application site to H3. The application site remains the original host site and the session acts as it does in a normal Metro Mirror Failover/Failback session.

While the session is in this mode, it remains capable of disaster recovery at Site 2 and will freeze all pairs to maintain consistency during an error. You can also suspend the session by command and recover to Site 2 if necessary. However, if recovering to Site 2 in this mode, it will not be possible to restart a full three-site solution. To get three sites back up and running, you need to copy back to H1 using the **Start H2->H1** command and then suspend and recover to H1. Once H1 is the application host again, you can issue the **Start H1->H2->H3** command to restart three-site support. If this is done, all data written to H3 during the practice is overwritten with the data that is on H1.

Overview of scenarios for practicing at H3 for Metro Global Mirror:

This topic provides an overview of the scenarios for practicing at H3 for Metro Global Mirror.

Note: For simplicity, the following scenarios are geared toward the configuration where the session is running H1->H2->H3. The same options apply when in the configuration is running H2->H1->H3.

There are two methods to set up a practice mode at the remote site, H3.

In the first method, you issue the following commands:
1. **Start H1->H2->H3**
2. **SuspendH2H3**
3. **RecoverH3**

When the session state is Target Available at H3, you can practice you disaster recovery steps on H3 while continuing to run the real production I/O to H1.

In the second method of setting up a practice mode, issue the following commands:
1. **Start H1->H2->H3**
2. **Suspend**
3. **RecoverH3**
4. **Start H1->H2**

When the session state is Target Available at H3, run practice I/O to H3. Continue running production I/O to H1.

The difference between the second method and the first method is that a **Suspend** command is issued instead of a **SuspendH2H3** command. When the session is in suspended mode (as opposed to suspended H2H3 mode), the session does not automatically go into practice mode. The application site switches to H3 and you must issue a **Start H1->H2** to enable the practice mode.

Practice scenario 1: While practicing, disaster occurs at Site 1, recover to Site 2:

This topic describes what to do in the event that a disaster occurs on Site 1 while practicing at Site 3.

If a disaster occurs at Site 1, the session immediately suspends itself. You then must complete the following steps:

1. Issue **RecoverH2**.
2. Move the production I/O to Site 2.
3. Issue **Start H2->H1** to make the session capable of disaster recovery again once Site 1 has been brought back up.

Practice scenario 2: While practicing, planned outage at Site 1, recover to Site 2:

This topic describes what to do if there is a planned outage on Site 1 while practicing at Site 3.

While practicing, if you want a planned outage with a switch to Site 2, complete the following steps:

1. Issue **Suspend**.

2. Issue **RecoverH2**.
3. Move the production I/O to Site 2.
4. Issue **Start H2->H1** to make the session capable of disaster recovery again once Site 1 has been brought back up.

Practice scenario 3: while practicing with production at site 2, move production back to site 1:

This topic describes an H3 practice scenario, with production running at site 2, followed by a move of production to site 1.

While practicing, if you moved the production to site 2, you need to move it back to site 1 before you can start three-site support. To move production back to site 1, do the following steps (assume the session is running H2->H1):

1. Issue **Suspend**.
2. Issue **Recover H1**.
3. Move the production I/O to site 1.
4. Issue **Start H1->H2** to make the session capable of disaster recovery again and move the production I/O to H3.
5. Issue **Start H1->H2->H3** when H3 is ready to receive a copy of the data at site 1. When this command completes, all writes that were done to H3 are overwritten by the data on site 1.

The following table shows the changing states of the role pairs after each step.

Table 37. Role-pair states at the completion of each step (participating and nonparticipating role pairs)

<table>
<tr><th rowspan="2">Steps</th><th colspan="6">Role-pair states</th></tr>
<tr><th>H1-H2</th><th>H2-H3</th><th>H2-J3</th><th>H3-J3</th><th>H1-H3</th><th>H1-J3</th></tr>
<tr><td>1</td><td>Suspended</td><td colspan="5" rowspan="2">Nonparticipating</td></tr>
<tr><td>2</td><td>Target Available</td></tr>
<tr><td>4</td><td rowspan="2">Prepared</td><td colspan="3">target-available</td><td colspan="2" rowspan="2">nonparticipating</td></tr>
<tr><td>5</td><td>Preparing</td><td>Prepared</td><td>Target Available</td></tr>
</table>

Chapter 4. Administering

Administer Tivoli Storage Productivity Center to authorize users, start, and use the console, which is the graphical user interface (GUI), start and stop services, and many other administrative tasks. Administer DB2 by using the DB2 command-line interface or IBM Data Studio.

Starting and stopping the Tivoli Storage Productivity Center servers

This section provides information about how to start and stop the Tivoli Storage Productivity Center services.

Tip: On AIX and Linux operating systems, the Jazz for Service Management server does not start automatically. You must manually start the server.

Stopping the Tivoli Storage Productivity Center servers

You can run scripts to stop the servers on the Windows, Linux, or AIX operating systems.

Stopping the servers on Windows

Tip: The default *TPC_install_directory* is `C:\Program Files\IBM\TPC`.

To stop the servers on the Windows operating system, enter the following commands in a Windows command prompt window in the following order:

Storage Resource agent

To stop the **Storage Resource Agent** service on Windows, complete the following steps:

1. Open Windows Services.
2. Stop the **IBM Storage Resource Agent - *directory*** service where *directory* is where the Storage Resource agent is installed.

Web server

TPC_install_directory`\scripts\stopTPCWeb.bat`

Data server

TPC_install_directory`\scripts\stopTPCData.bat`

Device server

TPC_install_directory`\scripts\stopTPCDevice.bat`

Replication server

TPC_install_directory`\scripts\stopTPCReplication.bat`

Jazz for Service Management and Tivoli Common Reporting

JazzSM_install_directory`\profile\bin\stopServer.bat` *server_name* `-username` *JazzSM_username* `-password` *JazzSM_password*

JazzSM_install_directory, *server_name*, *JazzSM_username*, and *JazzSM_password* were specified when you installed Jazz for Service Management. The default *JazzSM_install_directory* is `C:\Program Files\IBM\JazzSM`, and the default *server_name* is `server1`. You can also

© Copyright IBM Corp. 2001, 2013

return the name and status of the server by entering the following WebSphere Application Server command:

```
serverStatus -All
```

Stopping the servers on Linux or AIX

Tip: The default *TPC_install_directory* is `/opt/IBM/TPC`.

To stop the servers on Linux or AIX operating systems, enter the following commands in a command prompt window in the following order:

Storage Resource Agent

/SRA_install_directory/agent/bin/agent.sh stop

Web server

/TPC_install_directory/scripts/stopTPCWeb.sh

Data server

/TPC_install_directory/scripts/stopTPCData.sh.

Device server

/TPC_install_directory/scripts/stopTPCDevice.sh

Replication server

/TPC_install_directory/scripts/stopTPCReplication.sh

Jazz for Service Management and Tivoli Common Reporting

/JazzSM_install_directory/profile/bin/stopServer.sh *server_name* -username *JazzSM_username* -password *JazzSM_password*

JazzSM_install_directory, *server_name*, *JazzSM_username*, and *JazzSM_password* were specified when you installed Jazz for Service Management. The default *JazzSM_install_directory* is `/opt/IBM/JazzSM`, and the default *server_name* is `server1`. You can also return the name and status of the server by entering the following WebSphere Application Server command:

```
serverStatus -All
```

Related tasks:

"Checking the status of the Data Server" on page 174
This topic describes how to check the status of the Data server.

"Checking the status of the Device server" on page 174
This topic describes how to check the status of the Device server.

"Checking the status of agents" on page 174
Check the status of agents.

Starting the Tivoli Storage Productivity Center servers

You can run scripts to start the servers on the Windows, Linux, or AIX operating systems.

Starting the servers on Windows

Tip: The default *TPC_install_directory* is `C:\Program Files\IBM\TPC`.

To start the servers on the Windows operating system, enter the following commands in a Windows command prompt window in the following order:

Device server

TPC_install_directory\scripts\startTPCDevice.bat

Data server
TPC_install_directory\scripts\startTPCData.bat

Replication server
TPC_install_directory\scripts\startTPCReplication.bat

Web server
TPC_install_directory\scripts\startTPCWeb.bat

Storage Resource Agent - *directory*

Tip: The **Storage Resource Agent service** is started on the Windows operating system by using Windows Services.
To start the **Storage Resource Agent** service on Windows, complete the following steps:

1. Open Windows Services.
2. Start the **IBM Storage Resource Agent -** ***directory*** service where *directory* is where the Storage Resource agent is installed.

Jazz for Service Management and Tivoli Common Reporting
JazzSM_install_directory\profile\bin\startServer.bat *server_name*

JazzSM_install_directory and *server_name* were specified when you installed Jazz for Service Management. The default *JazzSM_install_directory* is `C:\Program Files\IBM\JazzSM`, and the default *server_name* is `server1`. You can also return the name and status of the server by entering the following WebSphere Application Server command:

```
serverStatus -All
```

Important: If you have not installed Jazz for Service Management by using Tivoli Storage Productivity Center, you must manually restart Jazz for Service Management each time you restart the computer or log in to the computer. When you restart Jazz for Service Management manually, Tivoli Common Reporting is also restarted. After you restart Jazz for Service Management, it might take some time before all components, including Tivoli Common Reporting, are running.

Starting the servers on Linux or AIX

Note: The default *TPC_install_directory* is `/opt/IBM/TPC`.

To start the servers on the Linux or AIX operating systems, enter the following commands in a command prompt window in the following order:

Device server
/*TPC_install_directory*/scripts/startTPCDevice.sh

Data server
/*TPC_install_directory*/scripts/startTPCData.sh

Replication server
/*TPC_install_directory*/scripts/startTPCReplication.sh

Web server
/*TPC_install_directory*/scripts/startTPCWeb.sh

Storage Resource Agent
/*SRA_install_directory*/agent/bin/agent.sh start

Jazz for Service Management and Tivoli Common Reporting

/JazzSM_install_directory/profile/bin/startServer.sh *server_name*

JazzSM_install_directory and *server_name* were specified when you installed Jazz for Service Management. The default *JazzSM_install_directory* is /opt/IBM/JazzSM, and the default *server_name* is server1. You can also return the name and status of the server by entering the following WebSphere Application Server command:

```
serverStatus -All
```

Important: If you restart your computer, you must manually restart Jazz for Service Management. When you restart Jazz for Service Management manually, Tivoli Common Reporting is also restarted. After you restart Jazz for Service Management, it might take some time before all components, including Tivoli Common Reporting, are running.

Related tasks:

"Checking the status of the Data Server" on page 174
This topic describes how to check the status of the Data server.

"Checking the status of the Device server" on page 174
This topic describes how to check the status of the Device server.

"Checking the status of agents" on page 174
Check the status of agents.

Default locations of log files

Check log files to view detailed information about Tivoli Storage Productivity Center processing and to troubleshoot problems.

The following list shows the default log file locations for Tivoli Storage Productivity Center and other components.

Device server:

The Liberty Profile logs are in the following directories:

Windows operating systems

TPC_installation_directory\wlp\usr\servers\deviceServer\logs

Linux or AIX operating systems:

TPC_installation_directory/wlp/usr/servers/deviceServer/logs

The operational logs are in the following directories:

Windows operating systems:

TPC_installation_directrory\device\log

Linux or AIX operating systems:

TPC_installation_directory/device/log

Data server:

Windows operating systems:

TPC_installation_directory\data\log

Linux or AIX operating systems:

TPC_installation_directory/data/log

Web server log files:

Windows operating systems:

TPC_installation_directory\ewas\profiles\WebServerProfile\logs\webServer

Linux or AIX operating systems:
TPC_installation_directory/ewas/profiles/WebServerProfile/logs/webServer

Web-based GUI:

Windows operating systems:
TPC_INSTALL_DIR\web\log

Linux or AIX operating systems:
TPC_INSTALL_DIR/web/log

Storage Resource agents:
TPC_installation_directory/agent/log/*name_of_server_SRA_communicates_with*

Tips:

- For Windows operating systems, the default *TPC_installation_directory* is C:\Program Files\IBM\TPC.
- For Linux or AIX operating systems, the default *TPC_installation_directory* is /opt/IBM/TPC.

Overview of required user names for initial logon to the Tivoli Storage Productivity Center family of products

All the graphical user interfaces (GUIs) in the Tivoli Storage Productivity Center family require a user ID and password. If you are logging on to a GUI immediately after you installed the software, the user ID that you must use differs depending on the type of installation and the GUI interface.

Immediately after you install the software, you must use the user name as described in the following tables. After you log on, you can assign roles for users in Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication. When users log on, their user roles determine their authorization level and the components that they can view in each GUI.

Required user name for initial logon after installation on a single server when only the common user is defined

After you install Tivoli Storage Productivity Center in a single-server environment, the required user name for the initial logon is the common user name that was defined for the Tivoli Storage Productivity Center installation.

You must use this user name to log on to the stand-alone GUI, the web-based GUI, the Tivoli Storage Productivity Center for Replication web-based GUI, and the Tivoli Storage Productivity Center command-line interface.

Required user name for initial logon after installation on multiple servers with a remote database schema

After you install the software on multiple servers with a remote database schema, you must use the user name as described in the following table.

Table 38. Required user name for initial logon after installation on multiple servers with a remote database schema

To log on to:	Use this user name:
Tivoli Storage Productivity Center stand-alone GUI	The user name that was defined for the installation of Tivoli Storage Productivity Center server.
Tivoli Storage Productivity Center web-based GUI	The user name that was defined for the installation of Tivoli Storage Productivity Center server.
Tivoli Storage Productivity Center for Replication web-based GUI	The user name that was defined for the installation of Tivoli Storage Productivity Center server.

Required user name for initial logon after installation on multiple servers and Tivoli Storage Productivity Center reports are remote

After you install the software on multiple servers, and Tivoli Storage Productivity Center reports are remote, you must use the user name as described in the following table.

Table 39. Required user name for initial logon after installation on multiple servers and Tivoli Storage Productivity Center reports are remote

To log on to:	Use this user name:
Tivoli Storage Productivity Center stand-alone GUI	The user name that was defined for the installation of Tivoli Storage Productivity Center server.
Tivoli Storage Productivity Center web-based GUI	The user name that was defined for the installation of Tivoli Storage Productivity Center server.
Tivoli Storage Productivity Center for Replication web-based GUI	The user name that was defined for the installation of Tivoli Storage Productivity Center server.

Role-based authorization

Roles determine the functions that are available to users of Tivoli Storage Productivity Center. When a user ID is authenticated to Tivoli Storage Productivity Center through the GUI, CLI, or APIs, membership in an operating system or LDAP group determines the authorization level of the user.

In Version 5.2, the roles that were previously defined in Tivoli Storage Productivity Center were consolidated into a smaller set of roles. The following table shows how the roles in versions earlier than 5.2 are mapped to the current set of roles:

Table 40. How roles in previous versions of Tivoli Storage Productivity Center are mapped to the roles in version 5.2 and later

Roles in previous versions	Roles in version 5.2 and later	Authorization level
Superuser Productivity Center administrator Disk administrator Fabric administrator Data administrator Tape administrator	Administrator	This role has full access to all monitoring and administrative functions. At least one group must have the Administrator role. **Note:** When Tivoli Storage Productivity Center is first installed, the following operating system groups are assigned the Administrator role: • Windows: Administrators • UNIX and Linux: root • AIX: system
Disk operator Fabric operator Data operator Tape operator	Monitor	In the web-based GUI, this role has access to the following read-only functions: • Viewing and exporting information about monitored resources • Viewing and acknowledging alerts • Viewing tasks and data collection jobs • Viewing data paths • Opening management GUIs • Opening logs • Tiering storage In the stand-alone GUI, this role has access to the following functions: • Viewing data that is collected by Tivoli Storage Productivity Center • Creating, generating, and saving reports
This role did not exist in versions 5.1 or earlier.	External Application	This role enables users of other applications to use the provisioning capability of Tivoli Storage Productivity Center to provision storage. For example, a VMware user with this role can provision storage in the vSphere GUI by using the vSphere Web Client extension for Tivoli Storage Productivity Center. If you assign the External Application role to the user, you must also assign one or more service classes to the user. This role does not enable users to log in to the Tivoli Storage Productivity Center GUIs.

Tips:

- If a user belongs to multiple groups and the groups have different roles, the role with the highest level of authorization is granted to the user. For example, if a

user belongs to a group that is assigned the Administrator role and also belongs to a group that is assigned a Monitor role, the user is granted the authorization of the Administrator role.

- If a user is not a member of a group that is assigned a Tivoli Storage Productivity Center role, no access is granted to that user.
- For rollup reports, you need Administrator authority to do the following actions:
  - Add, remove, or modify the Tivoli Storage Productivity Center subordinate server that the master server is monitoring.
  - Create or run probe jobs that include Tivoli Storage Productivity Center subordinate servers.

  Any Tivoli Storage Productivity Center role can generate rollup reports.
- The Tivoli Storage Productivity Center installation program adds the administrator, external application, and monitor roles to the Tivoli Storage Productivity Center installation user.

Functions available in the web-based GUI based on role and license

When you use the web-based GUI, your Tivoli Storage Productivity Center role and product license determine the functions that are available.

Users who are assigned the Administrator role or the Monitor role can use the web-based GUI. The functions that are available depend on the role that is assigned to the user:

Administrator role

Users who are assigned the Administrator role have access to all monitoring and administrative functions and are limited only by license restrictions.

Monitor role

Users who are assigned the Monitor role can view information about monitored resources and other objects such as tasks, alerts, and service classes. They can acknowledge alerts and resource statuses, open logs, and open management GUIs. Users who are assigned the Monitor role do not have access to administrative functions, with the following exceptions:

- They can assign storage pools to tiers.
- They can be granted permission in a service class to provision storage by using the service class. If so, they can use the Provision Storage wizard to create a provisioning task. Users can delete provisioning tasks that they create. If the service class specifies that administrator approval is not required, the users can run or schedule the provisioning tasks that they create.

To use some functions of the web-based GUI, you must have the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license. If you have the Tivoli Storage Productivity Center license or the Tivoli Storage Productivity Center Select license, the following functions are not available:

- Storage tier optimization. The Analyze Tiers wizard is not available.
- Pool balancing. The Balance Pools wizard is not available.
- Block storage provisioning. You are not able to provision volumes by using either the Provision Storage wizard or the vSphere Web Client extension. However, there are no license restrictions for file storage provisioning, therefore, Network Attached Storage (NAS) file shares can be provisioned.

The following table outlines the functions available by role. Some functions are restricted to the Administrator role. Functions that are restricted to the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license are also shown. In addition to the restrictions listed in this table, users who are assigned the Monitor role do not have access to user management functions.

Table 41. Role and license restrictions in the web-based GUI

Resource or object	**Functions available by role**	**Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license**
Performance monitors	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about the switch or storage system that is being monitored • Open performance monitor logs • Export performance monitor information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Start or stop a performance monitor • Schedule a performance monitor	All functions are available, depending on role.
Alerts	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about alerts • Acknowledge alerts • Export alert information to a file Users who are assigned the Administrator role can also remove alerts.	All functions are available, depending on role.
Tasks	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about tasks • Open task logs • Export task information to a file If a user assigned the Monitor role has permission to provision storage by using a service class, the user can delete provisioning tasks that the user creates. If the service class does not require administrator approval, the user can also schedule and run the provision tasks that the user creates. Users who are assigned the Administrator role can also complete the following administrative actions: • Run tasks • Schedule tasks • Delete tasks	All functions are available, depending on role.

Table 41. Role and license restrictions in the web-based GUI (continued)

Resource or object	Functions available by role	Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license
Storage systems	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about storage systems • Acknowledge storage system status • Open the management GUI or element manager for a storage system • Export storage system information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Add and remove storage systems • Schedule data collection jobs • Optimize storage tiering • Add storage systems to capacity pools • Administer connections	Optimize storage tiering
Volumes	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about volumes • Acknowledge volume status • Export volume information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Transform storage volumes • Optimize storage tiering	Optimize storage tiering
Pools	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about pools • Acknowledge pool status • Export pool information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Add pools to capacity pools • Balance pools • Optimize storage tiering	• Balance pools • Optimize storage tiering
Shares	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about shares • Export pool information to a file	All functions are available, depending on role.

Table 41. Role and license restrictions in the web-based GUI (continued)

Resource or object	Functions available by role	Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license
Servers	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about servers • Acknowledge server status • View and collect logs • Export server information to a file If a user assigned the Monitor role has permission to provision storage by using a service class, the user can provision storage to servers. Users who are assigned the Administrator role can also complete the following administrative actions: • Add and remove servers • Schedule data collection jobs • Provision storage to servers • Modify Storage Resource agents	• Provision block storage. There are no license restrictions for provisioning file storage. • Optimize storage tiering.
Hypervisors	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about hypervisors • Acknowledge hypervisor status • Export hypervisor information to a file If a user who is assigned the Monitor role has permission to provision storage by using a service class, the user can provision storage to hypervisors. Users who are assigned the Administrator role can also complete the following administrative actions: • Add and remove hypervisors • Schedule data collection jobs • Provision storage to hypervisors • Administer connections	Provision block storage. There are no license restrictions for provisioning file storage.

Table 41. Role and license restrictions in the web-based GUI (continued)

Resource or object	Functions available by role	Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license
Switches	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about switches • Acknowledge switch status • Open the management GUI or element manager for a switch • Export switch information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Add and remove switches • Schedule data collection jobs • Administer connections • Enable automatic zoning	Enable automatic zoning
Fabrics	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about fabrics • Acknowledge fabric status • Export fabric information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Add and remove fabrics • Schedule data collection jobs • Administer connections • Enable automatic zoning	Enable automatic zoning
Service classes	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about service classes • Export service class information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Create service classes • Modify service classes • Delete service classes	Create block-storage service classes. There are no license restrictions for working with file-storage service classes.

Table 41. Role and license restrictions in the web-based GUI (continued)

Resource or object	Functions available by role	Functions that require the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license
Capacity pools	Users who are assigned the Monitor or Administrator role can complete the following monitoring actions: • View information about capacity pools • Export capacity pool information to a file Users who are assigned the Administrator role can also complete the following administrative actions: • Create capacity pools • Modify capacity pools • Delete capacity pools	All functions are available, depending on role.

Functions available in the stand-alone GUI based on role and license

When you use the stand-alone GUI, your Tivoli Storage Productivity Center role and product license determine the functions that are available.

Users who are assigned the Administrator role or the Monitor role can use the stand-alone GUI. The functions that are available depend on the role that is assigned to the user:

Administrator role
Users who are assigned the Administrator role have access to all monitoring and administrative functions and are limited only by license restrictions.

Monitor role
Users who are assigned the Monitor role can view data that is collected by Tivoli Storage Productivity Center, and can create, generate, and save reports.

To use some functions of the stand-alone GUI, you must have the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license. The following table outlines the functions that are restricted to the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license. The table mirrors the navigation tree of the stand-alone GUI. The table shows nodes of the navigation tree and indicates whether that function is available with all licenses, or whether the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license is required.

Table 42. Functions with each Tivoli Storage Productivity Center license

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
Administrative Services		
Services	Yes	Yes
Data Sources[1]		

Table 42. Functions with each Tivoli Storage Productivity Center license (continued)

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
• CIMOM agents • Storage Resource agents • Storage Subsystems • TPC Servers • VMWare VI Data Source	• Yes • Yes • Yes • No • Yes	• Yes • Yes • Yes • Yes • Yes
Discovery	Yes	Yes
Configuration	Yes[2]	Yes
IBM Tivoli Storage Productivity Center		
Job Management	Yes	Yes
Reporting • My Reports (User ID reports) • System Reports • Rollup Reports • Data Source Reports	Yes See Table 8 on page 31 for a list of available reports.	Yes See Table 8 on page 31 for a list of available reports.
Topology Viewer	Yes	Yes
Monitoring (data collection) • Probe Alert Configurations • TPC Server Probes	 • Yes • No	 • Yes • Yes
Storage Resource Group Management	Yes[5]	Yes
Analytics • Configuration History • Configuration Analysis	 • No • No	 • Yes • Yes
Alerting: • Authentication Configuration Alerts • Alert Log	 • Yes • Yes[3]	 • Yes • Yes
Data Manager		
Monitoring: (data collection) • Groups • Pings • Scans • Profiles	 • Yes • Yes • No • No	 • Yes • Yes • Yes • Yes
Alerting	Yes[4]	Yes
Policy Management	No	Yes

Table 42. Functions with each Tivoli Storage Productivity Center license (continued)

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
Reporting	Yes Not all reports are available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license. See Table 8 on page 31 for a list of available reports.	Yes
Data Manager for Databases	No	Yes
Data Manager for Chargeback	No	Yes
Disk Manager		
Storage Subsystem	Yes	Yes
Monitoring (data collection) • Groups • Jobs • Subsystem Performance Alert Configurations	 • Yes • Yes • Yes	 • Yes • Yes • Yes
Alerting	Yes	Yes
Reporting • Groups • Storage Subsystems • Storage Subsystem Performance	 • Yes • Yes • Yes	 • Yes • Yes • Yes
Fabric Manager		
Monitoring (data collection) • Groups • Switch Performance Alert Configurations	 • Yes • Yes	 • Yes • Yes
Fabrics	Yes	Yes
Alerting	Yes	Yes
• Switch Performance (See "Creating switch performance reports" on page 1503)	Yes	Yes
Replication Manager	Yes	Yes

Table 42. Functions with each Tivoli Storage Productivity Center license (continued)

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
Note: 1. You cannot associate TPC Servers with Tivoli Storage Productivity Center in the Tivoli Storage Productivity Center or Storage Productivity Center Select license. VMware Data Sources are associated with the Tivoli Storage Productivity Center or Storage Productivity Center Select license. TPC Servers are associated with the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license. 2. The following nodes under **Configuration** are not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: • License Keys • Quota and Constraint e-mail Address Rules • Resource History Retention for Databases • Removed Resource Retention for Databases • Configuration History Settings 3. The following nodes under **Alert Log** are not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: • Directory • User • OS User Group • Configuration Analysis 4. The following node under **Data Manager** > **Alerting** is not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: Directory Alerts. Hypervisor Alerts and Other NAS Alerts are available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license. 5. You cannot create a storage resource group for a type of capacity pool.		

Table 43. Available reports with each IBM Tivoli Storage Productivity Center license

Reports	Tivoli Storage Productivity Center or Storage Productivity Center Select	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
Batch reports	Yes[1], [2]	Yes[1]
System Reports: • Data • Fabric • Disk	 • Yes[4] • Yes • Yes	 • Yes • Yes • Yes
Rollup Reports	No	Yes
Data Source Reports	Yes	Yes
Asset reports:		

Table 43. Available reports with each IBM Tivoli Storage Productivity Center license (continued)

Reports	Tivoli Storage Productivity Center or Storage Productivity Center Select	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
• By Cluster • By Computer • By Hypervisor • By IBM Storwize V7000 Unified/IBM SONAS • By OS Type • By Storage Subsystem • System-wide	• Yes • Yes • Yes • Yes • Yes • Yes • Yes[3]	• Yes • Yes • Yes • Yes • Yes • Yes • Yes
Availability reports	Yes	Yes
TPC-wide Storage Space reports: • Disk Space • File System Space • Consumed File System Space • Available File System Space	 • Yes • Yes • Yes • Yes	 • Yes • Yes • Yes • Yes
Usage reports	No	Yes (these reports are a result of a scan)
Usage Violation reports	No	Yes (these reports are a result of a scan)
Backup reports	No	Yes (these reports are a result of a scan)
Monitored Computer Storage Space reports	Yes	Yes
Storage Subsystem reports	Yes	Yes
Storage Subsystem Performance reports	Yes	Yes
Switch Performance reports	Yes	Yes

Table 43. Available reports with each IBM Tivoli Storage Productivity Center license (continued)

Reports	Tivoli Storage Productivity Center or Storage Productivity Center Select	IBM SmartCloud Virtual Storage Center Storage Analytics Engine
Note: 1. A Storage Resource agent must be deployed on the computer where you want to run batch reports. 2. The following batch reports are not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: • Rollup • Usage • Usage Violations • Backup 3. The following **Asset** > **System-wide** reports are available with the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license: • Monitored Directories • Unmanaged Computers • Users • OS User Groups 4. The following Data system reports are available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: • Disk Space Summary • Disk Defects • Computer Storage Availability • Computer Disk Space • Available File System Space		

Administering Tivoli Storage Productivity Center

Administer Tivoli Storage Productivity Center to authorize users, start the graphical user interfaces, manage services, change passwords, and many other administrative tasks.

Checking the version and license of Tivoli Storage Productivity Center

The version and license of Tivoli Storage Productivity Center that is installed on your system determine the Tivoli Storage Productivity Center functions that are available.

Check the version of Tivoli Storage Productivity Center that is installed on your system to verify that you are using the correct level of the documentation.

Check the license that is installed on your system if documented functions are not available. The functions might be restricted to a different license.

To check the version and license of Tivoli Storage Productivity Center that is installed, complete the following steps in the web-based GUI:

1. Click the question mark icon in the banner pane of the window to display a list of help topics.
2. From the list of help topics, select **About**.

Viewing and managing services

This topic describes viewing and managing services for Data server and Device server.

Data Server services

Data Server directs all monitoring operations. Data Server services include the various services that run in the IBM Tivoli Storage Productivity Center for Data Server component.

The Data Server is the control point for scheduling functions, configuration, event management, reporting, and GUI support. The Data Server directs all of the activities related to the monitoring performed on computers with agents installed and the reporting performed on the client GUI computers.

When you expand the Data Server node, five nodes appear which represent the running services within the server. Select any of the following nodes to display information on the server identification, run time statistics, requests processed, database information, and agent statistics:

Server The component that provides the functionality of the Data Server. The Server, sometimes called Data Server, provides monitoring functions for services, data, and facilities to the network.

Scheduler
The Scheduler is the process that receives requests to schedule operations, such as the running of scheduled pings, probes, and scans. It is a multithread background server process designed to handle the scheduling and launching of jobs, based on a simple timing scheme.

GUI The main graphical user interface that enables you to do monitoring operations, such as generate reports, check the status of agents, create and run jobs, setup monitoring for the Data Server, and so forth.

Device
Converts general commands for peripheral devices and passes them between the Server and the device. Defines I/O devices to the operating system and then controls the operation of these devices.

Agent The process on the server that sends and receives requests from agents for information gathering or processing. Agents can perform specific operations, such as scans or probes, on a regular schedule.

With Data Server services you can perform the following tasks:

- View a Data Server service log
- Configure tracing on a service
- Shut down a service
- Restart a service that is inactive

Additionally, through the Server service you can do the following:

- Configure audit logging
- Broadcast name and port information

To view the status of the server and services, expand **Administrative Services** > **Services** > **Data Server**. A green box icon next to the name of a service (for example, **Scheduler**) means that the service is active and running normally. If the service is not currently active, a red circle icon is displayed next to the name.

Refresh the Data Server list

To refresh the Data Server list, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** > **Services**.
2. Right-click **Data Server**, and click **Refresh**.

Viewing Data Server information with the Server service:

The Server service directs all Data Server information. You can view information about the Server service, including the name and port, disk space, connections, requests, and agent information.

To view Data Server information, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** > **Services** > **Data Server**.
2. Click **Server**. In the content pane, the Server service page opens displaying the following information:

   Server Displays the name of the server.

   Port Displays the port number on which the server is listening to agents.

   Start Time
   Displays the date and time when the server was started.

   Elapsed Time
   Displays the elapsed time when the server was last updated.

   VM Size
   Displays the disk space allocated to the virtual memory size of the server.

   Network Connections
   Displays the number of networks to which the server is connected.

   Maximum Connections
   Displays the maximum number of networks to which the server can connect.

   Pending Requests
   Displays the number of outstanding requests waiting for server response.

   Processed Requests
   Displays the number of requests processed by the server.

   Database
   The following fields display information about the Data Manager's repository database.

   JDBC URL
   Displays the location of the JDBC file used to connect to the database.

JDBC Driver
Displays the name of the JDBC driver used to connect to the database.

DB Connections
Displays the number of connections to the database.

DB Pool Count
Displays the database pool count.

Agents
The following fields provide information about the status of Data Manager agents.

Total The number of agents that are installed.

Online
The number of agents that are currently online (communicating with the server).

Offline
The number of agents that are currently offline (not communicating with the server, but were shut down normally).

Unreachable
The number of agents that are 1) not communicating with the server, and 2) provided no information to the server that they were shutting down.

Upgrading
The number of agents that are currently in the process of upgrading to the same version as the server.

Upgrade Required
The number of agents that must be upgraded to the same version as the server.

Viewing information about Data Server services:

Select any of the Data Server services to display information on that services' identification, run time statistics, requests processed, database information, and agent statistics.

To display information about a service, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** > **Services** > **Data Server**.
2. Click one of the following services:

   Scheduler
   Controls the running of scheduled pings, probes, and scans.

   GUI Enables you to generate report requests and to create and run jobs.

   Device
   Converts commands for peripheral devices.

   Agent Sends and receives requests from agents.
3. In the content pane, the following information about the selected Data server service is displayed:

   Service
   Whether the service is active or down.

Pending Requests
Number of requests waiting to be processed by the scheduler for the current session.

Processed Requests
Number of requests processed by the scheduler for the current session.

Idle Threads
Number of threads waiting for a process request.

Viewing a Data Server service log:

Use the content of a service log file to troubleshoot any errors that might occur during startup, processing, or shutdown of the Data Server service.

To display a service log, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** > **Services** > **Data Server**.
2. Right-click the name of the service and click **View Log**. The service log will open in the content pane.

Service alerts, if any, are displayed in red.

Logs are searchable by search string. Enter the string in the field provided and click the up or down arrow to search.

Configuring tracing for Data Server services:

Configure tracing for the Scheduler, GUI, Device, and Agent services by setting the debug level. Tracing changes can only be made by users with administrator authority.

To configure tracing for the Scheduler, GUI, Device, and Agent services, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** > **Services** > **Data Server**.
2. Right-click the name of the service and click **Configure Tracing**.
3. In the Server Tracing Configuration window, check the **Enable Trace** box to enable tracing for the service. If this box is unchecked, tracing will not be performed.
4. Once tracing has been enabled, select the level of debugging you want to perform. You can select one of the following:
   - DEBUG_MIN - Provides minimal information on errors. This is the default value.
   - DEBUG_MID - Provides more information on errors.
   - DEBUG_MAX - Lists all debugging trace information.
5. Click **OK**.

Tracing can be turned off by de-selecting the **Enable Trace** option in the Server Tracing Configuration window.

Configuring audit logging for the Server:

Use audit logging to record a history of user actions. This function is available for the Server only.

To configure audit logging, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** > **Services** > **Data Server**.
2. Right-click the **Server** node and click **Configure Audit Logging**. The Server Audit Logging Configuration window opens.
3. In the Server Audit Logging Configuration window, check the **Enable Trace** check box to enable tracing for the server component. If this box is unchecked, tracing will not be performed.
4. Specify the audit logging options:

   Maximum Number of Files
   Select the maximum number of trace files that will be created. When the maximum number of files has been reached, tracing will rollover and start writing to the first file. (default is 5)

   The value in this field also applies to the other services for which tracing is enabled.

   Maximum File Size (kilobytes)
   Select the maximum size of each log file. When the current trace file has reached its maximum capacity, tracing will continue to the next available trace file as defined in the Maximum Number of Files field. (default is 20 MB).

   The value in this field also applies to the other services for which tracing is enabled.
5. Click **OK**.

Broadcasting the Data Server location:

Use broadcasting to inform agents of the Data Server location. Broadcasting is available with the Server service only.

Broadcasting informs agents of the server location

To broadcast the Data Server location, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** > **Services** > **Data Server**.
2. Right-click the **Server** node and click **Broadcast**.

Shutting down a Data Server service:

Shut down the active process or processes associated with a service (for example, processes on the server that receive requests from the scheduler). If you shut down a service, you must restart it before the server will resume receiving requests from it.

To shut down a running service, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** > **Services** > **Data Server**.
2. Right-click the name of the service and click **Shutdown**.
3. Select one of the following shutdown methods:

   Normal
   Let all running and scheduled processes complete and then shut down cleanly. During this shutdown process, new jobs are allowed to run if

they were scheduled during the time that the shutdown is taking place. Once all jobs have completed and no more jobs are scheduled to start during the shutdown time, the service will be shut down.

Immediate
Let all running processes complete but not allow scheduled jobs to start. During this shutdown process, new jobs are *not* allowed to start if they were scheduled during the time that the shutdown is taking place. Once all running jobs have completed, the service will be shut down.

Abort Stops whatever processes are currently running and then shuts down the service.

Starting a Data Server service:

Reactivate a service after it has been stopped. The start action is available only for services that are not running.

To start a service, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** > **Services** > **Data Server**.
2. Right-click the name of the service and click **Start**.

Configuring tracing

Configure tracing for the Data server, agent, scheduler, Device server, and GUI by setting the debug level and log file characteristics. The level of information saved in a log file is determined by the tracing level. Tracing settings can only be changed by users with administrator authority.

To configure tracing for IBM Tivoli Storage Productivity Center, complete the following steps:

1. In the navigation tree pane, expand **Administrative Services** > **Services** > **Data Server**.
2. Right-click on one of the following components:
   - Server
   - Agent
   - Scheduler
   - Device
   - GUI

   Click **Configure Tracing**.
3. In the Server Tracing Configuration window, check the **Enable Trace** check box. Specify additional tracing options:

   Level Select the level of debugging you want to perform. Select one of the following debugging level:

   DEBUG_MIN
   Provides minimal information about errors. This option is the default value.

   DEBUG_MID
   Provides more information about errors.

   DEBUG_MAX
   Lists all debugging trace information.

Maximum number of files
Select the maximum number of trace files that is created. When the maximum number of files has been reached, tracing rolls over and start writing to the first file. (default is five)

The value in this field also applies to the other services for the Data server for which tracing is enabled.

Maximum File Size (kilobytes)
Select the maximum size of each log file. When the current trace file has reached its maximum capacity, tracing continues to the next available trace file as defined in the Maximum Number of Files field. (default is 20 MB).

4. Click **OK**.

Tracing can be turned off by clearing the **Enable Trace** option in the Server Tracing Configuration window.

Device server services

A **Device server** node is present in the navigation tree only if a Device server is deployed in IBM Tivoli Storage Productivity Center your environment.

In the Navigation Tree pane, expand **Administrative Services** > **Services** > **Device server** to display a list of the services associated with the Device server. The following services are displayed:

Device server
Detects all the operations of the Device server.

Event Service that monitors external events.

Log Service for configuring logging and tracing.

Configuration Analysis and History
Monitors the SAN configuration tools.

Control
Service that provides control functions.

Fabric Manager
Service that provides fabric management functions.

Disk Manager
Service that provides storage subsystem management functions.

Disk Monitor
Service that monitors the storage subsystems.

Out of Band Agent
Service that maintains the list of out-of-band Fabric agents.

Performance
Service that monitors device performance.

Configuration
Service for configuring Device server properties.

Discovery
Service that provides discovery and probe functions.

Message
Service that provides internal messaging functions.

Data Access functions
Service providing data access and control functions.

Data Server Monitor
Service that provides Data server monitoring functions.

Service Manager
Service that manages the other services.

Server Virtualization Management Service
Service providing server virtualization management functions.

Device server Monitor
Service that provides Device server monitoring functions.

A green box icon next to a service name means that the service is active and running normally. If the service is not currently active, a red circle icon displays next to the name.

Changing the Device server password:

This section provides information about changing the Device server password.

To change the password, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services** > **Services** > **Device server**.
2. Right-click the Device server node and select **Change Password**.
3. In the Change Password window, type the name of the new password twice and click **OK**.

Displaying information about a Device server service

You can view service information including start time, version, and status.

To display information about a service, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services** > **Services** > **Device server**.
2. Click the name of the service.

Viewing a Device server log

You can use the content of the log file to troubleshoot any errors that might occur during startup, processing or shut down. Service alerts, if any, are displayed in red.

To display the log of a service, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services** > **Services** > **Device server**.
2. Right-click the name of the service and select **View Log**.

Changing the user authentication configuration

The Tivoli Storage Productivity Center installation program establishes a default authentication configuration by using the federated repositories feature of the IBM WebSphere Application Server. You can change this authentication configuration.

In the federated repositories framework, the Tivoli Storage Productivity Center installation program creates the following repositories:

File-based user repository

This repository contains the user `tpcFileRegistryUser`. This user password is the same as the common user password that was entered during the Tivoli Storage Productivity Center installation.

Operating system repository

In the federated repositories framework, the Tivoli Storage Productivity Center installation program creates two repositories on the Tivoli Storage Productivity Center WebSphere Application Server web server. This server, which is located in the `TPC_INSTALL_DIR/ewas/profiles/WebServerProfile` directory, is used as the primary WebSphere Application Server for user authentication in Tivoli Storage Productivity Center.

The Device and Replication servers run on the WebSphere Application Server Liberty Profile, and these servers are only configured with File-based user repository. If the web server is down, the Device server and Replication server are used as the backup servers to perform the user authentication and allow the user name that was provided during the Tivoli Storage Productivity Center installation and the `tpcFileRegistryUser` to log in to the Tivoli Storage Productivity Center stand-alone GUI and Tivoli Storage Productivity Center web-based GUIs.

You can add an LDAP repository after you install Tivoli Storage Productivity Center. This configuration is completed on the WebSphere Integrated Solutions Console server. The LDAP repository configuration settings are not propagated to the Device and Replication servers. Therefore, if the web server is down, the authorized LDAP users cannot log in to Tivoli Storage Productivity Center stand-alone GUI and the Tivoli Storage Productivity Center for Replication web-based GUI. The backup user authentication mechanism that is based on Device and Replication servers allows the user name that was provided during the Tivoli Storage Productivity Center installation and the `tpcFileRegistryUser` to log in to the Tivoli Storage Productivity Center stand-alone GUI and the Tivoli Storage Productivity Center web-based GUIs.

On computers that are members of a Windows domain, the local OS repository also contains the domain users and groups that are managed by the Windows domain, if the computer is correctly configured with the Windows domain.

The LDAP repositories that are supported by Tivoli Storage Productivity Center depend on the WebSphere Application Server support. For more information about the supported LDAP repositories, see the relevant topic for your operating system and search for *LDAP Servers using Federated Repository Configuration*:

- For the Windows operating system, see http://www.ibm.com/support/docview.wss?rs=180&uid=swg27012421 .
- For the AIX operating system, see http://www.ibm.com/support/docview.wss?rs=180&uid=swg27012389.
- For the Linux operating system, see http://www.ibm.com/support/docview.wss?rs=180&uid=swg27012415.

When you change the user authentication configuration by adding or removing an LDAP repository in the federated repositories framework, you must first back up the existing WebSphere Application Server configuration files. You must also back up the WebSphere configuration files after you add an LDAP repository to the user

authentication configuration and want to later change the LDAP authentication settings. You also need to backup, and then restore the `isc.ear` file in the `TIP_installation_directory`.

If these users or groups are present in more than one repository in the federated repositories framework, the WebSphere Application Server cannot resolve duplicated users or groups. An example of a duplicated user is, for example, when an Administrator user exists in both the local OS and LDAP repository. You must ensure that the duplicated users (or groups) are not used during the configuration or to manage Tivoli Storage Productivity Center.

Adding an LDAP repository to the federated repositories

You can configure Tivoli Storage Productivity Center and Jazz for Service Management to communicate with an external Lightweight Directory Access Protocol (LDAP) repository, such as IBM Tivoli Directory Server or Microsoft Active Directory. When you change the authentication configuration, Tivoli Storage Productivity Center is available to users and groups in other repositories.

Important: When you install Tivoli Storage Productivity Center on a computer that is a member of a Windows domain, Active Directory users and groups also exist in the local OS repository. To determine whether the Active Directory users exist in the local OS repository, log in to the WebSphere Integrated Solutions Console and click **Users and Groups** > **Manage users**. If the Active Directory users exist in the local OS repository, you should not add that same Active Directory as an LDAP repository to the federated repositories in Tivoli Storage Productivity Center Version 5.2.

Important: The IBM WebSphere Application Server cannot resolve duplicated users or groups when these users or groups are present in more than one repository in the federated repositories framework. For example, an Administrator user can exist in both the local OS and LDAP repository. You must ensure that the duplicated users (or groups) are not used during the configuration or to manage Tivoli Storage Productivity Center.

Tivoli Storage Productivity Center and Jazz for Service Management each have their own WebSphere Application Server instance. You must configure both of these WebSphere instances to communicate with the LDAP repository. The procedure to configure these instances are almost identical, so the steps are provided only once. You must repeat the steps to configure both of these WebSphere Application Server instances with the LDAP repository. The WebSphere Application Server instance in Tivoli Storage Productivity Center is also called the web server.

This procedure uses the variable name *WebSphere_Directory* to indicate where WebSphere Application Server is located.

The location of the WebSphere Application Server directory is different for each instance:

- The Jazz for Service Management WebSphere Application Server directory: `JAZZSM_INSTALL_DIR/profile`
- The Tivoli Productivity WebSphere Application Server or web server directory: `TPC_INSTALL_DIR/ewas/profiles/WebServerProfile`

Important: Some of the field names can be different between the Tivoli Storage Productivity Center WebSphere instance and the Jazz for Service Management WebSphere instance.

To add an LDAP repository to the federated repositories in Tivoli Storage Productivity Center or Jazz for Service Management, complete the following steps:

Tip: If you need assistance, contact your LDAP server administrator.

1. Before you add an LDAP repository to the federated repositories, complete the following steps:

   Important: If you log in by using a Windows domain user name, before you run the backup commands, click **Start** > **Command Prompt** and select **Run as administrator**.

   a. Back up the WebSphere Application Server configuration for the Tivoli Storage Productivity Center instance of WebSphere Application Server and back up the WebSphere Application Server configuration for the Jazz for Service Management instance of WebSphere Application Server.

      - On Windows operating systems, run the following command:

        ```
        WebSphere_Directory\bin\backupConfig.bat -username
        adminuser -password adminpassword -nostop
        ```

        In the WebSphere Application Server configurations for Tivoli Storage Productivity Center and Jazz for Service Management, *adminuser* is the user name that was used to install Tivoli Storage Productivity Center (for example, `db2admin`) or Jazz for Service Management (for example `smadmin`), and *adminpassword* is the password that is associated with *adminuser*.

        Tip: If you completed a default installation on the Windows operating system:

        – The Tivoli Storage Productivity Center WebSphere Application Server directory is located here:

          ```
          C:\Program Files\IBM\TPC\ewas\profiles\
          WebServerProfile\
          ```

        – The Jazz for Service Management WebSphere Application Server directory is located here:

          ```
          C:\Program Files\IBM\JazzSM\profile\
          ```

      - On AIX or Linux operating systems, run the following command:

        ```
        WebSphere_Directory/bin/backupConfig.sh -username
        adminuser -password adminpassword -nostop
        ```

        In the WebSphere Application Server configurations for Tivoli Storage Productivity Center and Jazz for Service Management, *adminuser* is the user name that was used to install Tivoli Storage Productivity Center (for example, `db2inst1`) or Jazz for Service Management (for example `smadmin`), and *adminpassword* is the password that is associated with *adminuser*.

        Tip: If you completed a default installation on the AIX or Linux operating system:

        – The Tivoli Storage Productivity Center WebSphere Application Server directory is located here:

          ```
          /opt/IBM/TPC/ewas/profiles/WebServerProfile/
          ```

        – Jazz for Service Management WebSphere Application Server directory is located here:

/opt/IBM/JazzSM/profile/

b. Back up the soap.client.props file for the Tivoli Storage Productivity Center instance of WebSphere Application Server.
   - On the Windows operating system, run the **copy** command to back up this file:

     WebSphere_Directory\properties\soap.client.props

     Tip: If you completed a default installation on the Windows operating system, the Tivoli Storage Productivity Center Version 5.2 soap.client.props file is in this directory:

     C:\Program Files\IBM\TPC\ewas\profiles\WebServerProfile \properties\
   - On the AIX or Linux operating systems, run the **cp** command to back up this file:

     WebSphere_Directory/properties/soap.client.props

     Tip: If you completed a default installation on the AIX or Linux operating systems, the Tivoli Storage Productivity Center Version 5.2 soap.client.props file is in this directory:

     /opt/IBM/TPC/ewas/profiles/WebServerProfile/properties

c. Back up the IBM Cognos.ear directory for the Jazz for Service Management WebSphere Application Server.
   - On the Windows operating system, run the **xcopy** command to back up this directory:

     WebSphere_Directory\installedApps\JazzSMNode01Cell \IBM Cognos.ear

     Tip: If you completed a default installation on the Windows operating system, the Jazz for Service Management WebSphere Application Server directory is located here:

     C:\Program Files\IBM\JazzSM\profile\
   - On the AIX or Linux operating systems, run the **cp** command to back up this directory:

     WebSphere_Directory/installedApps/JazzSMNode01Cell /IBM Cognos.ear

     Tip: If you completed a default installation on the AIX or Linux operating systems, the Jazz for Service Management WebSphere Application Server directory is located here:

     /opt/IBM/JazzSM/profile/

   Tip: The default cell name JazzSMNode01Cell is used in the sample commands. If you did not use the default cell name when you installed Jazz for Service Management, you must specify the correct cell name in this step.

2. In the event of a problem when you add an LDAP repository, complete the following steps to restore the items you backed up in step 1:
   a. Run the restoreConfig.sh or the restoreConfig.bat command to restore the WebSphere Application Server configuration for the Tivoli Storage Productivity Center instance of WebSphere Application Server and restore the WebSphere Application Server configuration for the Jazz for Service Management instance of WebSphere Application Server. The **restoreConfig** command is in one of the following directories:

- On the Windows operating system, go to this directory:

  WebSphere_directory/bin/restoreConfig.bat

  Tip: If you completed a default installation on the Windows operating system:
  - The Tivoli Storage Productivity Center WebSphere Application Server directory is located here:

    C:\Program Files\IBM\TPC\ewas\profiles\WebServerProfile\
  - The Jazz for Service Management WebSphere Application Server directory is located here:

    C:\Program Files\IBM\JazzSM\profile\
- On the AIX or Linux operating system, go to this directory:

  WebSphere_directory/bin/restoreConfig.sh

  Tip: If you completed a default installation on the AIX or Linux operating system, the Tivoli Storage Productivity Center Version 5.2 soap.client.propsis in this directory:

  /opt/IBM/TPC/ewas/profiles/WebServerProfile/properties

b. Restore the soap.client.props file for the Tivoli Storage Productivity Center instance of WebSphere Application Server.
   - On the Windows operating system, run the **copy** command to restore this file:

     WebSphere_Directory\properties\soap.client.props

     Tip: If you completed a default installation on the Windows operating system, the Tivoli Storage Productivity Center Version 5.2 soap.client.props is in this directory:

     C:\Program Files\IBM\TPC\ewas\profiles\WebServerProfile\properties\
   - On the AIX or Linux operating system, run the **cp** command to restore this file:

     WebSphere_Directory/properties/soap.client.props

     Tip: If you completed a default installation on the AIX or Linux operating systems, the Tivoli Storage Productivity Center Version 5.2 soap.client.props is in this directory:

     /opt/IBM/TPC/ewas/profiles/WebServerProfile/
c. Restore the IBM Cognos.ear directory for the Jazz for Service Management WebSphere Application Server.
   - On the Windows operating system, run the **xcopy** command to restore this directory:

     WebSphere_Directory\installedApps\JazzSMNode01Cell\IBM Cognos.ear

     Tip: If you completed a default installation on the Windows operating system, the Jazz for Service Management WebSphere Application Server directory is located here:

     C:\Program Files\IBM\JazzSM\profile\
   - On the AIX or Linux operating systems, run the **cp** command to restore this directory:

     WebSphere_Directory/installedApps/JazzSMNode01Cell/IBM Cognos.ear

Note: The default cell name `JazzSMNode01Cell` is used in the sample commands. If you did not use the default cell name when you installed Jazz for Service Management, you must specify the correct cell name in this step.

Tip: If you completed a default installation on the AIX or Linux operating systems, the Jazz for Service Management WebSphere Application Server directory is located here:

`/opt/IBM/JazzSM/profile/`

d. Restart the Tivoli Storage Productivity Center web server and the Jazz for Service Management server. The execution of the **restoreConfig** command stops those WebSphere Application Server instances.

For information about starting the Tivoli Storage Productivity Center web server or Jazz for Service Management server, see the Tivoli Storage Productivity Center information center. Search for *Starting the Tivoli Storage Productivity Center services*.

3. To access the WebSphere Integrated Solutions Console, open a web browser and enter one of the following web addresses in the address field:
   - `http://hostname:port/ibm/console/logon.jsp`
   - `https://hostname:port/ibm/console/logon.jsp`

   The hostname is the server that is running WebSphere Application Server, such as the server name or IP address, and port is the port number for the WebSphere Application Server. The port number differs depending on which protocol you used (http or https) and the options that you selected when you installed Tivoli Storage Productivity Center.

   To determine the port number, complete the following steps:

   a. Open the *WebSphere_Directory*`/properties/portdef.props` file.
   b. The port number is the value that is assigned to one of the following keys:
      - For protocols that are not secure (for example, `http://`):

        `WC_adminhost`
      - For protocols that are secure (for example, `https://`):

        `WC_adminhost_secure`
4. Log in to the WebSphere Integrated Solutions Console. To complete this procedure, your user name must have Administrator authorization in the WebSphere Integrated Solutions Console.

   Important: For new installations of Tivoli Storage Productivity Center, *Administrator* is the Tivoli Storage Productivity Center common user name. You can log in to the Tivoli Storage Productivity Center GUI by using most of the user names in the Administrator group, but you can only use the Tivoli Storage Productivity Center common user name to log in to the WebSphere Integrated Solutions Console.
5. In the WebSphere Integrated Solutions Console navigation tree, click **Security** > **Global security**.
6. On the Global security page, in the User account repository section, click **Configure** next to the **Available realm definitions** menu.

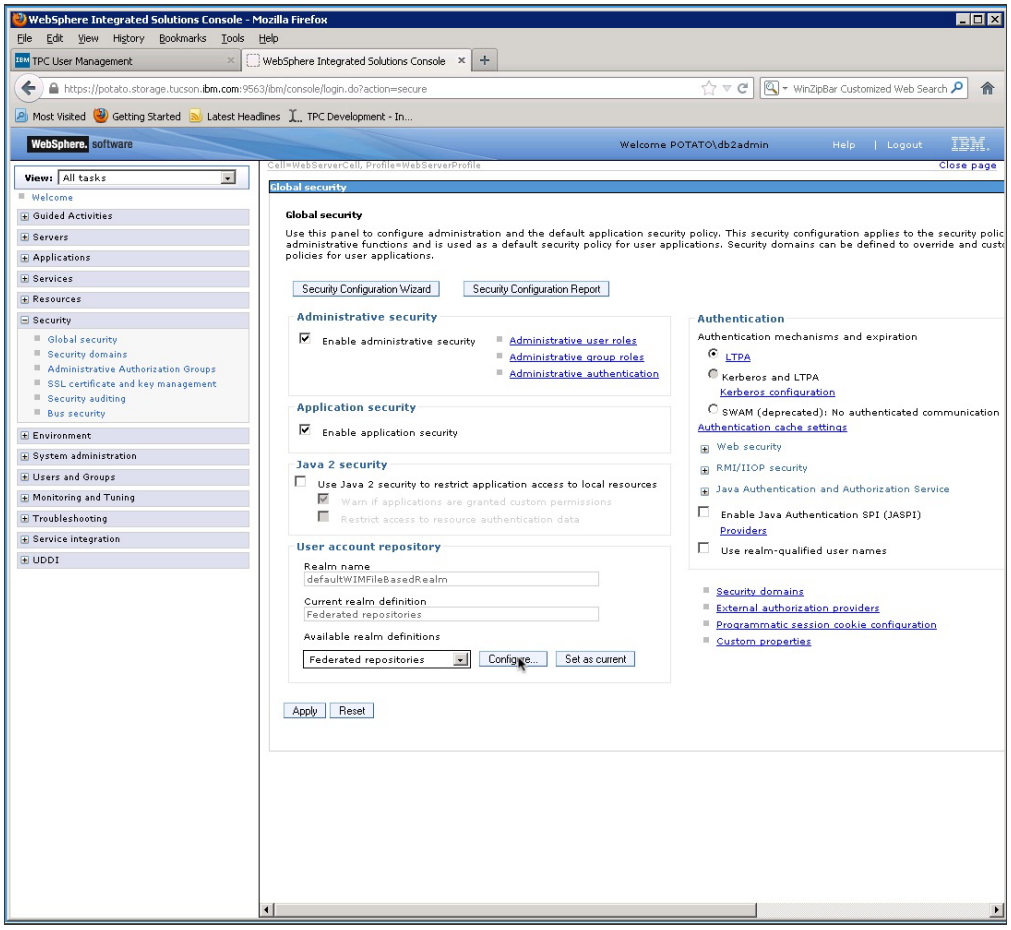

Figure 2. Global security page, configure federated repositories

7. On the **Global security** > **Federated repositories** page, under **Related Items**, click the **Manage repositories** link.

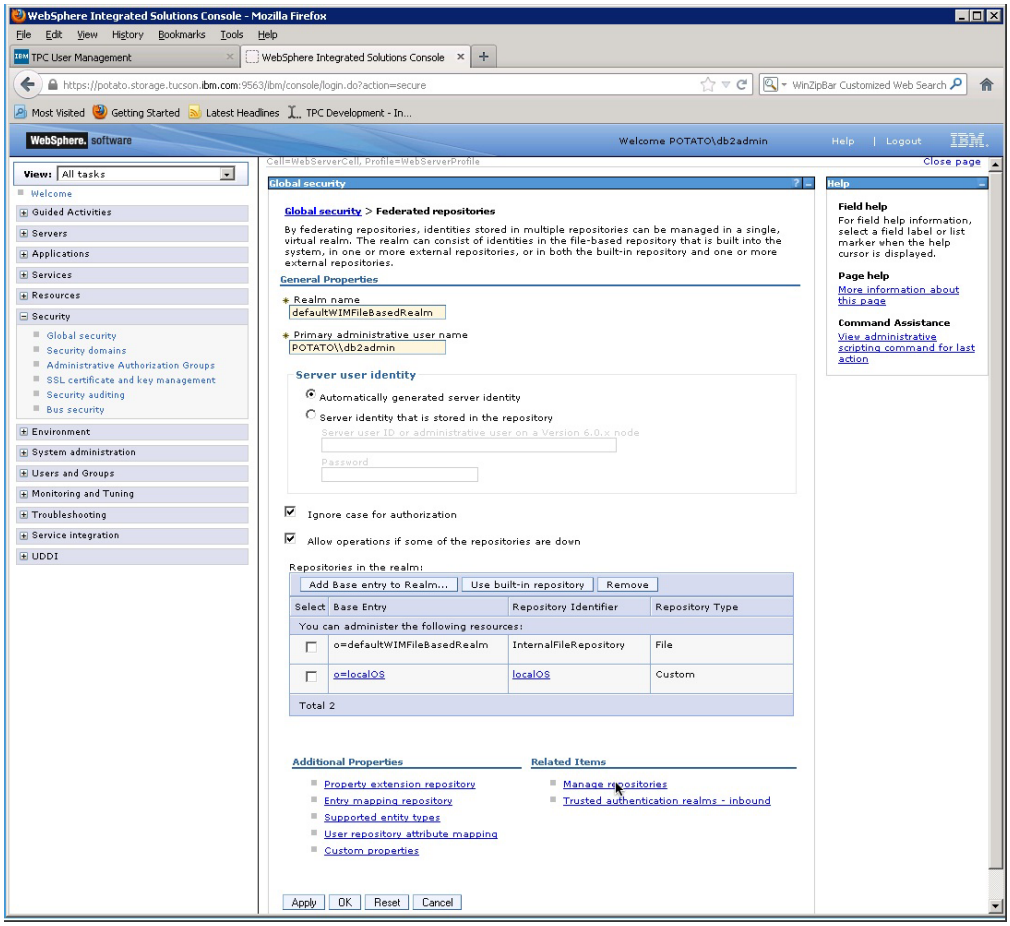

Figure 3. Federated repositories page, Manage repositories

8. On the **Global security** > **Federated repositories** > **Manage repositories** page, add the LDAP repository that you want to use for authentication.

   To add the LDAP repository, complete these steps:

   a. Click **Add** > **LDAP repository** to add a new repository.

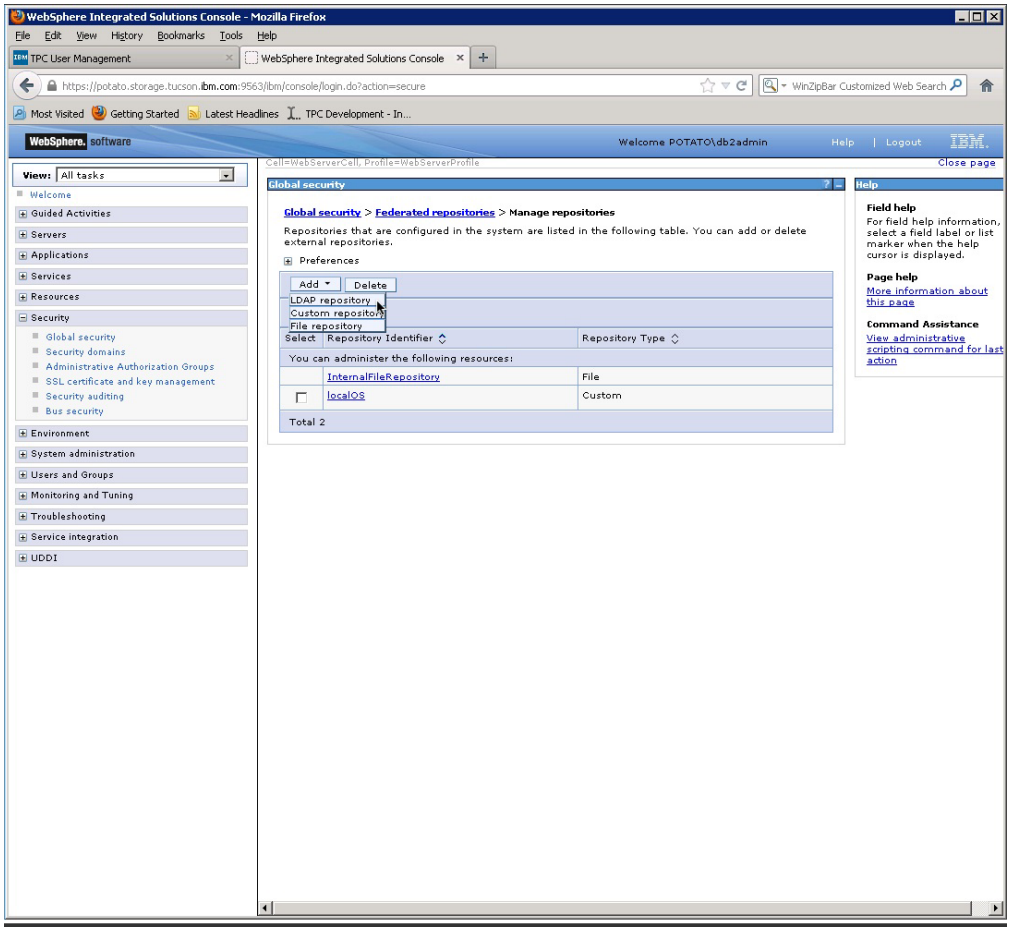

Figure 4. Manage repositories page, Add a new repository

b. Enter the values for the following fields:

Repository identifier
A unique identifier for the LDAP repository, which identifies the repository in the realm, for example, LDAP1.

Directory type
The type of LDAP server to which you want to connect.

Primary host name
The host name of the primary LDAP server. This host name is either an IP address or a domain name service (DNS) name.

Port The LDAP server port. The default value is 389. Depending on your LDAP server configuration, you can specify a different port. If you do not know which port to use, contact your LDAP server administrator.

Bind distinguished name
The distinguished name (DN) for WebSphere Application Server to use when it binds to the LDAP repository. If no name is specified, WebSphere Application Server binds anonymously to the LDAP repository. In most cases, bind DN and bind password are required. However, when an anonymous bind can satisfy all of the required functions, bind DN and bind password are not required.

If you are not sure whether an anonymous bind has satisfied the required functions, contact your LDAP server administrator.

Attention: There is no single value for the **Bind distinguished name** field that is correct for every Active Directory Server or for every LDAP server. The correct value for the **Bind distinguished name** field depends on the configuration of your Active Directory Server or your LDAP server. If you are unsure about the correct value to use for the **Bind distinguished name** field, contact your LDAP server administrator.

Bind password
The password for WebSphere Application Server to use when you bind to the LDAP repository.

Login properties
The authentication properties that are used to log on to WebSphere Application Server. Type `uid;cn` in this field. This value enables WebSphere Application Server to use the property that is required for the directory type.

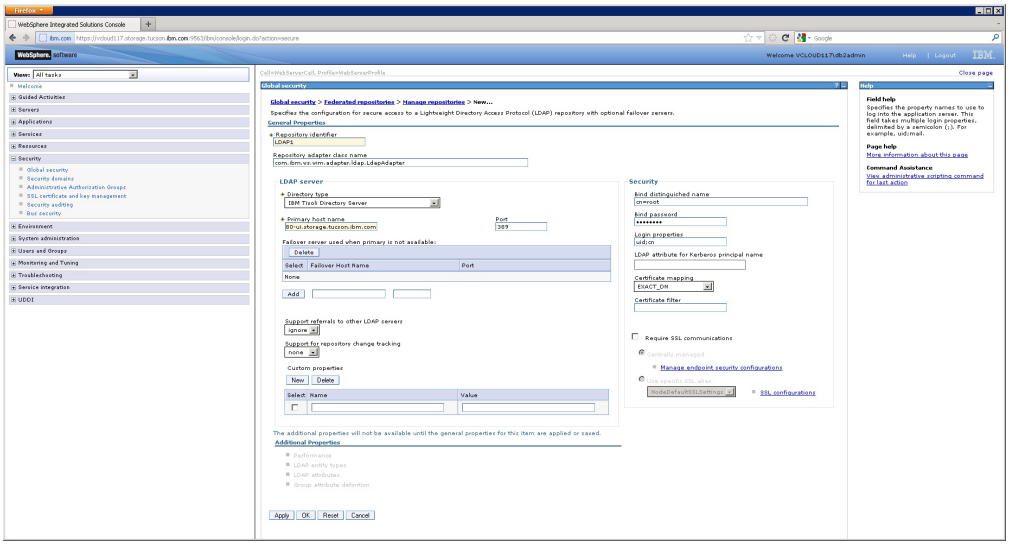

Figure 5. Manage repositories > New page

c. Click **OK**.

d. In the messages dialog box that is displayed on the Manage repositories page, click the **Save** link in **Save directly to the master configuration**.

Important: On the **Global security** > **Federated repositories** > **Manage repositories** page, do not delete the local OS repository.

9. From the Manage repositories page, return to the **Global security** > **Federated repositories** page.
10. In the **Repositories in the realm** panel, click **Add Base entry to Realm**.

In the Jazz for Service Management WebSphere Integrated Solutions Console, the label for this button is **Add repositories (LDAP, custom, etc)**.

Important: Do not change the Primary Administrator user name.

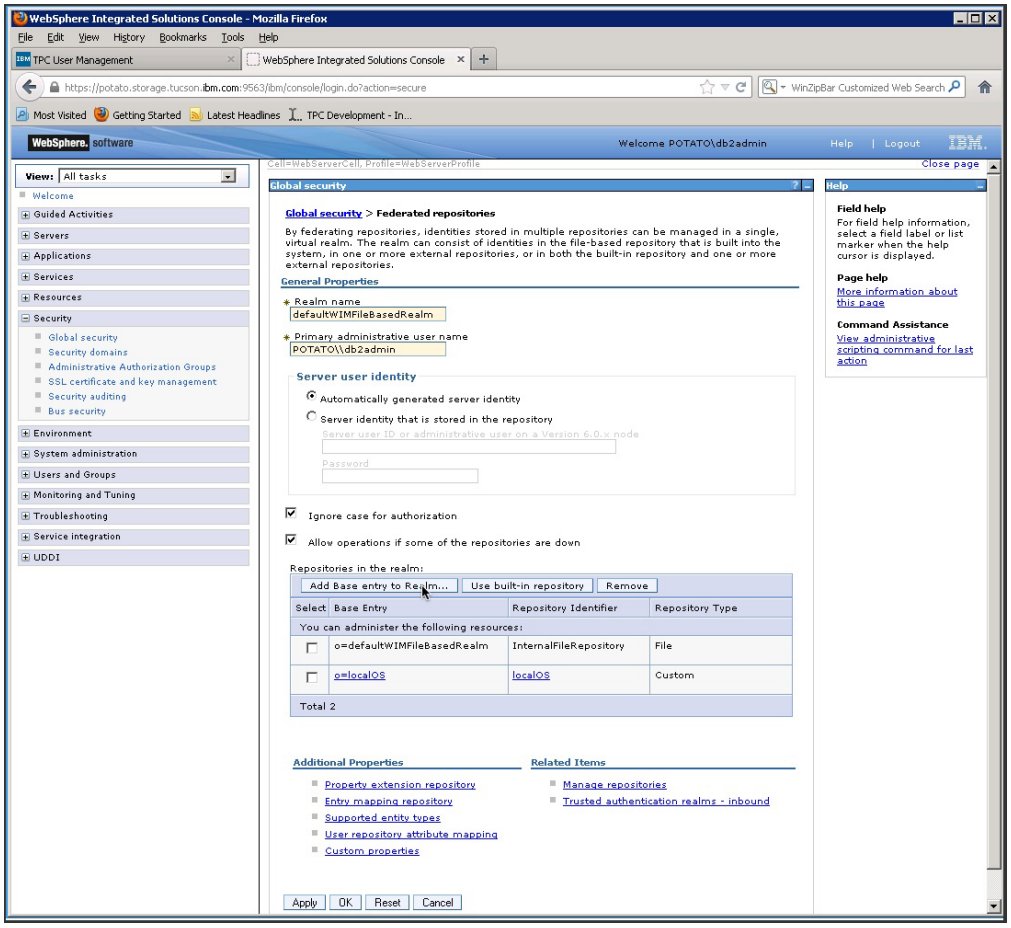

Figure 6. Global security > Federated repositories page, Add base entry to realm

11. If the Messages dialog box is displayed on the Repository reference page, click the **Save** link in **Save directly to the master configuration**.
12. On the Repository reference page, configure the following items:
    a. In the **Repository list**, select the repository that you created in step 8 on page 154.
    b. In the **Distinguished name of a base entry that uniquely identifies this set of entries in the realm** field, enter a DN for the repository. This DN maps to the DN of the base entry in the LDAP repository that you entered in the **Distinguished name of a base entry in this repository** field.

       Tip: In Jazz for Service Management WebSphere Integrated Solutions Console, the label for this field is **Unique distinguished name of the base (or parent) entry in federated repositories**.

       To avoid duplicate results during searches, the DN must uniquely identify the base entry in the repository. If multiple repositories are included in the realm and the repositories have the same base entry, use this field to define a DN that uniquely identifies each base entry. For example, repositories LDAP1 and LDAP2 might both use `o=ibm,c=us` as the base entry in the repository. Enter a DN in this field that distinguishes the base entries for each repository. For example: `o=ibm,c=us` for LDAP1 and `o=ibm2,c=us` for LDAP2.
    c. In the **Distinguished name of a base entry in this repository** field, enter the DN of the base entry in the LDAP repository that you want to map to

the DN that you entered in the **Distinguished name of a base entry that uniquely identifies this set of entries in the realm** field. In most instances, the value is the same in both fields.

The value in this field indicates the starting point for searches in the LDAP directory server. For example, for a user with a DN of `cn=John Doe, ou=rochester, o=ibm, c=us`, you can specify the LDAP base entry as any of the following options:

- `ou=rochester, o=ibm, c=us`
- `o=ibm, c=us`
- `c=us`

Important: The DN value that is entered in this field must be broad enough to include both users and the groups to which the users belong. For example, if a user in `ou=rochester, o=ibm, c=us` is also a member of groups that are in `dc=stategroups, ou=rochester, o=ibm, c=us`, enter `o=ibm, c=us` in this field.

In the Jazz for Service Management WebSphere Integrated Solutions Console, the label for this field is **Distinguished name of a subtree in the main repository**. You must first select the **Distinguished name in the repository is different** check box and then enter a value in the **Distinguished name of a subtree in the main repository** field. In most instances, you will enter the same value in the **Unique distinguished name of the base (or parent) entry in federated repositories** field and the **Distinguished name of a subtree in the main repository** field.

d. Click **OK**.

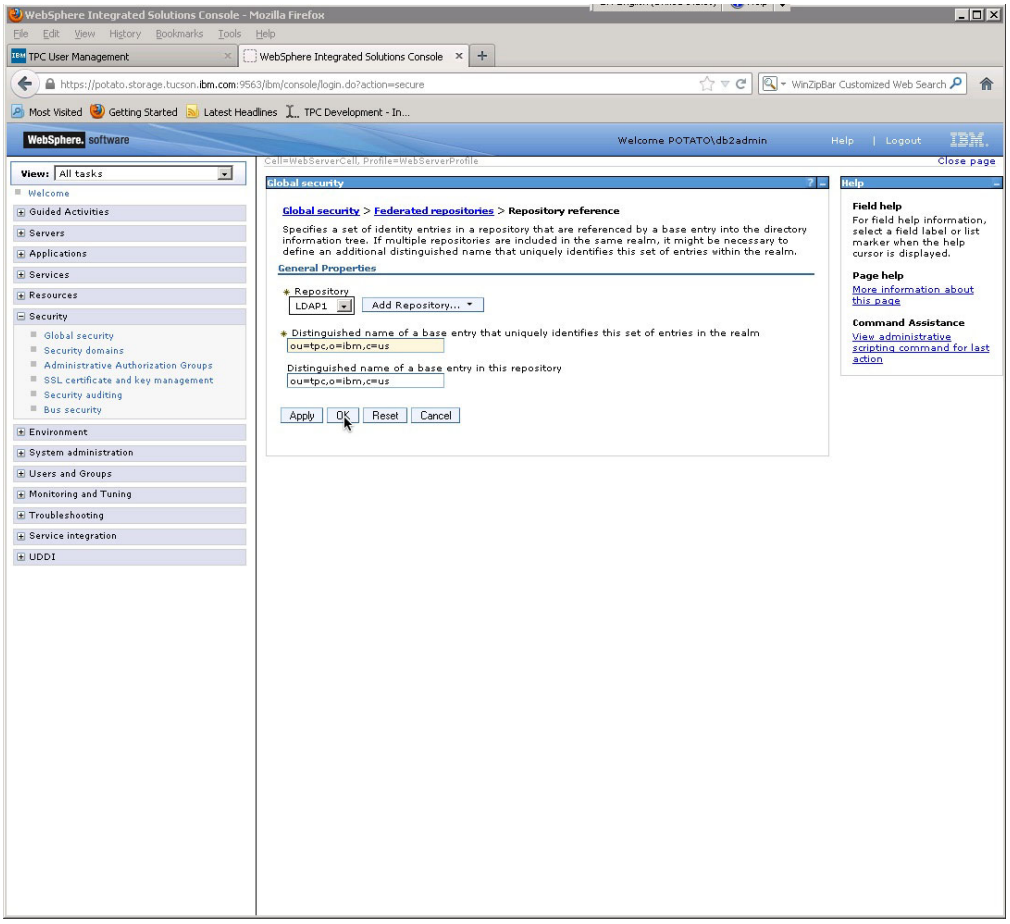

Figure 7. Global security > Federated repositories > Repository reference page, Add base entry to realm

e. In the messages dialog box that is displayed, click the **Save** link in **Save directly to the master configuration**.

   Important: In the Repositories in the realm table, do not remove the **localOS** entry or the **InternalFileRepository** entry.

13. Log out of the WebSphere Integrated Solutions Console.
14. Log out of the web-based GUI.
15. Stop and restart Tivoli Storage Productivity Center web server or Jazz for Service Management server.

    For information about stopping and starting the Tivoli Storage Productivity Center web server or Jazz for Service Management server, see the Tivoli Storage Productivity Center information center. Search for *Starting and stopping the Tivoli Storage Productivity Center services*.

    To determine the port number for Jazz for Service Management, complete the following steps:

    a. Open the *JazzSM_install_dir*`/profile/properties/portdef.props` file.
    b. Review the value for one of the following parameters:
       - `WC_adminhost` (for ports that are not secure)
       - `WC_adminhost_secure` (for secure ports)
    c. Configure Jazz for Service Management with LDAP:

1) Open a browser window and start the Jazz for Service Management server.
2) Specify the `http` or `https` port from step b to access Jazz for Service Management.

Note: You must repeat steps 5-15 to configure the Jazz for Service Management instance of WebSphere Application Server.

d. Stop and start the Jazz for Service Management server.

Note: If you installed Jazz for Service Management in the default installation location (for example, `C:\Program Files\IBM\JazzSM\profile\bin`), you can run the following commands to stop and start the Jazz for Service Management server:

To stop the Jazz for Service Management:

```
stopServer.bat server1 -username user_name -password password
```

To start the Jazz for Service Management server again:

```
startServer.bat server1
```

e. After the Jazz for Service Management server is restarted, wait 15 minutes before you log in to the Jazz for Service Management server.

To verify that the LDAP federated repository is configured correctly, complete the following steps:

1. Log in to the WebSphere Integrated Solutions Console by using the same user name and password from step 5 on page 152.

   Tip: If you try to log in to the WebSphere Integrated Solutions Console by using a local OS user name or a domain user name, an error message states that the user name or password is invalid. This error may occur because the user name that you are using to log in to the WebSphere Integrated Solutions Console exists in the LDAP federated repository that you just added.

   To resolve this issue, complete one of the following tasks:
   - For local OS user names, add the computer name as a prefix to the user name
   - For domain user names, add the domain name as a prefix to the user name
2. In the WebSphere Integrated Solutions Console navigation tree, click **Users and Groups** > **Manage Users**.
3. In the Search by list, select **User ID**.
4. Click **Search** to search for users in the federated repositories. The list of users includes users from the local file repository, the operating system repository, and the LDAP repository.
5. In the WebSphere Integrated Solutions Console navigation tree, click **Users and Groups** > **Manage Groups**.
6. In the Search by list, select **Group name**.
7. Click **Search** to search for groups in the federated repositories. The list of groups includes groups from the operating system repository and the LDAP repository.

   Important: Before you establish the authorization configuration for Jazz for Service Management, Tivoli Storage Productivity Center, and Tivoli Storage

Productivity Center for Replication, ensure that there are no duplicated user names or group names in the local file repository, the operating system repository, and the LDAP repository.

You should also assign Tivoli Storage Productivity Center for Replication roles, such as Monitor or Administrator, to the user groups in your operating system repository or your LDAP repository. In the Tivoli Storage Productivity Center web-based GUI, click **Settings** > **User Management** to assign the roles to the user groups.

For information about assigning roles to users and roles to groups for Tivoli Storage Productivity Center for Replication, go to the Tivoli Storage Productivity Center information center. Search for *Managing replication*. On the Managing replication page, click **Security**.

Enabling secure communication between Tivoli Storage Productivity Center and the LDAP repository:

You can use the Secure Socket Layer (SSL) protocol to secure the communication between Tivoli Storage Productivity Center and the LDAP repository that you are using for user authentication. The SSL protocol provides security and data integrity for communications over Transmission Control Protocol/Internet Protocol (TCP/IP) networks.

Tivoli Storage Productivity Center and Jazz for Service Management have their own WebSphere Application Server instance. You can configure both these WebSphere instances to communicate with LDAP repository. The procedure to configure these instances are almost identical, so the steps are provided only once. You must repeat the steps to configure both of these WebSphere Application Server instances with the LDAP repository. The WebSphere Application Server instance in Tivoli Storage Productivity Center is also called the web server.

The location of the WebSphere Application Server directory is different for each instance:

- The Jazz for Service Management WebSphere Application Server directory: *JAZZSM_INSTALL_DIR*/profile
- The Tivoli Productivity WebSphere Application Server or web server directory: *TPC_INSTALL_DIR*/ewas/profiles/WebServerProfile

To enable SSL for LDAP communications, you must complete the following steps in the WebSphere Integrated Solutions Console:

1. To access the WebSphere Integrated Solutions Console, open a web browser and one of the following web addresses in the address field:
   - http://hostname:port/ibm/console/logon.jsp
   - https://hostname:port/ibm/console/logon.jsp

   The hostname is the server that is running WebSphere Application Server, such as the server name or IP address, and port is the port number for the WebSphere Application Server. The port number differs depending on which protocol you used (http or https) and the options that you selected when you installed Tivoli Storage Productivity Center.

   To determine the port number, complete the following steps:

   a. Open the *WebSphere_Directory*/properties/portdef.props file.
   b. The port number is the value that is assigned to one of the following keys:

- For `http://` protocols:

  `WC_adminhost`
- For `https://` protocols:

  `WC_adminhost_secure`

2. Log in to the WebSphere Integrated Solutions Console. To complete this procedure, your user name must have Administrator authorization in the WebSphere Integrated Solutions Console.
3. In the WebSphere Integrated Solutions Console navigation tree, click **Security** > **SSL certificate and key management**.
4. On the SSL certificate and key management page, in the **Related Items** section, click **Key stores and certificates**.
5. On the **SSL certificate and key management** > **Key stores and certificates** page, in the table, click **NodeDefaultTrustStore**.
6. On the **SSL certificate and key management** > **Key stores and certificates** > **NodeDefaultTrustStore** page, in the **Additional Properties** section, click **Signer Certificates**.
7. On the **SSL certificate and key management** > **Key stores and certificates** > **NodeDefaultTrustStore** > **Signer Certificates** page, click **Retrieve from port**.
8. On the **SSL certificate and key management** > **Key stores and certificates** > **NodeDefaultTrustStore** > **Signer Certificates** > **Signer Certificates** > **Retrieve from port** page, enter the values for the following fields:

   Host The fully qualified host name and domain name of your LDAP-compliant repository.

   Port The port where the LDAP repository is listening for secure communications; this port is usually 636.

   SSL configuration for outbound connection
   Accept the default value.

   Alias An alias name for the retrieved certificate (for example, `LDAPCert`)
9. Click **Retrieve signer information**.
10. When the signer information is displayed, click **OK**.
11. In the messages dialog box that is displayed on the Signer certificates page, click the **Save** link in **Save directly to the master configuration**.
12. In the WebSphere Integrated Solutions Console navigation tree, click **Security** > **Global security**.
13. On the Global security page, in the **User account repository** section, click **Configure** next to the **Available realm definitions** menu.
14. On the **Global security** > **Federated repositories** page, under **Related Items**, click the **Manage repositories** link.
15. On the **Global security** > **Federated repositories** > **Manage repositories** page, click the identifier for the LDAP repository for which you want to enable the SSL protocol.
16. On the configuration page for the LDAP repository, configure the following items:
    a. In the **Port** field, enter the port where your LDAP repository is listening for secure communications; this port is usually 636.
    b. Select the **Require SSL communications** check box.
    c. Select **Use specific SSL alias** and select the alias that was created in step 8.

    Click **OK**.

17. On the **Global security** > **Federated repositories** > **Manage repositories** page, click the **Save** link in **Save directly to the master configuration**.
18. Log out from the WebSphere Integrated Solutions Console.
19. Stop and restart Tivoli Storage Productivity Center web server or Jazz for Service Management server. For information about stopping and starting the Tivoli Storage Productivity Center web server or Jazz for Service Management server, see the Tivoli Storage Productivity Center information center. Search for *Starting and stopping Tivoli Storage Productivity Center services*.

Disabling secure communication between Tivoli Storage Productivity Center and the LDAP repository:

You can disable the Secure Socket Layer (SSL) protocol between the LDAP repository and the Tivoli Storage Productivity Center system at any time.

Tivoli Storage Productivity Center and Jazz for Service Management have their own IBM WebSphere Application Server instance. You can disable the use of SSL for LDAP communications in both of these WebSphere instances. The procedure to modify these instances are almost identical, so the steps are provided only once. You must repeat the steps to disable the use of SSL for LDAP communications in both of these WebSphere Application Server instances.

The location of the WebSphere Application Server directory is different for each instance:

- The Jazz for Service Management WebSphere Application Server directory:
  JAZZSM_INSTALL_DIR/profile
- The Tivoli Productivity WebSphere Application Server or web server directory:
  TPC_INSTALL_DIR/ewas/profiles/WebServerProfile

To disable the use of SSL for LDAP communications, complete the following steps in the WebSphere Integrated Solutions Console:

1. To access the WebSphere Integrated Solutions Console, open a web browser and one of the following web addresses in the address field:
   - http://hostname:port/ibm/console/logon.jsp
   - https://hostname:port/ibm/console/logon.jsp

   The hostname is the server that is running WebSphere Application Server, such as the server name or IP address, and port is the port number for the WebSphere Application Server. The port number differs depending on which protocol you used (http or https) and the options that you selected when you installed Tivoli Storage Productivity Center.

   To determine the port number, complete the following steps:
   a. Open the *WebSphere_Directory*/properties/portdef.props file.
   b. The port number is the value that is assigned to one of the following keys:
      - For http:// protocols:
        WC_adminhost
      - For https:// protocols:
        WC_adminhost_secure
2. Log in to the WebSphere Integrated Solutions Console. To complete this procedure, your user name must have Administrator authorization in the WebSphere Integrated Solutions Console.
3. In the WebSphere Integrated Solutions Console navigation tree, click **Security** > **Global security**.

4. On the Global security page, in the **User account repository** section, click **Configure** next to the **Available realm definitions** menu.
5. On the **Global security** > **Federated repositories** page, under **Related Items**, click the **Manage repositories** link.
6. On the **Global security** > **Federated repositories** > **Manage repositories** page, click the identifier for the LDAP repository for which you want to disable the SSL protocol.
7. On the configuration page for the LDAP repository, configure the following items:
   a. In the **Port** field, enter the port where your LDAP repository is listening for nonsecure communications; this port is usually 389.
   b. Clear the **Require SSL communications** check box.

   Click **OK**.
8. In the messages dialog that is displayed on the **Global Security** > **Federated repositories** > **Manage repositories** page, click the, click the **Save** link in **Save directly to the master configuration**.
9. Log out from the WebSphere Integrated Solutions Console.
10. Stop and restart Tivoli Storage Productivity Center web server or Jazz for Service Management server. For information about stopping and starting the Tivoli Storage Productivity Center web server or Jazz for Service Management server, see the Tivoli Storage Productivity Center information center. Search for *Starting and stopping the Tivoli Storage Productivity Center services*.

Removing an LDAP repository from the federated repositories

To remove an LDAP repository from the federated repositories, you must use the IBM WebSphere Integrated Solutions Console.

If removing an LDAP repository from the federated repositories in Tivoli Storage Productivity Center leaves only the local OS repository and the file-based repository, the use of the Tivoli Storage Productivity Center single sign-on feature is limited. Storage system element managers do not support the local OS repository for single sign-on, even if the element manager is installed on the same system as Tivoli Storage Productivity Center.

The location of the WebSphere Application Server directory is different for each instance:

- The Jazz for Service Management WebSphere Application Server directory:
  `JAZZSM_INSTALL_DIR/profile`
- The Tivoli Storage Productivity Center WebSphere Application Server or web server directory:
  `TPC_INSTALL_DIR/ewas/profiles/WebServerProfile`

To remove an LDAP repository to the federated repositories in Tivoli Storage Productivity Center or Jazz for Service Management, complete the following steps:

1. Before you remove an LDAP repository to the federated repositories, complete the following steps:

   Important: If you log in by using a Windows domain user name, before you run the backup commands, click **Start** > **Command Prompt** and select **Run as administrator**.
   a. Back up the WebSphere Application Server configuration for the Tivoli Storage Productivity Center instance of WebSphere Application Server and

back up the WebSphere Application Server configuration for the Jazz for Service Management instance of WebSphere Application Server.

- On Windows operating systems, run the following command:

  ```
  WebSphere_Directory\bin\backupConfig.bat -username
  adminuser -password adminpassword -nostop
  ```

  In the WebSphere Application Server configurations for Tivoli Storage Productivity Center and Jazz for Service Management, *adminuser* is the user name that was used to install Tivoli Storage Productivity Center (for example, `db2admin`) or Jazz for Service Management (for example `smadmin`), and *adminpassword* is the password that is associated with *adminuser*.

  Tip: If you completed a default installation on the Windows operating system:

  – The Tivoli Storage Productivity Center WebSphere Application Server directory is located here:

  ```
  C:\Program Files\IBM\TPC\ewas\profiles\
  WebServerProfile\
  ```

  – The Jazz for Service Management WebSphere Application Server directory is located here:

  ```
  C:\Program Files\IBM\JazzSM\profile\
  ```

- On AIX or Linux operating systems, run the following command:

  ```
  WebSphere_Directory/bin/backupConfig.sh -username
  adminuser -password adminpassword -nostop
  ```

  In the WebSphere Application Server configurations for Tivoli Storage Productivity Center and Jazz for Service Management, *adminuser* is the user name that was used to install Tivoli Storage Productivity Center (for example, `db2inst1`) or Jazz for Service Management (for example `smadmin`), and *adminpassword* is the password that is associated with *adminuser*.

  Tip: If you completed a default installation on the AIX or Linux operating system:

  – The Tivoli Storage Productivity Center WebSphere Application Server directory is located here:

  ```
  /opt/IBM/TPC/ewas/profiles/WebServerProfile/
  ```

  – Jazz for Service Management WebSphere Application Server directory is located here:

  ```
  /opt/IBM/JazzSM/profile/
  ```

b. Back up the `soap.client.props` file for the Tivoli Storage Productivity Center instance of WebSphere Application Server.

- On the Windows operating system, run the **copy** command to back up this file:

  ```
  WebSphere_Directory\properties\soap.client.props
  ```

  Tip: If you completed a default installation on the Windows operating system, the Tivoli Storage Productivity Center Version 5.2 `soap.client.props` file is in this directory:

  ```
  C:\Program Files\IBM\TPC\ewas\profiles\WebServerProfile
  \properties\
  ```

- On the AIX or Linux operating systems, run the **cp** command to back up this file:

  ```
  WebSphere_Directory/properties/soap.client.props
  ```

Tip: If you completed a default installation on the AIX or Linux operating systems, the Tivoli Storage Productivity Center Version 5.2 `soap.client.props` file is in this directory:

```
/opt/IBM/TPC/ewas/profiles/WebServerProfile/properties
```

c. Back up the `IBM Cognos.ear` directory for the Jazz for Service Management WebSphere Application Server.

- On the Windows operating system, run the **xcopy** command to back up this directory:

```
WebSphere_Directory\installedApps\JazzSMNode01Cell
\IBM Cognos.ear
```

Tip: If you completed a default installation on the Windows operating system, the Jazz for Service Management WebSphere Application Server directory is located here:

```
C:\Program Files\IBM\JazzSM\profile\
```

- On the AIX or Linux operating systems, run the **cp** command to back up this directory:

```
WebSphere_Directory/installedApps/JazzSMNode01Cell
/IBM Cognos.ear
```

Tip: If you completed a default installation on the AIX or Linux operating systems, the Jazz for Service Management WebSphere Application Server directory is located here:

```
/opt/IBM/JazzSM/profile/
```

Tip: The default cell name `JazzSMNode01Cell` is used in the sample commands. If you did not use the default cell name when you installed Jazz for Service Management, you must specify the correct cell name in this step.

2. In the event of a problem when you remove an LDAP repository, complete the following steps to restore the items you backed up in step 1:

a. Run the `restoreConfig.sh` or the `restoreConfig.bat` command to restore the WebSphere Application Server configuration for the Tivoli Storage Productivity Center instance of WebSphere Application Server and restore the WebSphere Application Server configuration for the Jazz for Service Management instance of WebSphere Application Server. The **restoreConfig** command is in one of the following directories:

- On the Windows operating system, go to this directory:

```
WebSphere_directory/bin/restoreConfig.bat
```

Tip: If you completed a default installation on the Windows operating system:

  – The Tivoli Storage Productivity Center WebSphere Application Server directory is located here:

```
C:\Program Files\IBM\TPC\ewas\profiles\WebServerProfile\
```

  – The Jazz for Service Management WebSphere Application Server directory is located here:

```
C:\Program Files\IBM\JazzSM\profile\
```

- On the AIX or Linux operating system, go to this directory:

```
WebSphere_directory/bin/restoreConfig.sh
```

Tip: If you completed a default installation on the AIX or Linux operating system, the Tivoli Storage Productivity Center Version 5.2 `soap.client.propsis` in this directory:

`/opt/IBM/TPC/ewas/profiles/WebServerProfile/properties`

b. Restore the `soap.client.props` file for the Tivoli Storage Productivity Center instance of WebSphere Application Server.
   - On the Windows operating system, run the **copy** command to restore this file:

     WebSphere_Directory`\properties\soap.client.props`

     Tip: If you completed a default installation on the Windows operating system, the Tivoli Storage Productivity Center Version 5.2 `soap.client.props` is in this directory:

     `C:\Program Files\IBM\TPC\ewas\profiles\WebServerProfile\properties\`
   - On the AIX or Linux operating system, run the **cp** command to restore this file:

     WebSphere_Directory`/properties/soap.client.props`

     Tip: If you completed a default installation on the AIX or Linux operating systems, the Tivoli Storage Productivity Center Version 5.2 `soap.client.props` is in this directory:

     `/opt/IBM/TPC/ewas/profiles/WebServerProfile/`

c. Restore the `IBM Cognos.ear` directory for the Jazz for Service Management WebSphere Application Server.
   - On the Windows operating system, run the **xcopy** command to restore this directory:

     WebSphere_Directory`\installedApps\JazzSMNode01Cell\IBM Cognos.ear`

     Tip: If you completed a default installation on the Windows operating system, the Jazz for Service Management WebSphere Application Server directory is located here:

     `C:\Program Files\IBM\JazzSM\profile\`
   - On the AIX or Linux operating systems, run the **cp** command to restore this directory:

     WebSphere_Directory`/installedApps/JazzSMNode01Cell/IBM Cognos.ear`

     Note: The default cell name `JazzSMNode01Cell` is used in the sample commands. If you did not use the default cell name when you installed Jazz for Service Management, you must specify the correct cell name in this step.

     Tip: If you completed a default installation on the AIX or Linux operating systems, the Jazz for Service Management WebSphere Application Server directory is located here:

     `/opt/IBM/JazzSM/profile/`

d. Restart the Tivoli Storage Productivity Center web server and the Jazz for Service Management server. The execution of the **restoreConfig** command stops those WebSphere Application Server instances.

   For information about starting the Tivoli Storage Productivity Center web server or Jazz for Service Management server, see the Tivoli Storage Productivity Center information center. Search for *Starting the Tivoli Storage Productivity Center services*.

3. To open the WebSphere Integrated Solutions Console, open a web browser and enter one of the following addresses in the address field:
   - `http://hostname:port/ibm/console/logon.jsp`
   - `http://hostname:port/ibm/console/logon.jsp`

   The hostname is the server that is running WebSphere Application Server, such as the server name or IP address, and port is the port number for the WebSphere Application Server. The port number differs depending on which protocol you used (http or https) and the options that you selected when you installed Tivoli Storage Productivity Center.

   To determine the port number, complete the following steps:

   a. Open the *WebSphere_directory*`/properties/portdef.props` file.
   b. Determine the port number by looking for the value that is assigned to one of the following keys:
      - For `http:// protocols`, search for the `WC_adminhost` key
      - For `https:// protocols`, search for the `WC_adminhost_secure` key
4. On the WebSphere Integrated Solutions Console navigation tree, click **Security** > **Global Security**.

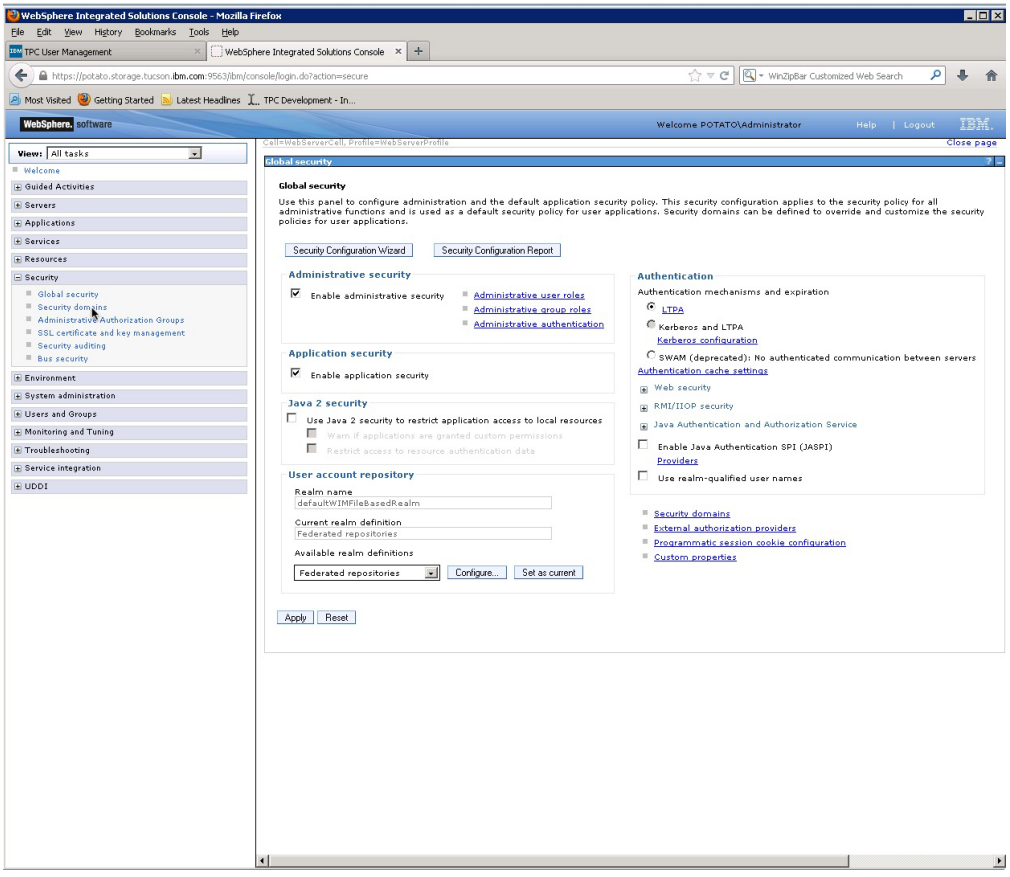

Figure 8. Opening the Global security page

5. On the Global security page, in the **User account repository** section, click **Configure** next to the **Available realm definitions** menu.

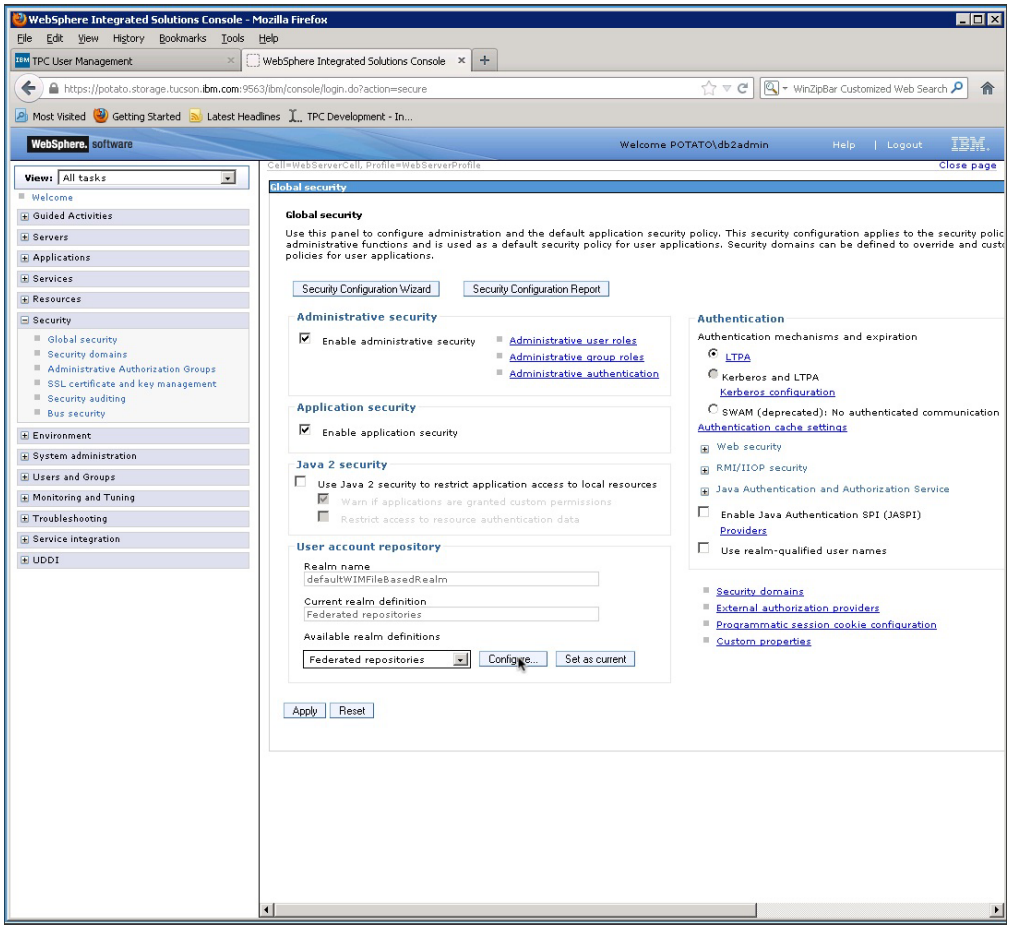

Figure 9. Configuring the available realm definitions

6. On the **Global security** > **Federated repositories** page, in the **Repositories in the realm** table, select the entry for the LDAP repository you previously added, and click **Remove**.

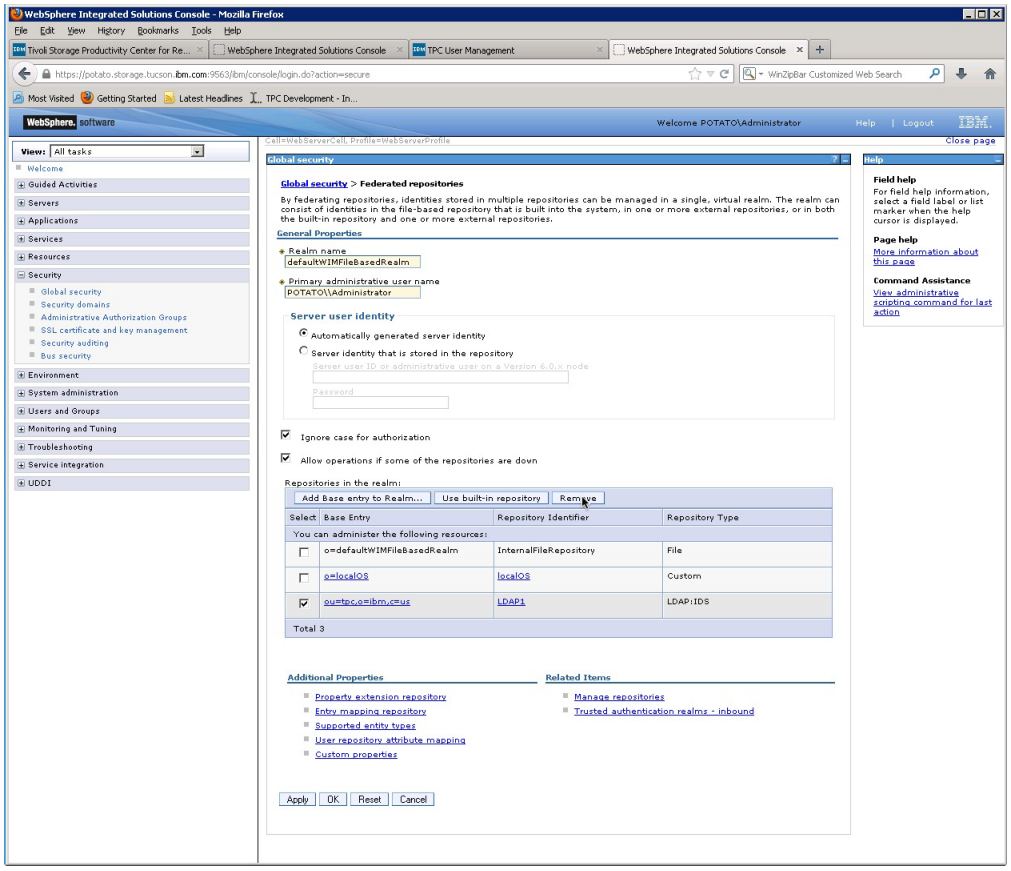

Figure 10. Removing the LDAP repository

7. In the message dialog box that is displayed on the Federated repositories page, click the **Save** link in **Save directly to the master configuration**.
8. On the **Global security** > **Federated repositories** page, under **Related items**, click **Manage repositories**.

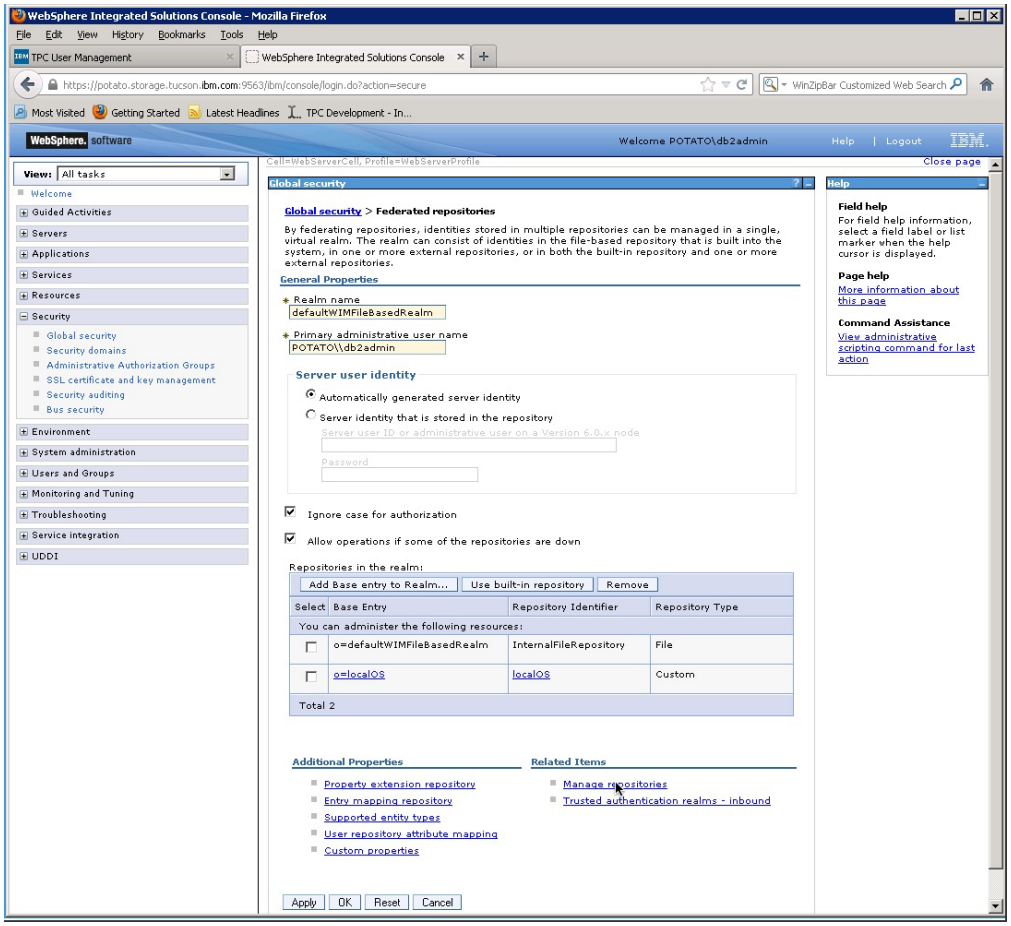

Figure 11. Managing the repositories

9. On the **Global security** > **Federated repositories** > **Manage repositories** page, select the entry for the LDAP repository that you previously added and click **Delete**.

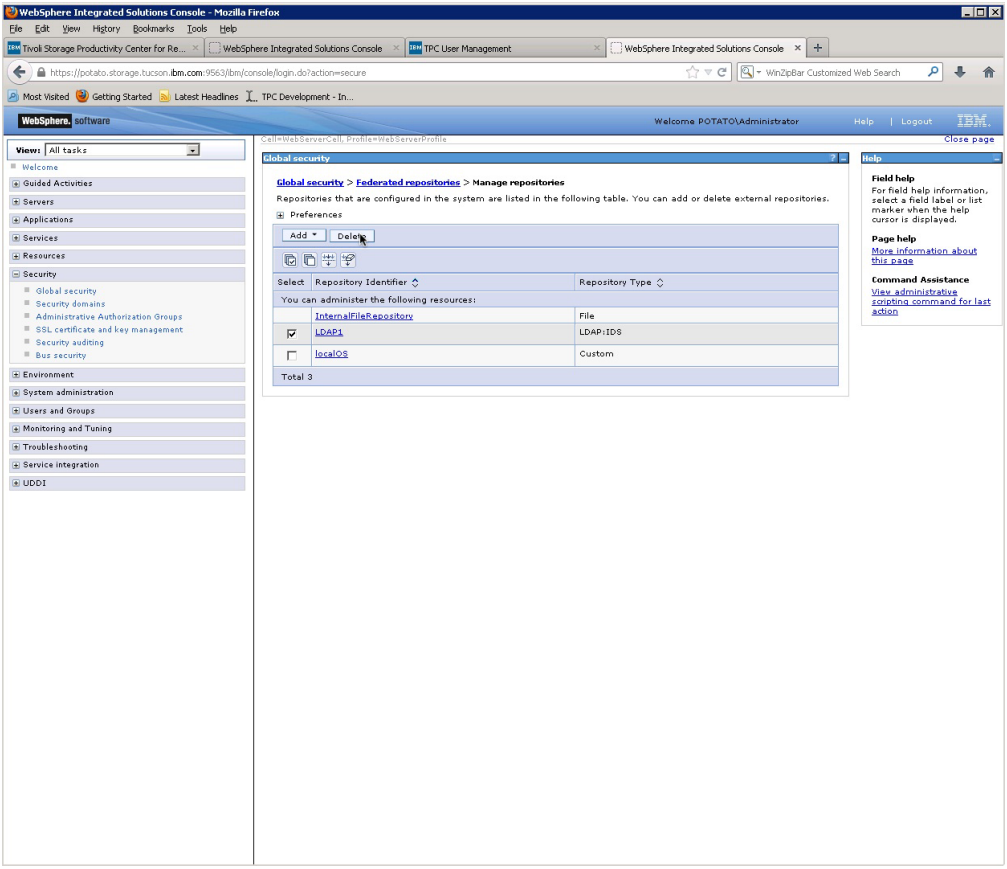

Figure 12. Deleting the LDAP repository

10. In the message dialog box that is displayed on the Manage repositories page, click the **Save** link in **Save directly to the master configuration**.
11. Log out from the WebSphere Integrated Solutions Console.
12. Log out of the web-based GUI.
13. Stop and restart the Tivoli Storage Productivity Center web server or Jazz for Service Management server. For information about stopping and starting the server, go to the Tivoli Storage Productivity Center information center. Search for *Starting and stopping Tivoli Storage Productivity Center services*.

    Important: Before you establish the authorization configuration for the WebSphere Integrated Solutions Console, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication, you must ensure that there are no duplicated user names or group names in the local file-based repository and the localOS repository.

    After you remove the LDAP repository from the Tivoli Storage Productivity Center federated repositories authentication configuration, you must establish the authorization configuration for the WebSphere Integrated Solutions Console.

    You should also assign Tivoli Storage Productivity Center roles to groups by using the **Role-to-Group Mapping** node in the stand-alone GUI.

    For information about assigning roles to users and roles to groups for Tivoli Storage Productivity Center for Replication, go to the Tivoli Storage Productivity Center information center. Search for *Managing replication*. On the Managing replication page, click **Security**.

Changing the LTPA token expiration for single sign-on

For single sign-on, the user authentication information is passed between applications by using Lightweight Third-Party Authentication (LTPA) tokens. By default, the tokens expire after eight hours. You can change the LTPA token expiration time by using the WebSphere Integrated Solutions Console.

To change the expiration time for the LTPA tokens, complete the following steps:

1. To start WebSphere Integrated Solutions Console, open a web browser and enter one of the following addresses in the address bar:
   - `http://hostname:port/ibm/console/logon.jsp`
   - `https://hostname:port/ibm/console/logon.jsp`

   Where hostname is the server that is running IBM WebSphere Application Server, such as the server name or IP address, and `port` is the port number for WebSphere Application Server.

   The port number differs depending on the protocol that you use (http or https) and the options that you selected when you installed Tivoli Storage Productivity Center.

   Determine the port number by looking for the value that is assigned to one of the following keys:

   a. Open the *WebSphere_directory*`/properties/portdef.props` file.

      Here are the locations for the WebSphere Application Server directory for each instance:

      - This is the location of the Jazz for Service Management WebSphere Application Server directory:

        JAZZSM_INSTALL_DIR`/profile`
      - This is the location of the Tivoli Storage Productivity Center WebSphere Application Server, or web server, directory:

        TPC_INSTALL_DIR`/ewas/profiles/WebServerProfile`

   b. Search for the value that is assigned to one of the following values:
      - For `http://` protocols, search for the `WC_adminhost` key
      - For `https://` protocols, search for the `WC_adminhost_secure` key
2. Log in to the WebSphere Integrated Solutions Console by using the appropriate user ID and password. Your user ID must have WebSphere Integrated Solutions Console administrator authorization.
3. In the WebSphere Integrated Solutions Console navigation tree, click **Security** > **Global security**.
4. On the Global security page, under **Authentication** and **Authentication mechanisms and expiration**, ensure that **LTPA** is selected and click the **LTPA** link.
5. On the **Global security** > **LTPA** page, under **LTPA Timeout**, type the expiration time in minutes in the **LTPA Timeout value for forwarded credentials between servers** field and click **OK**. The default expiration time is 480 minutes (eight hours). The expiration time must be greater than the authentication cache timeout value. The authentication cache timeout value can be found by clicking the **Authentication cache settings** link in the **Authentication** section of the Global security page. The default authentication cache timeout value is 10 minutes.
6. In the messages dialog box that is displayed on the Global security page, click the **Save** link in **Save directly to the master configuration**.
7. Log out from the WebSphere Integrated Solutions Console.

8. Stop and restart the WebSphere Application Server server. For information about stopping and starting the server, go to the Tivoli Storage Productivity Center information center. Search for *Starting and stopping the Tivoli Storage Productivity Center services*.

Checking Tivoli Storage Productivity Center status

This section provides information on checking the status of Tivoli Storage Productivity Center.

Checking the status of the Data Server

This topic describes how to check the status of the Data server.

To check the status of the Data Server, complete the following steps:

1. Open the IBM Tivoli Storage Productivity Center graphical user interface (GUI).
2. In the navigation tree, click **Administrative Services** > **Services** > **Data Server**. A status icon is displayed in the navigation tree for each service.

Checking the status of the Device server

This topic describes how to check the status of the Device server.

To check the status of the Device server, complete the following steps:

1. Open the IBM Tivoli Storage Productivity Center graphical user interface (GUI).
2. In the navigation tree, click **Administrative Services** > **Services** > **Device server**. A status icon is displayed in the navigation tree for each service.

Checking the status of agents

Check the status of agents.

In the navigation tree, expand **Administrative Services** > **Data Sources** and click one of the following nodes:

- **CIMOM Agents**
- **Storage Resource Agents**
- **Out of Band Fabric Agents**

In the content pane, a list of agents is displayed. If the agents are up and running, a green health status icon is shown in the State column. For CIM agents, a green health status icon is shown in the Connection Status column.

If the status icon is red for a Storage Resource agent, select that agent and click **View Log** to help determine the cause of the problem. If the reason that an agent is down is a connectivity problem, you might not be able to view the log files. A CIM agent has a red (critical) icon status until you enter valid credentials.

Increasing the memory allocation for the Data server

If the data memory that is allocated for your Data server is insufficient, you can increase the memory. The default maximum memory for the Data server is set to 1024 MB.

You cannot increase the memory for the Device server. The memory for the Device server is set to the maximum heap size for the JVM.

Increasing the memory allocation for the Data server that is running on AIX

Increase the memory allocation for the Data server that is running on AIX.

To increase the memory that is allocated for the Data server, complete the following steps:

1. Log on as a user with root authority.
2. Stop the Data server. From the command line, run the following command:
   ```
   /TPC_install_directory/scripts/stopTPCData.sh
   ```
   Where *TPC_install_directory* is the installation directory. The default directory is `/opt/IBM/TPC`.
3. Using a text editor, open the `/TPC_install_directory/data/server/tpcdsrv1` file.
4. Locate the following line:
   ```
   exec $JAVAEXE -Dsun.net.inetaddr.ttl=300 -Xrs -XmxXXXXm
   -cp $CLASSPATH com.tivoli.itsrm.server.Server &
   ```
   where *XXXX* is the memory allocated for the Data server. The default is `1024m` (1024 MB).
5. Increase the memory that is allocated for the Data server. For example, to increase the memory to 1536 MB, change the line to read as follows:
   ```
   exec $JAVAEXE -Dsun.net.inetaddr.ttl=300 -Xrs -Xmx1536m
   -cp $CLASSPATH com.tivoli.itsrm.server.Server &
   ```
6. Save the modified `tpcdsrv1` file.
7. Start the Data server by running the following command:
   ```
   /TPC_install_directory/scripts/startTPCData.sh
   ```

Increasing the memory allocation for the Data server that is running on Linux

Increase the memory allocation for the Data server that is running on Linux.

To increase the memory that is allocated for the Data server, complete the following steps:

1. Log on as a user with root authority.
2. Stop the Data server. From the command line, run the following command:
   ```
   /TPC_install_directory/scripts/stopTPCData.sh
   ```
   Where *TPC_install_directory* is the installation directory. The default directory is `/opt/IBM/TPC`.
3. Using a text editor, open the `/TPC_install_directory/data/server/tpcdsrv1` file.
4. Locate the following line:
   ```
   exec $JAVAEXE -Dsun.net.inetaddr.ttl=300 -Xrs -XmxXXXXm
   -cp $CLASSPATH com.tivoli.itsrm.server.Server &
   ```
   Where *XXXX* is the memory that is allocated for the Data server. The default is `1024m` (1024 MB).
5. Increase the memory that is allocated for the Data server. For example, to increase the memory to 1536 MB, change the line to read as follows:
   ```
   exec $JAVAEXE -Dsun.net.inetaddr.ttl=300 -Xrs -Xmx1536m
   -cp $CLASSPATH com.tivoli.itsrm.server.Server &
   ```
6. Save the modified `tpcdsrv1` file.
7. Start the Data server by running the following command:
   ```
   /TPC_install_directory/scripts/startTPCData.sh
   ```

Increasing memory allocation for Data server that is running on Windows

Increase the memory allocation for the Data server that is running on Windows.

To increase the memory that is allocated for the Data server, complete the following steps:

1. Open the Run window. For information about how to open the Run window, see "Accessing administration tools" on page 1735.
2. Type `regedit` and click **OK**. The Registry Editor window is displayed.
3. Expand the **HKEY_LOCAL_MACHINE** > **SOFTWARE** > **Wow6432Node** > **IBM** > **TSRM** > **1** in the Registry Editor window.
4. Right-click the 1 folder and click **New** > **String Value**.
5. Type `SRVJPARMS` as the name of the string.
6. Right-click the name of the string and click **Modify**.
7. Enter`-Xmx`*`XXXX`*`m` in the **Value data** field, where *XXXX* represents the number of megabytes for the server maximum heap size. Click **OK**. The default size is 1024 MB. The largest possible value for the maximum heap size is 1536 MB. If the value is set to something larger than 1536, that value is ignored and 1536 MB is used as the maximum heap size.
8. Stop and restart the server to have the changes take effect. To stop the server, complete the following steps:
   a. Open Windows Services. For information about how to open Services, see "Accessing administration tools" on page 1735.
   b. Right-click **IBM Tivoli Storage Productivity Center - Data Server** and click **Stop**. To restart the server, right-click the service and click **Start**.

Setting timeout values for the Device server

If a probe or discovery of a storage subsystem times out before the operation completes, you can increase the timeout values for the Device server.

If a probe or discovery of a storage subsystem times out before the operation completes, you receive the following error message:

```
HWN021650E Encountered timeout while connecting to CIMOM IP:port.
Check the CIMOM or increase timeout value.
```

where *IP* is the IP address, and *port* is the port number. If you determine that the Common Information Model Object Manager (CIMOM) is not the cause of the problem, you can use the command-line interface (CLI) to increase the timeout values for the Device server.

For those storage systems that use native interfaces to connect to Tivoli Storage Productivity Center you see this error message:

```
HWN020103E The external process exceeded the timeout limit and was cancelled.
```

The following storage systems use native interfaces to connect to Tivoli Storage Productivity Center:

- System Storage DS8000
- SAN Volume Controller
- The XIV system
- Storwize V3500
- Storwize V3700

- Storwize V7000
- Storwize V7000 Unified
- IBM SONAS

1. Run the **getdscfg** command to determine the current values of the timeout properties. From the command prompt, enter the following command:

   ```
   cli>tpctool getdscfg -user user -pwd password -url host:port
    -property timeout_property
   ```

   where:

   - *user* is an IBM Tivoli Storage Productivity Center user ID.
   - *password* is the password for the Tivoli Storage Productivity Center user ID.
   - *host* is the host name or IP address, and *port* is a valid port number for the HTTP service of the Device server. The default value for *port* is typically 9550.
   - *timeout_property* is one of the following strings:
     - http.timeout
     - CIMClientWrapper.Timeout
     - Probe.Timeout.Array
     - Probe.Timeout.LMM
     - Discovery.Timeout
     - CIMOMManager.TestConnectionTimeout

     Important: Timeout properties are displayed in milliseconds. If the value is **0** (zero), it means that there is no timeout.

     For the storage systems that use the native interface, the *timeout_property* strings are:
     - NAPI.Timeout.TestConnection
     - NAPI.Timeout.Probe
     - NAPI.Timeout.EventPoll
2. Run the **setdscfg** command to increase the timeout value. Run the following command:

   ```
   cli>tpctool setdscfg -user user -pwd password -url host:port
   -property timeout_property timeout_value
   ```

For more information about **tpctool**, see the *Tivoli Storage Productivity Center Command-Line Interface Reference*. You also can view help from the command line by issuing the command with the **-help** option.

Changing passwords

This section provides information about changing passwords.

If you installed IBM Tivoli Storage Productivity Center using the typical installation and used the same DB2 user ID and password for the items Tivoli Storage Productivity Center requires, then when you change the DB2 password, you must also change the passwords for the items that the DB2 password applies to. For example, the DB2 administrative password might also apply to the following items:

- Database administration user ID and password (for the Data Server or Device server to connect to the database)
- Database user ID and password to create the database schema

- Host authentication password (for the Storage Resource agents to communicate with the Device server)
- Storage Resource agent service login user ID and password (for Windows only, if this user ID does not exist)
- WebSphere administration user ID and password (for the Device server to communicate with embedded WebSphere if the user ID does not exist). This is the case only if you select OS authentication during the Tivoli Storage Productivity Center installation. If you chose LDAP authentication, then the WebSphere administration user ID and password are set to the LDAP TPC Administrator username and password.

Changing passwords by using the password tool

Use the password tool to change the passwords for DB2, Tivoli Storage Productivity Center, and Jazz for Service Management so that they can continue to authenticate to one another whenever you change a password.

If you are logged on to Tivoli Storage Productivity Center by using a domain user account, which is also a member of the local administrator group, when you run the change password tool, passwords are not updated. If you run the tool by using a local OS user account, and an error occurs, complete these steps.

To run the password tool when you log in by using a domain user account, choose one of the following steps:

- Right-click the `changepasswords.bat` file and select **Run as administrator**.
- Complete the following steps:
  1. Click **Start** > **All Programs** > **Open Administrative Tools** > **Local Security Policy**.
  2. On the Local Security Policy window, disable **User Account Control: Run all administrators in Admin Approval Mode**.
  3. Restart your computer.

Single server installation where components use the same logon credentials:

Use the password tool to change the password for DB2, Tivoli Storage Productivity Center, and Jazz for Service Management when these products are installed on the same server and use the same user name and password.

Before you use the password tool, ensure that you know the existing password or passwords that you want to change.

1. Complete the following steps:

   For Windows operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

   a. Open the Control Panel.
   b. Double-click **User Accounts**.
   c. Click **Change your password**.
   d. Enter the current password, enter, and confirm the new password.
   e. Click **Change password** and close the User Accounts window.

   For AIX or Linux operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

   Log in with the DB2 user ID, and run the **passwd** command.

Tip: To change other user credentials, run the **passwd** *username* command.

If you need to change the LDAP user password, see your LDAP server documentation.

2. Open a command prompt window and change the directory to the following default directory:

   On Windows operating systems:
   ```
   C:\Program Files\IBM\TPC\service
   ```

   On Linux or UNIX operating systems:
   ```
   /opt/IBM/TPC/service
   ```

3. Start the password tool by running the following program:

   On Windows operating systems:
   ```
   changepasswords.bat
   ```

   On Linux or AIX operating systems:
   ```
   changepasswords.sh
   ```

4. In the"Choose the password to change" window, click **Change Jazz for Service Management password** and, click **OK**.
5. Complete the following steps:
   a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1 on page 178.
   b. In response to the following message, click **Yes**:
      ```
      Are you sure you want to change the password?
      ```
   c. When the process is completed, click **Back to Main**.
6. To change the password for Tivoli Storage Productivity Center, click **Change Tivoli Storage Productivity Center server and DB2 passwords** and click **OK**.
   a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1 on page 178. Click **OK**.
   b. In response to the following message, click **Yes**:
      ```
      Are you sure you want to change the password?
      ```

      Tip: After you click **Yes**, if the DB2 servers are already started, the following messages might be displayed in a command prompt window. Ignore them and continue with the procedure.
      ```
      Changing DB2 User password for Data Server...
      Changing DB2 User password for Device Server...
      ```
   c. When the process is completed, click **Back to Main**.

      Tip: To verify that the password changes were completed successfully, review the `PWTool.log` file that is in the *`TPC_installation_directory`*`\service` directory, where *TPC_installation_directory* is the directory where Tivoli Storage Productivity Center is installed.
7. Exit the password tool.

Related tasks:

"Single-server installation where components use different logon credentials" on page 180
Use the password tool to change the passwords for DB2, Tivoli Storage Productivity Center, and Jazz for Service Management when these components are installed on the same server but use different user names and passwords.

"Multiple-server installation where Tivoli Storage Productivity Center reports are remote" on page 182
You can change the passwords for DB2, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center reports when these components are installed on two servers. Tivoli Storage Productivity Center reports are installed on *Server A*, and DB2 and the Tivoli Storage Productivity Center server are installed on *Server B*.

"Multiple-server installation where DB2 is remote" on page 184
Change the passwords for DB2, Tivoli Storage Productivity Center, and Jazz for Service Management when these components are installed on servers. Jazz for Service Management and Tivoli Storage Productivity Center are on *Server A* and DB2 is on *Server B*.

Single-server installation where components use different logon credentials:

Use the password tool to change the passwords for DB2, Tivoli Storage Productivity Center, and Jazz for Service Management when these components are installed on the same server but use different user names and passwords.

Before you use the password tool, ensure that you know the existing password or passwords that you want to change.

1. Complete the following steps:

   For Windows operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

   a. Open the Control Panel.
   b. Double-click **User Accounts**.
   c. Click **Change your password**.
   d. Enter the current password, enter, and confirm the new password.
   e. Click **Change password** and close the User Accounts window.

   For AIX or Linux operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

   When you log in by using DB2 user ID, run the **passwd** command. To change other user credentials, run the **passwd** *username* command.

   Tip: If you need to change the LDAP user password, see your LDAP server documentation.
2. Open a command prompt window and change the directory to the following default directory:

   On Windows operating systems:

   ```
   C:\Program Files\IBM\TPC\service
   ```

   On Linux or UNIX operating systems:

   ```
   /opt/IBM/TPC/service
   ```
3. Start the password tool by running the following program:

   On Windows operating systems:

   ```
   changepasswords.bat
   ```

   On Linux or AIX operating systems:

   ```
   changepasswords.sh
   ```
4. On the "Choose the password to change" window, click **Change Jazz for Service Management password** and, click **OK**.
5. To change the password, complete the following steps:

a. Enter, confirm the new password, and click **OK**.
b. In response to the following message, click **Yes**:
```
Are you sure you want to change the password?
```
c. When the process is completed, click **Back to Main**.

6. To change the password for Tivoli Storage Productivity Center, click **Change Tivoli Storage Productivity Center server passwords** and click **OK**.
7. Complete the following steps:
   a. Enter and confirm the new password and click **OK**.
   b. In response to the following message, click **Yes**:
   ```
   Are you sure you want to change the password?
   ```
   c. When the process is completed, click **Back to Main**.
8. To change the password for DB2, click **Change DB2 password** and click **OK**.
9. Complete the following steps:
   a. Enter and confirm the new password.
   b. Ensure that the **Restart servers** check box is selected and click **OK**.
   c. In response to the following message, click **Yes**:
   ```
   Are you sure you want to change the password?
   ```

   Tip: After you click **Yes**, if the DB2 servers are already started, the following errors might be displayed in a command prompt window. Ignore them and continue with the procedure.
   ```
   Changing DB2 User password for Data Server...
   Changing DB2 User password for Device Server...
   ```
   d. When the process is completed, click **Back to Main**.

   Tip: To verify that the password changes were completed successfully, review the `PWTool.log` file that is in the *`TPC_installation_directory`*`\service` directory, where *TPC_installation_directory* is the directory where Tivoli Storage Productivity Center is installed.

Related tasks:

"Single server installation where components use the same logon credentials" on page 178
Use the password tool to change the password for DB2, Tivoli Storage Productivity Center, and Jazz for Service Management when these products are installed on the same server and use the same user name and password.

"Multiple-server installation where Tivoli Storage Productivity Center reports are remote" on page 182
You can change the passwords for DB2, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center reports when these components are installed on two servers. Tivoli Storage Productivity Center reports are installed on *Server A*, and DB2 and the Tivoli Storage Productivity Center server are installed on *Server B*.

"Multiple-server installation where DB2 is remote" on page 184
Change the passwords for DB2, Tivoli Storage Productivity Center, and Jazz for Service Management when these components are installed on servers. Jazz for Service Management and Tivoli Storage Productivity Center are on *Server A* and DB2 is on *Server B*.

Multiple-server installation where Tivoli Storage Productivity Center reports are remote:

You can change the passwords for DB2, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center reports when these components are installed on two servers. Tivoli Storage Productivity Center reports are installed on *Server A*, and DB2 and the Tivoli Storage Productivity Center server are installed on *Server B*.

Before you use the password tool, ensure that you know the existing password or passwords that you want to change.

Important: You might need to update the Tivoli Common Reporting Cognos Content Store database password, in a multiple-server installation, where the Cognos Content Store database is installed on the different DB2 server instance to the Tivoli Storage Productivity Center database repository. Use the IBM Cognos Configuration tool to update the Cognos Content Store database password. This tool is found in the following directory:

- On Windows operating systems:
  JazzSM_INST_DIR\reporting\cognos\bin64\tcr_cogconfig.bat
- On AIX or Linux operating systems:
  JazzSM_INST_DIR/reporting/cognos/bin64/tcr_cogconfig.sh

After you start the tool, on the "IBM DB2 Content Store - Database - Resource Properties" window, click the pencil icon to change the password and then click **Save**.

After you run this tool, you must restart Jazz for Service Management.

Tivoli Storage Productivity Center is configured for OS authentication.

1. On the server where DB2 and Tivoli Storage Productivity Center are installed, complete the following steps:

   On Windows operating systems:

   a. Click **Start** > **Control Panel**.
   b. Double-click **User Accounts**.
   c. Click **Change your password**.
   d. Enter the current password, and enter the password again to confirm it.
   e. Click **Change password**.
   f. Close the User Accounts window.

   On AIX or Linux operating systems, if you logged in with the DB2 user ID, to change the DB2 and Tivoli Storage Productivity Center passwords, run the following command:

   `passwd username`

   To change other user credentials, run the *passwd* username command.

   Tip: If you need to change the LDAP user password, see your LDAP server documentation.
2. On *Server A* and *Server B*, open a command prompt window and go to the following default directory:
   - On Windows operating systems:
     C:\Program Files\IBM\TPC\service
   - On AIX or Linux operating systems:

```
/opt/IBM/TPC/service
```

3. On *Server A* and *Server B*, start the password tool by running the following program:
   - On Windows operating systems:
     ```
     changepasswords.bat
     ```
   - On the AIX or Linux operating systems:
     ```
     changepasswords.sh
     ```
4. On *Server B*, where Tivoli Storage Productivity Center and DB2 are installed, in the "Choose the password to change" window, click **Change Tivoli Storage Productivity Center server passwords** and click **OK**.
5. Complete the following steps:
   a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1 on page 182 and, click **OK**.
   b. In response to the following message, click **Yes**:
      ```
      Are you sure you want to change the password?
      ```
   c. When the process is completed, click **Back to Main**.
6. On *Server B*, click **Change DB2 password**, click **OK**.
7. Complete the following steps:
   a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1 on page 182, click **OK**.
   b. In response to the following message, click **Yes**:
      ```
      Are you sure you want to change the password?
      ```

      Tip: After you click **Yes**, if the DB2 servers are already started, the following messages might be displayed in a command prompt window.
      ```
      Changing DB2 User password for Data Server...
      Changing DB2 User password for Device Server...
      ```

      You can ignore these messages and continue with the procedure.
   c. When the process is completed, click **Back to Main**.
   d. Exit the password tool.

      Tip: To verify that the password changes were completed successfully, review the `PWTool.log` file that is in the *`TPC_installation_directory`*`\service` directory, where *TPC_installation_directory* is the directory in which Tivoli Storage Productivity Center is installed.
8. On *Server A*, in the "Choose the password to change" window, click **Change Jazz for Service Management password**, and click **OK**.
9. Complete the following steps:
   a. Enter and confirm the new password and click **OK**.
   b. In response to the following message, click **Yes**:
      ```
      Are you sure you want to change the password?
      ```
   c. When the process is completed, click **Back to Main**.
10. On *Server A*, in the "Choose the password to change" window, click **Change DB2 password** and click **OK**. Complete the following steps:
    a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1 on page 182.
    b. Ensure that the **Restart servers** check box is selected and click **OK**.
    c. In response to the following message, click **Yes**:

```
Are you sure you want to change the password?
```

Tip: After you click **Yes**, if the DB2 servers are already started, the following messages might be displayed in a command prompt window. Ignore them and continue with the procedure.

```
Updating Tivoli Storage Productivity Center Reports
configuration...
Restarting the Jazz for Service Management
server...
```

d. When the process is completed, click **Back to Main**.

Related tasks:

"Single server installation where components use the same logon credentials" on page 178
Use the password tool to change the password for DB2, Tivoli Storage Productivity Center, and Jazz for Service Management when these products are installed on the same server and use the same user name and password.

"Single-server installation where components use different logon credentials" on page 180
Use the password tool to change the passwords for DB2, Tivoli Storage Productivity Center, and Jazz for Service Management when these components are installed on the same server but use different user names and passwords.

"Multiple-server installation where DB2 is remote"
Change the passwords for DB2, Tivoli Storage Productivity Center, and Jazz for Service Management when these components are installed on servers. Jazz for Service Management and Tivoli Storage Productivity Center are on *Server A* and DB2 is on *Server B*.

Multiple-server installation where DB2 is remote:

Change the passwords for DB2, Tivoli Storage Productivity Center, and Jazz for Service Management when these components are installed on servers. Jazz for Service Management and Tivoli Storage Productivity Center are on *Server A* and DB2 is on *Server B*.

Before you use the password tool, ensure that you know the existing password or passwords that you want to change.

Important: You might need to update the Tivoli Common Reporting Cognos Content Store database password, in a multiple-server installation, where the Cognos Content Store database is installed on the different DB2 server instance to the Tivoli Storage Productivity Center database repository. Use the IBM Cognos Configuration tool to update the Cognos Content Store database password. This tool is found in the following directory:

- On Windows operating systems:

  JazzSM_INST_DIR\reporting\cognos\bin64\tcr_cogconfig.bat
- On AIX or Linux operating systems:

  JazzSM_INST_DIR/reporting/cognos/bin64/tcr_cogconfig.sh

After you start the tool, on the "IBM DB2 Content Store - Database - Resource Properties" window, click the pencil icon to change the password and then click **Save**.

After you run this tool, you must restart Jazz for Service Management.

Tivoli Storage Productivity Center is configured for OS authentication.

1. On *Server B*, where DB2 is installed, complete the following steps:

   For Windows operating system:

   a. Open the Control Panel.
   b. Double-click **User Accounts**.
   c. Click **Change your password**.
   d. Enter the current password, and enter and confirm the new password.
   e. Click **Change password**.
   f. Close the User Accounts window.

   For AIX or Linux operating system: when you are logged in with the DB2 user ID, run the **passwd** command. To change other user credentials, run the `passwd username` command.

   Tip: If you need to change the LDAP user password, see your LDAP server documentation.

2. On *Server A*, where Tivoli Storage Productivity Center and Jazz for Service Management are installed, complete the following steps:

   For Windows operating system:

   a. Open the Control Panel.
   b. Double-click **User Accounts**.
   c. Click **Change your password**.
   d. Enter the current password, and enter the password again to confirm it. The new password can be different from the password that you created in step 1.
   e. Click **Change password** and close the User Accounts window.

   For AIX or Linux operating systems, when you are logged in with the DB2 user ID, run the **passwd** command. To change other user credentials, run the **passwd** *username* command.

3. On *Server A* and *Server B*, open a command prompt window and go to the following default directory:

   - For Windows operating system:

     ```
     C:\Program Files\IBM\TPC\service
     ```

   - For Linux or AIX operating system:

     ```
     /opt/IBM/TPC/service
     ```

4. On *Server A* and *Server B*, start the password tool by running the following program:

   - For Windows operating systems:

     ```
     changepasswords.sh
     ```

   - For Linux or AIX operating system:

     ```
     changepasswords.bat
     ```

5. On *Server B*, where DB2 is installed, in the "Choose the password to change" window, click Change DB2 password, and click **OK**. Complete the following steps:

   a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1.
   b. Ensure that the **Restart servers** check box is selected and click **OK**.
   c. In response to the following message, click **Yes**:

      ```
      Are you sure you want to change the password?
      ```

Tip: After you click **Yes**, if the DB2 servers are already started, the following messages might be displayed in a command prompt window:

```
Change DB2 user password for Data server...
Changing DB2 user password for Device server...
```

You can ignore these messages and continue with the procedure.

d. When the process is completed, click **Back to Main**.

6. On Server B, click **Exit program** to exit the password tool.
7. On *Server A*, in the "Choose the password to change" window, click **Change Jazz for Service Management** and click **OK**. Complete the following steps:
   a. Enter and confirm the new password, ensure that it is the same as the new password that you entered in step 2 on page 185, and click **OK**.
   b. In response to the following message, click **Yes**:

      ```
      Are you sure you want to change the password?
      ```

   c. When the process is completed, click **Back to Main**.
8. On *Server A*, where Jazz for Service Management and Tivoli Storage Productivity Center are installed, in the "Choose the password to change" window, click **Change DB2 password** and click **OK**. Complete the following steps:
   a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1 on page 185.
   b. Ensure that the **Restart servers** check box is selected and click **OK**.
   c. In response to the following message, click **Yes**:

      ```
      Are you sure you want to change the password?
      ```

      Tip: After you click **Yes**, if the DB2 servers are already started, the following errors might be displayed in a command prompt window:

      ```
      Changing DB2 User password for Data Server...
      Changing DB2 User password for Device Server...
      ```

      You can ignore these messages and continue with the procedure.

   d. When the process is completed, click **Back to Main**.
9. On *Server A*, click **Exit program** to exit the password tool.

Important: For the changes to take effect, you may need to manually restart the Tivoli Storage Productivity Center servers, Jazz for Service Management, or both of these components.

Tip: To verify that the password changes were completed successfully, review the `PWTool.log` file that is generated on *Server A* and *Server B*. For each server, this log file is in the *`TPC_installation_directory`*`\service` directory, where *TPC_installation_directory* is the directory in which Tivoli Storage Productivity Center is installed.

Related tasks:

"Single server installation where components use the same logon credentials" on page 178

Use the password tool to change the password for DB2, Tivoli Storage Productivity Center, and Jazz for Service Management when these products are installed on the same server and use the same user name and password.

"Single-server installation where components use different logon credentials" on page 180
Use the password tool to change the passwords for DB2, Tivoli Storage Productivity Center, and Jazz for Service Management when these components are installed on the same server but use different user names and passwords.

"Multiple-server installation where Tivoli Storage Productivity Center reports are remote" on page 182
You can change the passwords for DB2, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center reports when these components are installed on two servers. Tivoli Storage Productivity Center reports are installed on *Server A*, and DB2 and the Tivoli Storage Productivity Center server are installed on *Server B*.

Changing passwords by editing configuration files

Changing the DB2 user password for the Device server:

This section provides information about changing the DB2 user password for the Device server.

If you change the DB2 user password, follow these steps to change it for the Device server:

1. Go the *TPC_installation_dir*\cli directory and run the following command:
   ```
   tpctool encrypt password
   ```
   This encrypts the password.
2. Copy the encrypted password to the following file: *TPC_installation_dir*\device\conf\tsnmdbparms.properties Replace this password:
   ```
   tivoli.sanmgmt.jdbc.dbPassword=<encrypted_password>
   ```
3. Restart the Device server.

Changing the DB2 password for the Data server:

You can change the DB2 password for the Data server.

Context: The following directory locations are assumed for the Tivoli Storage Productivity Center directory:

- On the Windows operating system, C:\Program Files\IBM\TPC.
- On the AIX or Linux operating systems, /opt/IBM/TPC.

To change the DB2 password, complete the following steps:

1. From a command prompt, go to the directory where the repository.config file is located in the Tivoli Storage Productivity Center installation path.
   For example:
   - On the Windows operating system:
     ```
     cd c:\Program Files\IBM\TPC\data\config\
     ```
   - On the AIX or Linux operating systems:
     ```
     cd /opt/IBM/TPC/data/config/
     ```
2. Run the following command to update the password for the Data server.
   - On the Windows operating system:

```
"C:\Program Files\IBM\TPC\jre\bin\java" -classpath
"C:\Program Files\IBM\TPC\data\server\lib\
TPC_Common.jar" com.tivoli.itsrm.repository.Transform -p new_password
repository.config
```

- On the AIX or Linux operating systems:

```
/opt/IBM/TPC/jre/bin/java -classpath /opt/IBM/TPC/data/server/lib/
TPC_Common.jar com.tivoli.itsrm.repository.Transform -p new_password
repository.config
```

3. Stop and start the Data Server again.
4. Start the Tivoli Storage Productivity Center GUI.

Checking for a fully qualified host name

IBM Tivoli Storage Productivity Center requires fully qualified host names. Some machines might be configured to return a short host name, such as tpcserver1 instead of a fully qualified host name, such as tpcserver1.tpc.example.com. This topic provides information on how to check for a fully qualified host name.

Checking for a fully qualified host name for Windows systems

Verify the fully qualified host name on Windows operating systems.

1. Choose one of these options:

Option	Description
Windows Server 2012	1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**. 2. Click **Control Panel**, and then click **System**. 3. Click **Change Settings**, click **Change**, and then click **Change** again.
Windows 7, Windows 2008	1. Click **Start** > **Control Panel** > **System and Security**. 2. Click **System**, and then click **Change Settings**. 3. On the **Computer Name** tab, click **Change**.
Windows XP, Windows Vista	1. On the desktop, right-click **My Computer**, and then click **Properties**. 2. On the **Computer Name** tab, click **Change**. .

2. In the **Computer name** field, enter the fully qualified host name, and then click **More**.
3. Verify that the **Primary DNS suffix** field contains a domain name, and then click **OK**.

Checking for a fully qualified host name for AIX systems

This topic provides information on how to verify a fully qualified host name for AIX.

The default domain name search order is as follows:

1. Domain Name System (DNS) server

2. Network Information Service (NIS)
3. Local /etc/hosts file.

If the /etc/resolv.conf file does not exist, the /etc/hosts file is used. If only the /etc/hosts file is used, the fully qualified computer name must be the first one that is listed after the IP address.

Verify that the /etc/resolv.conf file exists and contains the appropriate information, such as:

```
domain mydivision.mycompany.com
nameserver 123.123.123.123
```

If NIS is installed, the /etc/irs.conf file overrides the system default. It contains the following information:

```
hosts = bind,local
```

The /etc/netsvc.conf file, if it exists, overrides the /etc/irs.conf file and the system default. It contains the following information:

```
hosts = bind,local
```

If the NSORDER environment variable is set, it overrides all of the preceding files. It contains the following information:

```
export NDORDER=bind,local
```

Checking for a fully qualified host name for Linux systems

This topic provides information on how to verify a fully qualified host name for Linux.

Linux uses a resolver library to obtain the IP address corresponding to a host name. The /etc/host.conf file specifies how names are resolved. The entries in the /etc/host.conf file tell the resolver library what services to use, and in what order, to resolve names. Edit the host.conf file using the **vi** editor to add the following lines:

```
# Lookup names through DNS first then fall back to /etc/hosts.
order bind,hosts
# Machines with multiple IP addresses.
multi on
# Check for IP address spoofing.
nospoof on
```

The **order** option indicates the order of services. The sample entry specifies that the resolver library should first consult the name server to resolve a name and then check the /etc/hosts file. It is recommended to set the resolver library to first check the name server, bind file, and then the hosts file (hosts) for better performance and security on all your servers. You must have the DNS and BIND software installed for this configuration to work.

The **multi** option determines whether a host in the /etc/hosts file can have multiple IP addresses. Hosts that have more than one IP address are said to be multihomed, because the presence of multiple IP addresses implies that the host has several network interfaces.

The **nospoof** option takes care of not permitting spoofing on this machine. IP-Spoofing is a security exploit that works by tricking computers into a trust relationship that you are someone that you really are not. In this type of attack, a machine is set up to look like a legitimate server and then issue connections and

other types of network activities to legitimize end systems, other servers, or large data repository systems. This option must be set ON for all types of servers.

Checking for a fully qualified host name for Oracle Solaris

This topic provides information about how to verify a fully qualified host name for Oracle Solaris systems.

Verify that the `/etc/resolv.conf` file exists and contains the appropriate information, such as:

```
domain mydivision.mycompany.com
nameserver 123.123.123.123
```

A short name is used if the `/etc/nsswitch.conf` file contains a line that begins as follows and if the `/etc/hosts` file contains the short name for the computer:

```
hosts: files
```

To correct this problem, follow these steps:

1. Change the line in the `/etc/nsswitch.conf` file to the following:
   ```
   hosts: dns nis files
   ```
2. Enter the following command to stop the **inet** service:
   ```
   /etc/init.d/inetsvc stop
   ```
3. Enter the following command to restart the **inet** service:
   ```
   /etc/init.d/inetsvc start
   ```

Granting local administrative privileges to a domain account

Automatically grant administrative privileges to Windows domain accounts. The user account for the Storage Resource agent requires local administrative rights. Because these rights are not necessarily guaranteed for domain users in a Windows domain environment, you are shown how to grant local administrative rights to domain users. Using this procedure, you do not have to manually process each machine in the domain.

Note: These steps are for a Windows system that is a member of a Windows domain and not for the Windows Domain Primary Domain Controller.

To use Group Policy to grant local administrative privileges to a domain account, complete the following steps:

1. On the domain controller, go to **Administrative Tools** > **Active Directory Users and Computers** (you must be running with Domain Administrator privileges).
2. Right-click on the Organizational Unit (OU) upon which you want to apply the Group Policy. Click **Ptoperties**.
3. The Group Policy Properties panel is displayed. Select the Group Policy tab and click **New** to create a Group Policy.
4. Designate a name for the new Group Policy. Select the new Group Policy and click **Edit**.
5. The Group Policy Object Editor panel is displayed. Go to **New Group Policy Object** *your_policy* > **Computer Configuration** > **Windows Settings** > **Security Settings** > **Restricted Groups**. Right-click **Restricted Groups**, and then click **Add Group**.
6. For example, name the new group `Administrators`. Under `Properties`, add the user `Administrator`, and the domain accounts or groups upon which you want the Group Policy in effect for. For example, you can add `TPC\storageadmin`, `TPC\storagegroup`, and `TPC\TestGroup`. Click **OK**.

7. Add these user rights to the domain account:
   - Act as part of the operating system
   - Log on as a service

   In the Group Policy Object Editor, go to **New Group Policy Object** *your_policy* > **Computer Configuration** > **Windows Settings** > **Security Settings** > **Local Policies** > **User Rights Assignments**. In the content pane, select "Log on as a service" and double-click. Add the domain user for whom you are granting user rights and click **OK**. Repeat this step for "Act as part of the operating system."
8. The group policy is now enforced for the Organizational Unit to include the domain accounts and groups specified under the local Administrators group on each computer in the Organizational Unit. In addition, the domain user has been granted the necessary rights. To verify the user rights, log in to a domain computer and open the Computer Management console. Select **Groups**, double-click the Administrators group, and verify the membership of the domain users.

SQL access to Tivoli Storage Productivity Center views

IBM Tivoli Storage Productivity Center stores the information that is collected by data collection schedules in its DB2 database repository. This information is organized into a set of accessible views. You can use Structured Query Language (SQL) to retrieve the information from these views to create custom reports or for use within other applications.

Prerequisite: See the *Planning* section in the Tivoli Storage Productivity Center Information Center for information about the requirements when retrieving data from views in the database repository.

A view is a logical table that consists of data that is generated by a query. A view is based on an underlying set of base tables in the database repository. The data in a view is determined by a **SELECT** statement that is run on those base tables. The data collected Tivoli Storage Productivity Center is organized into the following sets of views that you can access in the database repository:

Storage entity views

These views include detailed information about the properties of storage entities. For example, the name, disk space, and available disk space for a storage subsystem. Information about the following storage entities is available in views in the database repository:

- Storage entity views: Hosts, volume groups, disks, virtual disks, external storage volumes, disk paths, virtual disk paths, controllers, logical volumes, file systems, clusters, cluster resource groups, hypervisors, virtual machines, virtual machine disks, hypervisor volumes, host virtual files, hypervisor partitions, MAC addresses, files, file types, directories, exports, OS users, OS user groups, IP addresses, Windows domains, and Storage Resource agents.
- Host relationship views:
  - Logical volume to host disk
  - Logical volume to logical volume (Parent and child logical volume relationships)
  - Computer to storage entity (disk or logical volume)
  - File system to logical volume

  (for example, the relationship between a logical volume and a host disk)

- Database views: Database instances, database containers, database archive logs, database control files, database data files, database devices, database log files, database log groups, database segments, database tables, database table spaces
- Storage subsystem views: Storage subsystems, storage virtualizers, storage pools, extents, volumes, disk drives, disk groups, element managers, nodes, backend disks, storage virtualizer I/O groups, volume paths, array sites, ranks
- Storage subsystem relationship views:
  - Disk drive to storage extent
  - Storage volume to storage extent
  - Volume to disk drive
  - Storage volume to backend disk
- NAS Filers
- Fabric views: Fabrics, nodes, peripheral entities, port to port connections, switches, switch blades, zones, zone members, zone sets, zone aliases, FC ports, host bus adapters
- Fabric relationship views:
  - Blade to FC port
  - Node to HBA, Node to port
  - Peripheral entity to node
  - Switch to port
  - Zone to alias, Zone to zone member, Zone to node, Zone to member
  - Zone set to zone
  - Zone alias to node
  - Alias to member
  - Fabric to switch, Fabric to zone set, Fabric to zone, Fabric to zone alias
- Performance Management views (for example, metrics and threshold violations)

Entities that are defined by Tivoli Storage Productivity Center
These views include information about entities that are defined by Tivoli Storage Productivity Center. These entities include Storage Resource agents, the alert log, the Tivoli Storage Productivity Center servers, scan profiles, monitoring groups, reporting groups, and storage resource groups.

Data sources
These views include information about the data sources that are associated with Tivoli Storage Productivity Center. These data sources include Storage Resource agents, Tivoli Storage Productivity Center, CIMOM agents, Out of Band agents, and VMWare Virtual Infrastructure.

Aggregated views
These views provide summary information for the database history, data in a database instance, and the Storage Resource agent file system.

Reporting views
These views combine several different entities into one view, such as aggregate database and file system summary views, file size, time, and age distribution views, enterprise-wide statistics views, directory size views, and historical views of statistics for database summaries, ping jobs, user space usage, and files in user groups.

Rollup views

These views include rollup report information from the master and subordinate Tivoli Storage Productivity Center servers, Storage Resource agents, host cluster data, computer group, host, database computer groups, fabric SAN assets, switch assets, storage subsystem group, storage subsystems, and Data Manager for Databases.

To view detailed descriptions of the views that are provided with Tivoli Storage Productivity Center, go to the Tivoli Storage Productivity Center support website.

This website contains the following documentation:

- Information about the TPCREPORT schema, including these items:
  - A list of the views that exist in the schema and descriptions of the columns with the views.
  - A list of the views that contain performance data and descriptions of the performance metrics within those views.
  - Diagrams that show the relationships between views.
- A diagram that traces the relationship between a computer and a storage subsystem disk. It also shows the views that contain information about the storage entities within that relationship.

Accessing views in the database repository

Use SQL to access the views in Tivoli Storage Productivity Center database repository. You can do this using an SQL command line or an application capable of using SQL to retrieve information from views. This topic provides an example of how to create and view a report based on those views.

To retrieve information from the database repository you must have the proper authority to access the views. You must also collect information about your storage assets to populate those views with information. See the *Planning* section in the Tivoli Storage Productivity Center Information Center for more information.

To perform the steps described in this topic, you must install the correct version BIRT. You can download BIRT v2.2.1 from the following Web page: http://www-01.ibm.com/software/brandcatalog/portal/opal/details?catalog.label=1TW10OT02. You must use this version of BIRT when working with Tivoli Storage Productivity Center and Tivoli Common Reporting.

You can use a number of SQL and reporting tools for accessing and viewing the content of IBM Tivoli Storage Productivity Center's exposed views. For example, you can use the DB2 Command Line Processor to run SQL commands against the views.

This topic describes how to use BIRT to create a report package based on the views and how to import that package into Tivoli Common Reporting for viewing. Tivoli Common Reporting is installed with Tivoli Storage Productivity Center.

See the following web sites for more information about BIRT and Tivoli Common Reporting:

- BIRT overview information: http://www.eclipse.org/birt/phoenix/intro/
- Tivoli Common Reporting overview information: https://www.ibm.com/developerworks/community/groups/service/html/communityview?communityUuid=9caf63c9-15a1-4a03-96b3-8fc700f3a364

Complete these steps to use BIRT and Tivoli Common Reporting to work with Tivoli Storage Productivity Center's views:

1. Install BIRT on a computer that has access to Tivoli Storage Productivity Center's database repository. The database repository is located on the computer where the Tivoli Storage Productivity Center server is installed.
2. Start the BIRT Designer tool.
3. Set the connection information for the Tivoli Storage Productivity Center database repository in Data Source section of the user interface.
4. Use the **Data Sets** section of the user interface to create a report based on SQL queries to the exposed views that are in the TPCREPORT schema. See http://www.eclipse.org/birt/phoenix/intro/ for information on how to use BIRT to create a report. When writing queries against the views make sure to qualify the view names with the TPCREPORT schema (for example, TPCREPORT.STORAGESUBSYSTEM).

   Note: BIRT limits the amount of views that are displayed to 100 by default. To change this, access the **Preferences** window for the Data Set Editor and enter a value greater than **275** in the **Maximum number of tables in each schema to display** field.
5. Define the layout of the report. Use BIRT to customize the appearance of a report, including customizing column names and creating charts, graphics, and aggregated totals for the data in the report. Use the preview function to view a snapshot of the report that you defined.
6. Create a report package (archive file) for the reports to import into Tivoli Common Reporting.
7. Import the report package into Tivoli Common Reporting.
   a. Start Tivoli Integrated Portal. For information about starting Tivoli Integrated Portal, go to the Tivoli Storage Productivity Center information center. Search for **Starting the Tivoli Storage Productivity Center GUI from Tivoli Integrated Portal**.
   b. Expand **Reporting** in the navigation tree and click **Common Reporting**.
   c. Right-click **Report Sets** and select **Import Report Package**.
   d. Select the report package you want to import and click **Import**. Click **Overwrite** in the **Advanced Options** section if the report package was previously imported. The name of the report package appears in the **Reports** section of the panel.
   e. Right-click the report package and select **Views As** > ***output_type***, where ***output_type*** represents the format of the generated report.
   f. View the report.

Repository copy tool

You can use the Repository copy tool, **repocopy**, to export all the tables in the IBM Tivoli Storage Productivity Center repository (Data Manager data) for purposes of debugging problems.

You can send the exported data to IBM Software Support to help debug problems.

Important:

- You must stop the Device server and Data Server before you use **repocopy** to import a repository.

- You cannot import **repocopy** data into the Tivoli Storage Productivity Center database if the migration tool was not run against the database. When you upgrade the Device server, you must run the migration tool before you import **repocopy** data.

Exporting repository data

Use the Repository copy tool to export data from an existing repository into a text file.

To export repository data, follow these steps:

1. Go to the following default directory:

   Windows operating systems:
   `c:\Program Files\IBM\TPC\data\server\tools`

   Linux or AIX operating systems:
   `/opt/IBM/TPC/data/server/tools`

2. Issue the **repocopy** command:

   Windows operating systems:
   repocopy.bat

   Linux or AIX operating systems:
   repocopy.sh

3. Select **Export data from repository tables** and click **Next**.
4. In the Options for Import/Export window, enter information in the following fields:

   Directory for Export
   Enter the directory where the comma-delimited file is saved.

   Delimiter
   Enter a delimiter for the delimited file format (a comma is the default).

   Quote Enter the symbol that contains string data (double quotation marks is the default).

   Tivoli Storage Productivity Center exports the data into the comma-delimited file that you specify, and places it in a file named *tablename*`.txt`. Click **Next**.
5. Select one of the following options and click **Next**.
   - Export by using DB2 native format.
   - Export by using text files (the preferred method).
6. Select one of the following options and click **Next**.
   - Export base tables (always export the base tables)
   - Export Performance Manager tables, if requested by IBM Software Support
   - Export history tables that are used for Tivoli Storage Productivity Center for Data history reports, only if requested by IBM Software Support

   The information that is detected in the `server.config` file is displayed in the Connection Properties window within the following fields:
   - Database Types
   - User name
   - Password
   - Driver Class
   - Driver URL
   - Database

- DB Creator
- Classpath

If you want to export data from a different database from the one listed in the `server.config` file, you can select the database from the **Database Types** list box. Manually enter the database information.

7. Click **Finish**.
8. Click **Run**.

As you progress through the export process, messages are written to a progress log that is displayed. You can track the steps through the progress log.

When the **repocopy** command is used with a remote database, the DB2 shared library is not available for loading the `libTSRMinsudb.so` file. You can ignore this message. Click **OK** and continue.

Importing repository data

Use the Repository Copy tool to import data from a text file. The exported data must be from an existing repository.

To import data into repository tables, follow these steps:

1. Stop the Device server and Data server.
2. Go to the following default directory:

   Windows operating systems:
   TPC_installation_directory\data\server\tools

   Linux or AIX operating systems:
   TPC_installation_directory/data/server/tools
3. Issue the **repocopy** command:

   Windows operating systems:
   repocopy.bat

   Linux or AIX operating systems:
   repocopy.sh
4. Select **Import data into repository tables**, and then click **Next**.
5. In the Options for Import/Export window, enter information in the following fields:

   Directory for Import
   Enter the directory where the comma-delimited files are stored.

   Delimiter
   Enter a delimiter that is used for the delimited file format (comma is the default). You must use the same delimiter that was used in the exported file.

   Quote Enter the symbol that is used to contain string data (double quotation marks is the default). You must use the same double quotation marks that were used in the exported file.

   Delete before inserting
   Check this option if you want to delete any existing data in the repository tables before you import new data.

   Tivoli Storage Productivity Center imports the data from the file that you specify. Click **Next**.

6. Enter the database and the access information of the database to which you want to import data. The Connection Properties window contains the following fields:
   - Database Types
   - User name
   - Password
   - Driver Class
   - Driver URL
   - Database
   - DB Creator
   - Classpath

   Click **Finish**.
7. Click **Run**.

As you progress through the import, messages are written to a progress log. You can use the progress log to track the steps as they happen.

When the **repocopy** command is used with a remote database, the DB2 shared library is not available for loading the `libTSRMinsudb.so` file. You can ignore this message. Click **OK** and continue.

Collecting diagnostic information about Tivoli Storage Productivity Center

You can use the service tool to collect diagnostic information about Tivoli Storage Productivity Center. The tool detects the system configuration, collects the applicable information, and creates a compressed file that can be sent to IBM Software Support.

Related tasks:

"Creating a compressed file for Storage Resource agent" on page 200
Run the service tool on Storage Resource agents that were deployed by using the web-based GUI to create a compressed file that can be sent to IBM Software Support.

"Creating a compressed file for servers" on page 199
Run the service tool for the Data Server and the Device server to create a compressed file that can be sent to IBM Software Support.

Related reference:

"Service tool overview" on page 198
The service tool collects information from all installed Tivoli Storage Productivity Center components. The tool detects the system configuration, collects the applicable information, and creates a compressed file that can be sent to IBM Software Support.

"How to customize the service tool" on page 201
You can customize the service tool to collect data about specific Tivoli Storage Productivity Center components, which you can send to IBM Software Support. You can also use command-line parameters to specify a location to place the data that is collected, specify that the data is compressed, or to specify both.

Service tool overview

The service tool collects information from all installed Tivoli Storage Productivity Center components. The tool detects the system configuration, collects the applicable information, and creates a compressed file that can be sent to IBM Software Support.

The service tool collects the following information:

- Host name
- IP address and configuration information
- Operating system and version. On the Windows operating system, a `msInfo.txt` report is also generated
- Java home, version, and class path
- Java Virtual Machine (JVM) implementation name and version
- Protocol statistics
- Internet Protocol network connections for Tivoli Storage Productivity Center, including listening ports
- Diagnostic information about the system and its services
- Listing of all library files, for example, server and library and agent and library
- HOSTS file
- Tivoli Storage Productivity Center version and license files

When the service tool is run on the system where the Data Server or the Device server are installed, it also collects the following information:

- For the Data Server, information about all of the remote and local graphical user interfaces (GUIs) that are associated with it
- For Device and Replication servers, information about their profiles in IBM WebSphere Application Server Liberty Profile
- All applied interim fixes
- Installation logs
- The contents of the `log` and `logs` directory, including subdirectories
- The contents of the `conf` and `config` directory
- Directory listing of the `lib` and `bin` directory
- The contents of the `log` and `conf` subdirectories of the `web` directory
- For the web-based GUI, information about its profile in the embedded WebSphere Application Server
- Information from the **ipconfig /all** command on Windows operating systems
- Information from the **ipconfig -a** command on Linux and AIX operating systems
- Information from the **netstat -an** command on all operating systems

When the service tool is run on the system where the database repository is installed, it also collects the DB2 support information.

When the service tool is run on the system where Tivoli Common Reporting component is installed, it also collects the following information:

- Logs for the deployment engine
- A listing of the Composite Offering Installer (COI) components that are installed
- InstallAnywhere logs for Tivoli Common Reporting
- Logs, configuration, and profile files for Jazz for Service Management

- Logs and configuration files for Tivoli Common Reporting that are located in the `reporting` directory in the Jazz for Service Management installation directory
- Information from the **ipconfig /all** command on Windows operating systems
- Information from the **ifconfig -a** command on Linux and AIX operating systems
- Information from the **netstat -an** command on all operating systems

The service tool does not collect service logs or other service information for Tivoli Storage Productivity Center for Replication. To collect the log and trace files for the Replication server, go to the Tivoli Storage Productivity Center information center and select **IBM Tivoli Storage Productivity Center** > **Troubleshooting** > **Tivoli Storage Productivity Center for Replication** > **Logs**.

When the service tool runs on the Storage Resource agent computer, it collects the following information:

- All applied interim fixes
- Everything in the `config`, `log`, `nls`, `output`, and `service` directories, including subdirectories
- Everything in the `opt/IBM/CAP` directory on Linux and AIX operating systems
- Directory listing of the `ProgramData\Application Data\IBM\CAP` directory on Windows operating systems
- Directory listing of the `agent` directory
- Directory listing of the `bin` directory
- Listing of version numbers for the Storage Resource agent component
- Information from the **ipconfig /all** command on Windows operating systems
- Information from the **ifconfig -a** command on Linux and AIX operating systems
- Information from the **netstat -an** command on all operating systems

By default, the service data is collected in one of the following directories:

Windows operating systems:

TPC_installation_directory`\service\data`

Linux and AIX operating systems:

TPC_installation_directory`/service/data`

For more information about changing the default directory, see "How to customize the service tool" on page 201.

You can run the service tool on Tivoli Storage Productivity Center regardless of whether you configured it on a single server or on multiple servers. The service tool automatically recognizes the installed components and collects service data about them. For more information about running the service tool for servers, see "Creating a compressed file for servers."

Creating a compressed file for servers

Run the service tool for the Data Server and the Device server to create a compressed file that can be sent to IBM Software Support.

You must have administrator authority on Windows operating systems or root authority on AIX and Linux operating systems.

To run the tool for the servers, follow these steps:

1. Log on to the system where Tivoli Storage Productivity Center is installed.
2. If you used the default directory, go to the following directory:

   Windows operating systems:
   `C:\Program Files\IBM\TPC\service\`

   Linux or AIX operating systems:
   `/user_or_opt/IBM/TPC/service/`
3. Run the following program:

   Windows operating systems:
   `service.bat`

   Linux or AIX operating systems:
   `service.sh`

A compressed file, `TPCServiceFiles.zip`, is created in the following directory:

Windows operating systems:
`TPC_installation_directory\service\data`

Linux or AIX operating systems:
`TPC_installation_directory/service/data`

Tip: You can customize the data that is collected by the service tool. For more information, see "How to customize the service tool" on page 201.

Related tasks:

"Creating a compressed file for Storage Resource agent"
Run the service tool on Storage Resource agents that were deployed by using the web-based GUI to create a compressed file that can be sent to IBM Software Support.

Creating a compressed file for Storage Resource agent

Run the service tool on Storage Resource agents that were deployed by using the web-based GUI to create a compressed file that can be sent to IBM Software Support.

You must have administrator authority on Windows operating systems or root authority on AIX and Linux operating systems.

To run the service tool on the Storage Resource agents, complete the following steps:

1. Log on to the web-based GUI.
2. In the navigation pane, select **Server Resources** > **Servers**.
3. Right-click the server where the Storage Resource agent is deployed, and select **Logs** > **Collect Agent Logs**.

The following compressed files are created:

Windows operating systems:
`C:\Program Files\IBM\TPC\data\log\SRATraces\SRA_computer_name\TPCServiceInfo.zip`

AIX and Linux operating systems:
`/opt/IBM/TPC/data/log/SRATraces/SRA_computer_name/TPCServiceInfo.zip`

where *SRA_computer_name* represents the name of the computer on which the Storage Resource agent is located.

If the compressed file cannot be created, a message indicates that the job was unsuccessful.

For more information about the error, see the server log file or the services script information file. The files are in one of the following default directories:

Server log file
This file is on the system where Tivoli Storage Productivity Center is installed:

Windows operating systems:
`c:\Program Files\IBM\TPC\data\log`

AIX or Linux operating systems:
`/opt/IBM/TPC/data/log`

Services script information file
This file is on the computer on which the Storage Resource agent is installed:

Windows operating systems:
`c:\Program Files\IBM\TPC\`*`SRA_computer_name`*`\services\TPCServiceInfo.log`

AIX or Linux operating systems:
`/opt/IBM/TPC/`*`SRA_computer_name`*`/services/TPCServiceInfo.log`

For more information about customizing the data that is collected by the service tool, see "How to customize the service tool."

Related tasks:

"Creating a compressed file for servers" on page 199
Run the service tool for the Data Server and the Device server to create a compressed file that can be sent to IBM Software Support.

How to customize the service tool

You can customize the service tool to collect data about specific Tivoli Storage Productivity Center components, which you can send to IBM Software Support. You can also use command-line parameters to specify a location to place the data that is collected, specify that the data is compressed, or to specify both.

Specifying help and output command-line parameters

To obtain information about the service tool usage, use the **-help** command-line parameter.

To specify the data that is collected by the service tool, use the following command-line parameters when you run service tool:

-output *directory_path*
Places the files that contain the data that was collected in a directory that you specify. If you specify a directory that does not exist on your system, that directory is created. If you do not use the **-output** *directory_path* parameter, the files are placed in the default directory:

Windows operating systems
`TPC_installation_directory``\service\data`

Linux or AIX operating systems
TPC_installation_directory/service/data

Restriction: If you specify a directory, the directory path cannot contain spaces. This restriction refers to the **-output** option.

-nozip
When you use this parameter, the compressed collected data archives are no longer created. The service tool creates separate directories for each component for which data was collected. You can then create compressed archives for the collected files. In this way, you can control the size and content of each compressed file. To specify a directory other than the default directory, use the **-output** *directory_path* parameter.

Tip: You can specify more than one parameter, for example, `C:\Program Files\IBM\TPC\service>service -install -nozip`.

Collecting data for specific Tivoli Storage Productivity Center components

You can use the service tool to collect data about specific Tivoli Storage Productivity Center components.

Use the following parameters to specify the components:

-all
All components.

-install
Installation component files.

-data
Data Server component files.

-device
Device server component files.

-repl
Replication server component files.

-sra
Storage Resource agent component files.

-db
DB2 files.

-cli
Command-line interface files.

-gui
Collects data about web-based GUI files.

-tcr
Jazz for Service Management and Tivoli Common Reporting data.

Tip: You can specify to collect data for more than one component, for example, `C:\Program Files\IBM\TPC\service>service -tcr -install -gui`.

If you collect information about a particular component, and you do not specify the **-nozip** parameter, you can identify the contents of a compressed service file from its name. For example, if you specify the **-tcr -install -gui** parameters but

did not specify the **-nozip** parameter, a file named TPCServiceFiles_tcr_install_gui.zip is created.

Note: An existing compressed file is overwritten when another file of the same name is created. For example, if you run C:\Program Files\IBM\TPC\service>service **-db -gui**, a file named TPCServiceFiles_db_gui.zip is created. If you rerun the tool with the same component options, a new file named TPCServiceFiles_device_db.zip is created. This new file overwrites the previously created file unless you specify the **-nozip** parameter, or use the **-output** parameter to specify a different path.

If you specify the **-nozip** parameter, a directory is created for the components that you specified in the parameter. If you did not specify a specific component, data is collected for all installed components, and the data is placed in files in the following directory:

Windows operating systems:
TPC_installation_directory\service\data

Linux or AIX operating systems:
TPC_installation_directory/service/data

Files for particular components are then placed in a directory that corresponds to that component. Common files, such as license.txt, are placed in the following directory:

Windows operating systems:
TPC_installation_directory\service\data

Linux or AIX operating systems:
TPC_installation_directory/service/data

Restriction: You cannot specify a specific component as a command-line parameter when you specify the **-all** parameter. Also, when you specify a component that is not installed on the computer, the service tool displays an error message.

Administering resources and data sources

Administer monitored resources and the data sources that are associated with those resources. Data sources can be agents that manage resources or VMware vCenter servers. An agent might be a CIM agent or a Storage Resource agent.

In Tivoli Storage Productivity Center V5.2, many of the tasks for administering resources were moved to the web-based GUI. Use the following table as a roadmap to the functions for administering resources that are available in each GUI:

Table 44. Tasks for administering resources in each GUI

Tasks	Web-based GUI	Stand-alone GUI
Adding resources for monitoring Learn more...	You can add the following resources: • Storage systems (other than IBM SONAS storage systems) • Servers (without a Storage Resource agent • Hypervisors • Fabrics • Switches	You can add the following resources: • Servers (with a Storage Resource agent) • IBM SONAS storage systems • Tivoli Storage Productivity Center servers as subordinate servers
Removing resources (data sources)	You can remove the following resources: • Storage systems • Servers • Hypervisors • Fabrics • Switches • NAS filers	You can remove the following resources: • CIM agents • VMware data sources • Tivoli Storage Productivity Center servers that are subordinate servers
Testing connections	You can test the connections of all monitored resources to ensure that Tivoli Storage Productivity Center is communicating with those resources. **Restriction:** You cannot use this GUI to test the connection of Tivoli Storage Productivity Center servers that are subordinate servers.	You can test the connection of the following resources: • Storage systems • Tivoli Storage Productivity Center servers that are subordinate servers
Updating user credentials	Update the credentials for a resource when you want to change the user name and password that Tivoli Storage Productivity Center uses to communicate with that resource.	Not available
Administering Storage Resource agents that are deployed on servers Learn more ...	You can take the following actions: • View information about an agent • Update user credentials and certificate location • Test connection • Enable and disable • Enable and disable fabric functions • Enable and disable the running of scripts • Upgrade an agent • View a log of actions that were taken by an agent • Collect service logs • Modify ports	You can take the following actions: • View information about an agent • Deploy an agent • View a log of actions that were taken by an agent • Register an agent with a different server

Table 44. Tasks for administering resources in each GUI (continued)

Tasks	Web-based GUI	Stand-alone GUI
Showing the resources that are managed by a data source or agent	Not available	You can view a list of the resources that are managed by the following data sources: • CIM agents • VMware vCenter servers

Checking the status data sources

Check the status of data sources in the stand-alone GUI.

1. In the navigation tree of the stand-alone GUI, expand **Administrative Services** > **Data Sources**.
2. Click any of the following nodes to view the status of the related data sources:
   - **CIMOM Agents**
   - **Storage Resource agents**
   - **Out of Band Fabric Agents**
   - **Storage Subsystems**
   - **TPC Servers**
   - **VMware VI Data Source**
3. For CIMOM Agents and Storage Subsystems, view the icon in the Connection Status column. For other data sources, view the icon in the State column. If the data sources are up and running, a green health status icon is displayed.

For information about how to view the status of resources in the web-based GUI, see "Viewing the status of top-level resources" on page 379.

CIM agents

Administer CIM agents that are associated with storage resources that are monitored by Tivoli Storage Productivity Center. Common Information Model (CIM) agents provide a CIM interface for collecting information about certain types of storage systems and switches.

Tivoli Storage Productivity Center communicates with CIM agents to collect information about the following resources:

- IBM TotalStorage Enterprise Storage Server
- IBM System Storage DS4000
- IBM System Storage DS5000 Storage Manager
- IBM System Storage DS6000
- Non-IBM storage systems that are managed by SMI-S certified Common Information Model Object Manager (CIMOM), such as EMC, Hitachi, and NetApp
- Switches: Brocade, Cisco, QLogic

Tivoli Storage Productivity Center communicates directly with the following storage systems and does not require CIM agents:

- System Storage DS8000
- SAN Volume Controller
- The XIV system

- Storwize V7000
- Storwize V3500
- Storwize V3700
- Storwize V7000 Unified
- IBM SONAS

Adding a CIM agent

Use the web-based GUI to add storage systems and switches that are managed by CIM agents. When you enter information about a storage system or switch, you must enter connection information for the CIM agent that manages it.

You must specify connection information for CIM agents when you add the following resources:

- IBM TotalStorage Enterprise Storage Server
- IBM System Storage DS4000
- IBM System Storage DS5000 Storage Manager
- IBM System Storage DS6000
- Non-IBM storage systems that are managed by SMI-S certified Common Information Model Object Manager (CIMOM), such as EMC, Hitachi, and NetApp
- Switches: Brocade, Cisco, QLogic

Tips:

- In 5.2 and later, the ability to add CIM agents as data sources was moved from the stand-alone GUI to the web-based GUI. You define a connection to a CIM agent when you add resources that are managed by that agent.
- If a CIM agent manages multiple resources, all the resources that it manages are added to Tivoli Storage Productivity Center.
- CIM agents must comply with SMI-S standards.
- For a complete list of resources and their CIM agents that you can add, see the *Switches & Directors* and *Storage* sections in the *Tivoli Storage Productivity Center Interoperability Matrix*.

To add storage system or switch that is managed by a CIM agent, complete the following steps:

1. In the navigation pane of the web-based GUI, select the resource that you want to add. To add a storage system, select **Storage Resources** > **Storage Systems**. To add a switch, select **Network Resources** > **Switches**.
2. Click **Add *resource***, where *resource* represents the type of resource that you want to add. To add a storage system, click **Add Storage System** and select the type of storage system that you want to add. To add a switch, click **Add Switch**.
3. Specify information about the CIM agent that manages the resource.

   Help tips in the GUI: To view descriptions of the information that you must enter for a CIM agent, hover the mouse pointer over the related help icons in the wizard.
4. Complete the wizard.

This procedure does not physically add a CIM agent, but enables Tivoli Storage Productivity Center to communicate with that agent and collect data about its managed resources.

Viewing information about a CIM agent

View detailed information about CIM agents that manage the resources that are monitored by Tivoli Storage Productivity Center.

Use the following locations in the Tivoli Storage Productivity Center GUIs to view information about a CIM agent:

Web-based GUI

To view information about a resource that is managed by a CIM agent, complete the following steps:

1. In the navigation pane of the web-based GUI, select the type of resource that you want to view. To view a storage system that is managed by a CIM agent, select **Storage Resources** > **Storage Systems**. To view a switch that is managed by a CIM agent, select **Network Resources** > **Switches**.
2. On the resource list page, right-click a storage system or switch that is managed by a CIM agent and select **View Properties**. The key properties of the resource and its CIM agent is displayed.

Stand-alone GUI

To view information about a CIM agent, use either of the following methods:

- In the navigation tree of the stand-alone GUI, expand **IBM Tivoli Storage Productivity Center** > **Reporting** > **Data Source Reports** > **CIMOM Agents** and click **By CIMOM Agent** or **By Managed Device**. For more information about these reports, see "Data source reports - CIMOM agents" on page 1633.
- In the navigation tree, expand **Administrative Services** > **Data Sources** and click **CIMOM Agents**. A list of CIM agents is displayed. Click the magnifying glass icon next to a CIM agent to view the following information about that agent:

  Service URL
  The service URL of a CIM agent that contains its IP address, the port on which a CIM agent is listening, and the protocol that is used for communication. This URL has a protocol [http|https], an IP address or host name, and a port number.

  Display Name
  The name of the CIM agent as specified by the CIM provider.

  Description
  The optional description that was entered when the CIM agent was added.

  Username
  The user ID that is used for authentication when Tivoli Storage Productivity Center logs in to a CIM agent.

  Interoperability Namespace
  Enter the interoperability namespace of a CIM agent. This namespace within the CIMOM allows for accessing a CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIM agent when it retrieves information.

Truststore Location
The location (path on this computer) of a certificate file for certificate-based authentication in the https protocol.

User Interface Description
The name of the Human Interface Service (if any) supported by this CIM agent.

Software Level
The software version level of the CIM agent.

Protocol Version
The version of the cim-xml protocol.

Authentication Mechanism
The authentication mechanism that is supported by a CIM agent. This field can contain the following values: Unknown, None, Other, Basic, Digest.

Protocol
The communication protocol that is used for the CIM agent. Possible values are http and https.

SLP Attributes
The standard set of attributes for this CIM agent. The attributes are retrieved via SLP.

Connection Status
The status of a CIM agent as it relates to Tivoli Storage Productivity Center. Possible values are SUCCESS, UNCONFIGURED, UNKNOWN, INVALID_NAMESPACE, TIMEOUT, REFUSED, LOGIN_FAILED, SSL_HANDSHAKE_ERROR, SSL_REGISTRATION_INVALID, CIMCLIENT_ERROR.

Status Timestamp
The date and time when the Connection Status information was last collected.

Updating the credentials for a CIM agent

Change the credentials that Tivoli Storage Productivity Center uses to authenticate to a CIM agent.

To update the credentials for a CIM agent, complete the following steps:

1. In the navigation pane of the web-based GUI, select the resource that is managed by a CIM agent. To update the credentials of a CIM agent that manages a storage system, select **Storage Resources** > **Storage Systems**. To update the credentials of a CIM agent that manages a switch, select **Network Resources** > **Switches**.
2. Right-click the storage system or switch that is managed by a CIM agent and select **Connections** > **Update Credentials**. If a resource is managed by multiple CIM agents, the menu is displayed as **Connections** > **Update Credentials** > ***data sources***. Select the data source for which you want to update the credentials.
3. Update the user name and password for a CIM agent. The following information is shown for a CIM agent:

   CIM agent host name or IP address
   The host name or IP address for the computer on which the CIM agent that manages the resouce is installed. Depending on what is supported

in your environment, you can enter an Internet Protocol version 4 (IPv4) or IPv6 address. If you enter an IPv6 address, the preferred representation is written as eight groups of four hexadecimal digits. Example: 2001:DB95:0000:1234:0000:0000:5678:ABCD.

User name, password
The user name and password that are used to authenticate to the CIM agent.

4. Click **OK** to apply the changes.

Testing the connection to a CIM agent

Ensure that Tivoli Storage Productivity Center is communicating properly with a CIM agent.

To test the connection between Tivoli Storage Productivity Center and a CIM agent, complete the following steps:

1. In the navigation pane of the web-based GUI, select the resource that is managed by a CIM agent. To test the connection of a CIM agent that manages a storage system, select **Storage Resources** > **Storage Systems**. To test the connection of a CIM agent that manages a switch, select **Network Resources** > **Switches**.
2. Right-click the storage system or switch that is managed by a CIM agent and select **Connections** > **Test Connection**. If a resource is managed by multiple CIM agents, the menu is displayed as **Connections** > **Test Connection** > ***data sources***. Select the data source for which you want to test the connection. This process might take some time to complete. In this case, you can choose to continue the operation in the background. The status of the resource is automatically updated when the process is complete.

Removing a CIM agent

Remove a CIM agent from Tivoli Storage Productivity Center. Data that was discovered by a CIM agent is not removed from the database repository.

To remove a CIM agent, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Administrative Services** > **Data Sources**.
2. Click **CIMOM Agents** to view a list of agents that are monitored by Tivoli Storage Productivity Center.
3. Select a CIM agent.
4. Click **Remove CIMOM**.
5. Click **OK** to confirm the removal of the CIM agent.

Viewing the resources that are managed by a CIM agent

View the resources that are managed by a CIM agent. A CIM agent can manage multiple storage systems or switches at the same time.

To view the resources that are managed by a CIM agent, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Administrative Services** > **Data Sources**.
2. Click **CIMOM Agents** to view a list of agents that are monitored by Tivoli Storage Productivity Center.
3. Select a CIM agent.
4. Click **Show Managed Devices**. A list of managed resources is displayed.

Collecting CIM agent logs

You can collect logs for certain IBM CIM agents using the command line interface.

Remember: Storage systems that use the native interfaces (for example, DS8000, the XIV system, SAN Volume Controller, and Storwize V7000) do not use CIM agents.

1. Change to the directory where the CIM agent is installed.
   - On Linux operating systems, DS3000, DS4000, DS5000, DS6000 are installed at `/opt/IBM/cimagent/cimom`
   - On Windows operating systems, DS3000, DS4000, DS5000, DS6000 are installed at `C:\Program Files\IBM\cimagent\cimom`
2. Run one of the following commands:

   On Linux operating systems

   ```
   collectLogs.sh
   ```

   On Windows operating systems

   ```
   collectLogs.bat
   ```

   A **collectedLogs.zip** file is created.

   Important: This file is overwritten if you run the script again.

Verifying that a CIM agent is running

You can verify that a CIM agent is running from the command line interface.

To verify that a CIM agent is up and running, run the following command:

```
telnet <IP> <port>
```

Where <IP> is the IP address of the system where the CIM agent is installed, and <port> is the port number. By default, this is 5989 for a secure connection and 5988 for an unsecure connection.

Servers and Storage Resource agents

Administer servers and the Storage Resource agents that collect asset, fabric, file, and file system information about servers.

Storage Resource agents can also gather information about database managers installed on the server, and NAS device information. You can create ping, probe, and scan jobs to run against the servers that have Storage Resource agents installed.

Use the stand-alone GUI to deploy Storage Resource agents to servers. Use the web-based GUI to complete all other tasks for Storage Resource agents.

For information about deploying Storage Resource agents, see “Deployment guidelines and limitations for Storage Resource agents” on page 83.

Adding servers with Storage Resource agents

Use the stand-alone GUI to add Storage Resource agents to your system. You can add Storage Resource agents by scheduling a job or by adding an agent immediately.

To add a Storage Resource agent, complete the following steps:

1. In the navigation tree, click **Administrative Services** > **Configuration**. Right-click **Storage Resource Agent Deployments**. On the menu, click **Create Storage Resource Agent Deployments**.
2. The Create Storage Resource agent Deployments window opens displaying the Computers page. Enter the following information:

   Header fields

   Creator
   Predefined creator name.

   Name Predefined name.

   Description (optional)
   Enter a user-defined description name for the job.

   Enabled
   Check this box to enable Storage Resource agent deployments.

   Add Host List
   Click this button to enter names and logon credentials for the computers on which you want to deploy Storage Resource agents.

   Remember: When you use an IPv6 address to add a server, the TPC Server must also have an IPv6 address. Both of these servers must be in the same domain name system (DNS).

   You can enter information in one of the following ways:

   - Click **Add Agents from MS Directory** to install Storage Resource agents on one or more Windows computers that are members of a Windows domain. Enter the domain controller name and login information about the displayed window. On the next window, select the domain and click **Get List of Domain Computers**. The host names of computers that you add by using this method are displayed in the **Remote Agent Machines** section of the window. In that section, you can also specify an installation location for each computer and indicate whether you are reinstalling a Storage Resource agent on a computer.
   - Click this button to retrieve a list of computer names from a flat text file. The computers listed in the file are automatically added to the deployment schedule and displayed in the **Remote Agent Machines** section. In that section, you can also specify an installation location for each computer and indicate whether you are reinstalling a Storage Resource agent on a computer.

     The file that you retrieve must contain one host name or IP address per line. For example:

     ```
     tpcserver1.tpc.example.com
     9.79.179.179
     tpcserver1.tpc.example.com
     9.89.189.189
     ```

     If Tivoli Storage Productivity Center discovers any syntax problems in the file, none of the host entries in the file are added to the deployment job.
   - In the **Remote Agent Machines** section, enter the names or IP addresses of the computers where you want to deploy the Storage Resource agents. Specify the installation location for the Storage

Resource agent runtime files. If you do not enter an installation location, the default location is used. The default installation directories are:

- For Windows: `C:\Program Files\IBM\TPC\agent`
- For Linux and UNIX: `/opt/IBM/TPC/agent`

- Select **Force** under the following circumstances:
  - If an earlier Storage Resource agent installation failed and there are damaged agent files on the computer that cause further installations to fail. If you select this option, Tivoli Storage Productivity Center attempts to overwrite the previous failed deployment on the computer with a new Storage Resource agent.
  - If you want an existing Storage Resource agent to communicate with an additional Tivoli Storage Productivity Center server. To do this, you must create the deployment job from the additional Tivoli Storage Productivity Center server to which you want the Storage Resource agent to communicate.

    To install the Storage Resource agent to communicate with two different servers using the same location, follow the example in these steps:

    a. From server 1 (S1), install the Storage Resource agent to the location `C:\Program Files\IBM\TPC\agent` using port 9510.
    b. From server 2 (S2), install the Storage Resource agent in the same location using the same port 9510 and specify the **Force** option. This option is required when you install the Storage Resource agent in the same location.

    If you want to install two Storage Resource agents on the same system, each with a different port, install the agents using different port numbers and different locations.

  Note:

  - You cannot change the communication type for a Storage Resource agent (daemon or non-daemon) when you select the **Force** option. Make sure to select the same communication type as the existing Storage Resource agent when you create a deployment job.
  - If a Storage Resource agent exists on a target computer and you do not select **Force**, an error occurs during validation and the Storage Resource agent is not installed.
- Select **Initial Scan** to perform a file system scan on the computer where a Storage Resource agent is deployed. If this box is checked, the scan is run immediately after a Storage Resource agent is deployed.
- Click **Windows Services Information** to enter login information that Tivoli Storage Productivity Center uses to create a Windows service for the Storage Resource agent. Use the Windows Service Information window to select the local system account, enter information about an existing user account, or create a new local account. If you do not specify a Windows service account, the default itcauser account is created. This action is only available for agents that are deployed to Windows computers as daemon services. This button is disabled if Use Daemon Service for Runtime Operation is not selected.

Enter the following information:

User Enter the user ID that Tivoli Storage Productivity Center uses

when logging in to the host computer to install a Storage Resource agent. The value in the field is applied to all the computers that appear in the **Remote Agent Machines** table.

Attention: This ID must have administrative or root privileges on the target computer.

Password
Enter the password for the user ID.

Certificate Location
Enter the fully qualified path of the certificate on the computer where the Data server is located. This file is used for certificate-based authentication. If you do not enter a value in this field the default location is used.

Passphrase
Enter the passphrase for the certificate file. The passphrase was created when the certificate was generated.

Port Enter the port number on which a Storage Resource agent listens for requests. The default is 9510.

Tip: This value is required if you run the Storage Resource agent as a daemon service.

Use Daemon Service for Runtime Operation
Select this option to run a Storage Resource agent on the monitored computer as a daemon service.

See Storage Resource agent Deployment Considerations for more information about deploying agents as daemon and non-daemon services.

Restriction: Do not select this option to run the Storage Resource agent as a stand-alone executable file on the monitored computer.

Validate before Save
Check this box to indicate that Tivoli Storage Productivity Center should attempt to communicate with the computers that are added to the deployment job when you click **Save** on this page. If the validation fails for a computer, the agent is not installed on that computer.

Edit Selected Entries
Click this button to edit settings for the selected Storage Resource agents. This information includes user ID, password, certificate location, passphrase, and port. You must submit the deployment job for the changes to take effect.

Remove
Click this button to remove a Storage Resource agent from the deployment job.

Tip: This button does not remove the agent from Tivoli Storage Productivity Center if it is already installed; it only removes it from the deployment job.

3. Click the **When to run** tab to enter the following information:

How often to run

Specify a time to run:

- **Run now**
- **Run once at** (specify a date and time to run)

How to handle time zones

Specify a time zone to use:

- **Use the time zone that the server runs in**
- **Use this time zone** (select a time zone)

4. Click the **Alert** tab to specify the following information:

   Triggering-Condition

   The triggering condition you can specify is: Storage Resource agent Deployment Failed

   Triggered-Actions

   Choose from the following check boxes:

   - **SNMP Trap**
   - **TEC / OMNIbus Event**
   - **Login Notification**
   - **Windows Event Log**
   - **Email**

   Depending on what action you select, you might have other choices to make. For example, if you select the **Windows Event Log** check box, the **Event Type** field becomes active so that you can specify a severity for the event in the Windows event log.

5. Click **File > Save**.

For related tasks, see the *Planning > General Planning > Planning for Storage Resource agents* section in the *Tivoli Storage Productivity Center Installation and Configuration Guide* for additional configuration information about certificates used after deployment.

Check the TPCD_#####.log file in the Tivoli Storage Productivity Center installation directory to view the error messages related to failed validations. The following are the default locations for this log file:

- Windows: `C:\Program Files\IBM\TPC\data\log`
- UNIX/Linux: `/opt/IBM/TPC/data/log`

If multiple Storage Resource agents are set up to probe or scan the same storage resources, the Storage Resource agent that was added to Tivoli Storage Productivity Center first is used for the probe or scan. Therefore, only data that is gathered by the first Storage Resource agent is shown.

Viewing information about Storage Resource agents

View detailed information about the Storage Resource agents that are deployed on monitored resources.

To view information about a Storage Resource agent, complete the following steps:

1. From the navigation pane in the web-based GUI, click **Server Resources** > **Servers**.
2. On the Servers page, right-click the server where the agent is deployed and select **View Properties**.

3. In the properties notebook, click the **Agent** tab.

Detailed information about the agent is shown, such as the agent state and version, and the date and time when the agent was last updated.

If the Storage Resource agent has a state of Upgrade needed, the agent must be upgraded to the same version level as the Tivoli Storage Productivity Center server to which it is communicating.

Viewing Storage Resource agent log files

The log files for a Storage Resource agent contain informational, warning, and error messages for the actions that were taken by the agent. You can use the content of the log files to troubleshoot any errors that might occur when a Storage Resource agent is started, processing data, or shut down.

By default, the log files are located in the following directories on the server where an agent is deployed:

Windows

```
C:\Program Files\IBM\TPC\agent\log\SRV1\agent.log
```

Linux, UNIX, and AIX

```
/opt/IBM/TPC/agent/log/computer_name/agent.log
```

where *computer_name* represents the name of the server where Tivoli Storage Productivity Center is installed. If an agent communicates with more than one installation of Tivoli Storage Productivity Center, a subfolder is created for each installation. For example, if the agent communicates with servers named SRV1 and SRV2, the following folders are created:

- `C:\Program Files\IBM\TPC\agent\log\SRV1\agent.log`
- `C:\Program Files\IBM\TPC\agent\log\SRV2\agent.log`

To view the log file for a Storage Resource agent, complete the following steps:

1. From the navigation pane in the web-based GUI, click **Server Resources** > **Servers**.
2. On the Servers page, locate the server that contains the Storage Resource agent that you want to analyze.
3. Right-click the server row and select **Logs** > **View Agent Log**.
4. Optional: To view only the log entries that have a Warning or Error status, select an option from the **Show all** list. You can choose to view only entries that have the following statuses:
   - Only error entries
   - Only warning entries
   - Error and warning entries
5. Optional: To view an explanation of the message that is associated with a log entry, click the link in the **ID** column.

Upgrading Storage Resource agents

Upgrade Storage Resource agents to ensure that they are at the same release level as the Tivoli Storage Productivity Center server.

When you apply maintenance to Tivoli Storage Productivity Center, you can upgrade Storage Resource agents immediately or at a later time. To ensure that all your agents are at the current release level and to manage your network load, schedule upgrades regularly.

If a Storage Resource agent is not at the same level as the Tivoli Storage Productivity Center server, the following limitations occur:

- New functions in the current release might not be available for the resources that are monitored by the agent.
- Problem fixes are not applied to the agent.

You can upgrade Storage Resource agents by using the following methods:

- Use the **Modify Agents** > **Upgrade** action on the Servers page in the web-based GUI.
- Use a Storage Resource agent command.

To determine if a Storage Resource agent must be upgraded, complete the following steps:

1. In the navigation pane of the web-based GUI, go to **Server Resources** > **Servers**.
2. View the values in the Agent State column. If the state of the agent is **Upgraded needed**, the Storage Resource agent for the related server must be upgraded.

Starting agent upgrades:

Upgrade a Storage Resource agent to the same release level as the Tivoli Storage Productivity Center server.

The ability to start the upgrade process for a Storage Resource agent is available when the following conditions are met:

- A Storage Resource agent must be deployed on the server that you want Tivoli Storage Productivity Center to monitor.
- An agent upgrade is not currently running for the server.
- The version of the agent that is deployed on the server is earlier than the Tivoli Storage Productivity Center server version.

To upgrade a Storage Resource agent that was not upgraded at maintenance time, complete the following steps:

1. In the navigation pane of the web-based GUI, go to **Server Resources** > **Servers**.
2. On the Servers page, right-click the server that contains the Storage Resource agent to upgrade and select **Modify Agents** > **Upgrade**.
3. Select **Immediate** from the **Agent Upgrade** list on the Upgrade Agent window.
4. Click **Upgrade** to start the upgrade process.

Scheduling agent upgrades:

Schedule the upgrade process for a Storage Resource agent.

You can schedule the upgrade process for a Storage Resource agent when the following conditions are met:

- A Storage Resource agent must be deployed on the server that you want Tivoli Storage Productivity Center to monitor.
- An agent upgrade is not currently running for the server.
- The version of the agent that is deployed on the server is earlier than the Tivoli Storage Productivity Center server version.

Tips:

- To manage the workload for a server and the network, schedule the agent upgrade for a time when the server and network are not busy.
- The scheduled time for an agent upgrade is based on the time zone of the Tivoli Storage Productivity Center server, not the time zone of the server where the Storage Resource agent is installed. For example, if an agent is installed on a server in the Central (CST) time zone, but the Tivoli Storage Productivity Center server is in the Pacific (PST) time zone, the time that is shown in the web-based GUI when you schedule the upgrade is PST.

To schedule the upgrade of a Storage Resource agent that was not upgraded at maintenance time, complete the following steps:

1. In the navigation pane of the web-based GUI, go to **Server Resources** > **Servers**.
2. On the Servers page, right-click the server that contains the Storage Resource agent to upgrade and select **Modify Agents** > **Upgrade**.
3. Select **Specific** from the **Agent Upgrade** list on the Upgrade Agent window.
4. Select the date and time and click **Upgrade** to schedule the agent upgrade.

Upgrading Storage Resource agents by using a command:

You can manually upgrade Storage Resource agents.

To manually upgrade the Storage Resource agent, complete the following steps:

1. Go to the DVD location of the installation program (by using the `Storage Resource Agent` image) and go to the `bin` directory:
   `cd` *`DVD_image_location`*`/data/sra/`*`operating_system_name`*

   Where *DVD_image_location* is the location of the installation image for the Storage Resource agent.
2. Run the upgrade command:

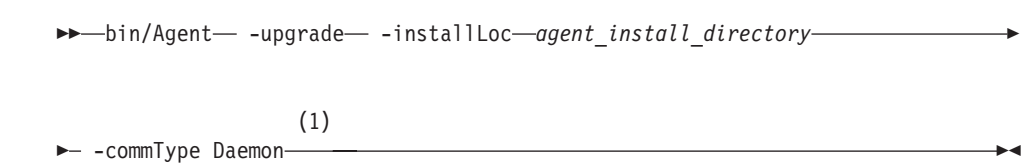

Notes:

1 Parameter when the agent is run as a daemon service.

The parameters are:

-installLoc "*agent_install_directory***"**
Location where the agent is installed. Enclose the directory name in quotation marks, for example, `"C:\Program Files\IBM\TPC_SRA\"`.

-commType Daemon
If the agent is run as a daemon service, then this parameter must be specified.

Here is an example for a daemon-based service by using the default location:

```
bin/Agent -upgrade
-installLoc "/opt/IBM/TPC/"
-commType Daemon
```

Here is an example for a non-daemon service by using a non-default location:

```
bin/Agent
-upgrade -installLoc "C:\Program Files\IBM\TPC_SRA\"
```

Tip: If you run the upgrade program outside of the *DVD_image_location* installation directory, then you must specify the full path.

If the upgrade fails, see the return codes in the Tivoli Storage Productivity Center Information Center. Search for *Return codes used by Storage Resource agent*.

Disabling Storage Resource agents

Disable Storage Resource agents so that they no longer collect data or run Tivoli Storage Productivity Center jobs.

You might want to disable a Storage Resource agent under the following conditions:

- The monitored server is undergoing maintenance and is unavailable. This action prevents Tivoli Storage Productivity Center from flagging the agent as "down" if it cannot reach the agent. The number of times that the server tries to contact the agent is defined by the **agentErrorLimit** parameter in the server.config file.
- The monitored server is busy with resource-intensive processing and you do not want to add any Tivoli Storage Productivity Center jobs to that processing load.

To disable a Storage Resource agent, complete the following steps:

1. From the navigation pane in the web-based GUI, click **Server Resources** > **Servers**.
2. On the Servers page, right-click the server where the agent is deployed and select **Modify Agents** > **Disable**.
3. Click **OK** to confirm that you want to disable the agent. The state of the agent is changed to Disabled and remains in that state until it is enabled again. You can disable agents on multiple servers at the same time.

When you disable a Storage Resource agent that is deployed as a daemon service, the service is shut down, and the agent is disabled. Tivoli Storage Productivity Center no longer sends requests to the agent or contacts it for job processing. A Storage Resource agent that is deployed as a non-daemon agent runs as a stand-alone process. Because a service is not required for this type of agent, it is not necessary to shut down the agent before it is disabled.

Enabling Storage Resource agents

You can enable Storage Resource agents that are in a Disabled or Down state. After an agent is enabled, the Tivoli Storage Productivity Center server resumes communication with that agent.

If the Tivoli Storage Productivity Center server cannot contact an agent, the agent is automatically flagged as "down". You can use the **Enable** action to reestablish

communication between the agent and the Tivoli Storage Productivity Center server. The number of times that the Tivoli Storage Productivity Center server tries to contact the agent is specified in the **agentErrorLimit** parameter in the `server.config` file. The default value for the **agentErrorLimit** parameter is 3.

By default, the `server.config` file is located in the following directory:

Windows

```
C:\Program Files\IBM\TPC\Data\config
```

Linux or UNIX

```
/opt/IBM/TPC/Data/config
```

1. From the navigation pane in the web-based GUI, click **Server Resources** > **Servers**.
2. On the Servers page, right-click the server where the agent is deployed and select **Modify Agents** > **Enable**. You can enable agents on multiple servers at the same time.
3. Click **OK** to confirm that you want to enable the agent.
4. If the agent is running as a daemon service, enter the user ID, password, and other credentials for the server where the agent is deployed. Click **OK** to start the service and enable the agent.

The agent is enabled and the state of the agent is updated to reflect its current condition, such as Up or Upgrade needed. If the agent is deployed as a daemon service, the service is started when you enable the agent.

Testing the connection with a Storage Resource agent

Verify that the Tivoli Storage Productivity Center server can communicate with the server where a Storage Resource agent is deployed.

Use the **Test Connection** action in the web-based GUI to verify the state of the Storage Resource agent. For example, if the agent has a state of Down or Unreachable on the Servers page, you can test the connection to verify the state of the agent.

1. From the navigation pane in the web-based GUI, click **Server Resources** > **Servers**.
2. On the Servers page, right-click the server where the Storage Resource agent is deployed and select **Modify Agents** > **Test Connection**.
3. Optional: If the process is slow, click **Close** in the Testing Agent Connection window to run the operation in the background.

When the operation is complete, the server status and the agent state are automatically updated on the Servers page.

Changing credentials for Storage Resource agents

You can change Storage Resource agent credentials, such as the user name and password that Tivoli Storage Productivity Center uses for logging on to the server where the agent is deployed.

1. From the navigation pane in the web-based GUI, click **Server Resources** > **Servers**.
2. On the Servers page, right-click the server where the agent is deployed and select **Modify Agents** > **Update Credentials**.
3. In the Enter User Credentials window, change the credentials for logging on to the server where the agent is installed. You can change the following credentials:

User name, Password
The user name and password that Tivoli Storage Productivity Center uses for logging on to the server where the Storage Resource agent is deployed. The user name must have administrative or root privileges on the server. This action is available only for Storage Resource agents that were deployed as non-daemon services.

The user name and password must contain valid characters. You can enter the following characters:

- A - Z (uppercase characters)
- a - z (lowercase characters)
- 0 - 9 (numeric characters)
- Series of punctuation marks or special characters: ! # % & * + - / = ? ^ _ { } ( ) . ,

Restrictions:

- User names and passwords cannot contain spaces and must have at least one character.
- The maximum length of a user name or password is 128 characters.

Certificate location
The fully qualified path of the certificate file for the Storage Resource agent, for example, `TPC_install_directory/data/sra/operating_system/certs/sra.pem`. This file is on the computer where the Tivoli Storage Productivity Center server is installed. If the agent uses Secure Shell (SSH) protocol for communication, the certificate location field is displayed.

Passphrase
The passphrase for the certificate file. The passphrase was created when the certificate was generated.

4. Click **OK** to save the changes.

For more information about using certificates after you install a Storage Resource agent, see the Tivoli Storage Productivity Center information center. Search for *Storage Resource agent protocol support*.

Collecting service data

Collect service data about the selected Storage Resource agent. Service data includes diagnostic information such as logs, trace files, configuration information, and computer details. Use this information to troubleshoot any errors that might occur during startup, processing, or shutdown of a Storage Resource agent.

To collect service data for a Storage Resource agent, complete the following steps:

1. In the navigation pane of the web-based GUI, select **Server Resources** > **Servers**.
2. Right-click a server and select **Logs** > **Collect Agent Logs**. A message is displayed that shows the location where the service file is stored on the Tivoli Storage Productivity Center server.
3. In a command line or other navigation tool, go to the directory where the service file is located and unpack its contents.

If the collection of service data is successful, a message is displayed that shows the location of the resulting service file (.zip). The file is stored in a directory on the same computer as the Tivoli Storage Productivity Center server. The file is in the following default directories:

- Windows operating system: `C:\Program Files\IBM\TPC\data\log\SRATraces\`*`agent_computer_name`*`\TPCServiceInfo.zip`
- UNIX or Linux operating system: `/opt/IBM/TPC/data/log/SRATraces/`*`agent_computer_name`*`/TPCServiceInfo.zip`

Where *agent_computer_name* represents the name of the server on which a Storage Resource agent is deployed. If an agent communicates with more that one installation of Tivoli Storage Productivity Center, a subfolder is created for each installation.

If the collection of service data fails, an error message is displayed. For more information about why a data collection failed, see the server log file or the services script. These files are in the following default directories:

- Server log file (on the computer where the Tivoli Storage Productivity Center server is installed):
  - Windows operating system: `c:\Program Files\IBM\TPC\data\log`
  - UNIX or Linux operating system: `/opt/IBM/TPC/data/log`
- Services script file (on the server where the Storage Resource agent is deployed):
  - Windows operating system: `C:\Program Files\IBM\TPC\agent\service\`*`agent_computer_name`*`\TPCServiceInfo.html`
  - UNIX or Linux operating system: `/opt/IBM/TPC/agent/service/`*`agent_computer_name`*`/TPCServiceInfo.html`

  Where *agent_computer_name* represents the name of the server on which the Storage Resource agent is deployed.

Enabling or disabling scripts for Storage Resource agents

You can enable or disable scripts that are sent from the Tivoli Storage Productivity Center server to Storage Resource agents.

If you enable scripts to run, the Storage Resource agent runs the scripts that are sent from the Tivoli Storage Productivity Center server.

If you disable scripts from running, the Storage Resource agent only runs the scripts that are stored on the server where the agent is deployed. The agent does not run scripts that are sent from the Tivoli Storage Productivity Center server.

1. From the navigation pane in the web-based GUI, click **Server Resources** > **Servers**.
2. On the Servers page, right-click the server where the agent is deployed. Select **Modify Agents** > **Enable running scripts on agent** or **Modify Agents** > **Disable running scripts on agent** to enable or disable scripts from running.

Enabling or disabling the monitoring of fabrics by Storage Resource agents

You can enable or disable fabric monitoring by Storage Resource agents. Fabric monitoring is enabled by default. When you enable fabric monitoring, the agent collects information about fabrics that the server is connected to.

After you install a Storage Resource agent on a server, you can enable or disable the monitoring of fabrics that the server is connected to. If you enable fabric monitoring, the agent collects information about the SAN and zoning.

If you disable fabric monitoring, the agent cannot collect fabric information or monitor fabrics that the server is connected to. If the agent is the only data source that is managing the fabric, the fabric is no longer managed. A state of Unreachable is shown for the fabric on the Fabrics page.

1. From the navigation pane in the web-based GUI, click **Server Resources** > **Servers**.
2. On the Servers page, right-click the server where the agent is deployed. Select **Modify Agents** > **Enable Fabric Functions** or **Modify Agents** > **Disable Fabric Functions** to enable or disable fabric monitoring.

Using the help command for Storage Resource agents

The **help** command for Storage Resource agents provides information about the parameters for installing, uninstalling, and upgrading Storage Resource agents.

For information about the Storage Resource agent commands, run the **help** command. Follow these steps:

1. Go to the installation location for the Storage Resource agent:

   ```
   cd <installation_location>
   ```

2. Run the following command:

   ```
   bin/Agent -help
   ```

3. The output from the **help** command is as follows:

   ```
   Usage:
   Agent -INSTALL
         [-COMMTYPE DAEMON -AGENTPORT portnumber]
         [-FORCE]
         -INSTALLLOC  pathname
         -SERVERIP    address[,address,...]
         -SERVERPORT  portnumber
         [-USERID username -PASSWORD password -CERT file -PASSPHRASE phrase]


   Agent -UNINSTALL
         [-FORCE]
         -SERVERNAME servername


   Agent -UPGRADE
         -INSTALLLOC  pathname
   ```

For information about how to install the Storage Resource agent using commands, see the information about installing agents in the *Tivoli Storage Productivity Center Installation and Configuration Guide.*

For information about how to uninstall the Storage Resource agent using commands, see the information about uninstalling agents in the *Tivoli Storage Productivity Center Installation and Configuration Guide.*

For information about how to upgrade the Storage Resource agent using commands, see the information about upgrading agents in the *Tivoli Storage Productivity Center Installation and Configuration Guide.*

Removing servers

Remove servers that you no longer want to monitor with Tivoli Storage Productivity Center.

You can use the web-based GUI to remove servers. If a Storage Resource agent is deployed to the server, the agent is uninstalled.

When the server is removed, it is no longer monitored by Tivoli Storage Productivity Center. All the data that was collected about the server is removed from the database repository.

Tip: When you remove a server, it is only removed from Tivoli Storage Productivity Center. The server is not physically deleted from the storage environment.

To remove a server, complete the following steps:

1. From the navigation pane in the web-based GUI, click **Server Resources** > **Servers**.
2. On the Servers page, right-click the server where the agent is deployed and select **Remove**.
3. Click **Remove** to confirm that you want to remove the server.

Registering the Storage Resource agent with a different server

This topic provides information on how to register the Storage Resource agent with a different server.

To register the Storage Resource agent with a different server, complete the following steps:

1. You have a Storage Resource agent that is installed on Server A and you want the Storage Resource agent to also point to Server B.
2. From Server B, in the navigation tree, expand **Administrative Services** > **Configuration** > **Storage Resource Agent Deployments**. Right-click **Storage Resource Agent Deployments**.
3. Specify the same port and location with the **force** option.

   Note: For an RXA Storage Resource agent, just specify the same location with the **force** option.
4. Run the deployment job. The Storage Resource agent now points to Server A and Server B. If you no longer want the Storage Resource agent to point to Server A, just go to the GUI on Server A and delete the Storage Resource agent.

Manually changing the Windows service logon

Change the Windows service logon for a Storage Resource agent.

To change the Windows service logon for a Storage Resource agent, complete the following steps:

1. Start Windows Services.
2. On the Services window, right-click **IBM Tivoli Storage Resource agent - 'C:\Program Files\IBM\TPC\'**.
3. Select **Properties**.
4. Click the **Log On** tab.
5. Change the values for **This account**, **Password**, and **Confirm password** to the login credentials that you want to use. If your Tivoli Storage Productivity

Center server is part of a Windows domain, change this logon to <domain>\<account>. For example: mydomain\myaccount.

Important: The Storage Resource agent requires that the domain account has local administrator privileges and the "Log on as a service" and "Act as part of the operating system" user rights.

6. Click **Apply** and then **OK** to save your changes.

Managing the daemon Storage Resource agent service on the Virtual I/O Server

Use this information to start and stop the daemon Storage Resource agent service for the Virtual I/O Server.

Starting and stopping the daemon Storage Resource agent service

To start or stop the daemon service, follow these steps:

1. Log in to the Virtual I/O Server using the **padmin** user ID.
2. Run the following command to set up the AIX environment:
   ```
   oem_setup_env
   ```
3. Change to the base directory where the Storage Resource agent is located. For example:
   - To stop the service, run this command:
     ```
     /SRA_install_directory/agent/bin/agent.sh stop
     ```
   - To start the service, run this command:
     ```
     /SRA_install_directory/agent/bin/agent.sh start
     ```

Deployment guidelines and limitations for Storage Resource agents

There are a number of guidelines and limitations that you must consider when managing the Storage Resource agents in your environment.

Capacity guidelines for Storage Resource agents: For the capacity guidelines for Storage Resource agents by Tivoli Storage Productivity Center version, see http://www.ibm.com/support/docview.wss?uid=swg21424912.

Use the following information when you deploy Storage Resource agents:

- Multiple Storage Resource agents that are probing or scanning the same storage resources
- Platforms that support the deployment of Storage Resource agents
- Product functions that are not available for storage devices monitored by Storage Resource agents
- Required authority for deploying Storage Resource agents
- Orphan zones
- Firewalls and Storage Resource agent deployments
- Multiple computers in a deployment
- Communication between the Tivoli Storage Productivity Center server and a Storage Resource agent
- Daemon and non-daemon services
- Authentication between the Tivoli Storage Productivity Center server and a Storage Resource agent
- Storage Resource agents on the same computer

- Time zones for computers monitored by Storage Resource agents
- AIX systems using RSH protocol
- Deployments on Windows - NetBIOS setting
- Deployments on Windows - User Account Control (UAC) remote restrictions
- Additional information about Storage Resource agents

Multiple Storage Resource agents that are probing or scanning the same storage resources

If multiple Storage Resource agents are set up to probe or scan the same storage resources, the Storage Resource agent that was added to Tivoli Storage Productivity Center first is used for the probe or scan. Therefore, only data that is gathered by the first Storage Resource agent is shown.

Platforms that support the deployment of Storage Resource agents

For a list of platforms on which you can deploy Storage Resource agents, go to the Tivoli Storage Productivity Center support site at http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition, click the **Documentation** link, and enter `Platform Support: Agents, Servers and Browsers` in the **Search** box. Click the link to the document for the appropriate release.

Product functions that are unavailable for resources that are monitored by Storage Resource agents

Before you deploy a Storage Resource agent, make sure that the product functions you want to use on the monitored devices are available for those agents. The following functions are not available for resources that are monitored by Storage Resource agents:

- Certain relational database monitoring. For list of relational databases that can be monitored by Storage Resource agents, go to the Tivoli Storage Productivity Center Interoperability Matrix and click the link for *Agents, Servers and Browsers*.
- The reporting of HBA, fabric topology, or zoning information for fabrics that are connected to hosts that are running Linux on IBM System z hardware. These limitations also apply to Storage Resource agents on all guest operating systems for VMware configurations.

Required authorities for deploying and running Storage Resource agents

Before you can create deployment schedules and deploy Storage Resource agents on target computers, you must meet the following requirements:

- To create deployment schedules, you must be logged in to Tivoli Storage Productivity Center with a user ID that has the **Administrator** role. For more information about user roles, see "Role-based authorization" on page 77.
- To deploy Storage Resource agents on target computers, you must provide a user ID that has administrative rights on those computers. You enter this ID when you create a deployment schedule. Tivoli Storage Productivity Center uses this ID to log on to the target computers and install and configure the necessary runtime files for the agents.

The user under which a Storage Resource agent (daemon or non-daemon) runs must have the following authorities on the target computers:

- On the Linux or AIX operating systems, the user must have root authority. By default, an agent runs under the user 'root'.

- On the Windows operating systems, the user must have Administrator authority and be a member of the Administrators group. By default, a Storage Resource agent runs under the 'Local System' account.

Orphan zones

Storage Resource agents do not collect information about orphan zones. An orphan zone is a zone that does not belong to at least one zoneset.

Firewalls and Storage Resource agent deployments

Before you can deploy a Storage Resource agent on a computer, you must turn off the firewall on that computer. If you do not turn off the firewall, the deployment fails.

To turn off the firewall on a Windows 2008 computer, complete the following steps:

1. Open Administrative Tools.
2. Click **Windows Firewall with Advanced Security**.
3. Click **Windows Firewall Properties**.
4. Change the **Firewall state** field to **Off** on the following tabs:
   - **Domain Profile**
   - **Private Profile**
   - **Public Profile**
5. Click **OK** to accept the changes and exit.
6. Deploy a Storage Resource agent to the Windows 2008 computer.

Multiple computers in a deployment

You can define a deployment schedule to include multiple computers on which to install Storage Resource agents. The computers that you include in a deployment schedule must share administrative user ID and password. Tivoli Storage Productivity Center uses these user credentials to log on to the computers when you installing Storage Resource agents. If the computers in a deployment schedule do not share administrative user credentials, you must create separate deployment jobs for them.

Tip: When you deploy a Storage Resource agent to multiple computers, a globally unique identifier (GUID) is created for each computer (if one does not exist).

Communication between the Tivoli Storage Productivity Center server and a Storage Resource agent

The Tivoli Storage Productivity Center server connects to a monitored computer when a Storage Resource agent is deployed and whenever a data collection schedule runs against that agent.

During deployment, the server communicates with the target computer by using one of the following protocols:

- Windows server message block protocol (SMB)
- Secure Shell protocol (SSH)
- Remote execution protocol (REXEC)
- Remote shell protocol (RSH)

After deployment, the type of communication between the server and agent on that computer depends on whether you deployed the agent as daemon service or non-daemon service. See the *IBM Tivoli Storage*

Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide for more information about these protocols.

Daemon and non-daemon services

You can deploy a Storage Resource agent agent as a daemon or non-daemon service:

- A Storage Resource agent that is deployed as a daemon service runs in the background on the monitored computer and listens for requests from the Tivoli Storage Productivity Center server. Connectivity between the server and agent is established by using SSL. The server and agent have their respective certificates and no additional information is required besides those certificates and the security that is provided by the SSL protocol.
- A Storage Resource agent deployed as a service on demand (non-daemon service) runs as a stand-alone executable file on the monitored computer. Communication from the server to the agent uses the same protocol that was used during the deployment of the agent. Communication from the agent to the server uses SSL.

Authentication between the Tivoli Storage Productivity Center server and a Storage Resource agent

Tivoli Storage Productivity Center requires the correct authentication information (user name, password, port, certificate location, or passphrase) for monitored computers each time it communicates with Storage Resource agents on those computers. If the authentication information changes for a host computer on which a Storage Resource agent is deployed, the authentication information for that agent must be updated by using the **Modify Agents** > **Update Credentials** action on the Servers page in the web-based GUI.

Storage Resource agents on the same computer

You cannot deploy a Storage Resource agent on a computer where a Storage Resource agent is already installed *and* pointing to the same Data server. You can deploy a Storage Resource agent on the same computer as another Storage Resource agent if those agents communicate with different Data servers and use different ports when you listen for requests.

Time zones for computers that are monitored by Storage Resource agents

The time zones of computers that are monitored by Storage Resource agents are shown as Greenwich Mean Time (GMT) offsets in Tivoli Storage Productivity Center reports. For example, a computer in Los Angeles shows the following time zones in the By Computer report in Asset reporting:

```
(GMT-8:00) GMT-8:00
```

Linux and AIX operating systems by using Remote Shell protocol (RSH)

If RSH is configured to use a user ID and password, the connection fails. To successfully connect to a system by using RSH, you must set up the `.rhosts` file (in the home directory of the account). RSH must be configured to accept a login from the system that is running your application.

Deployments on Windows operating systems - NetBIOS setting

If you want to install a Storage Resource agent on Windows targets, the **Enable NetBIOS over TCP/IP** option must be selected in the Control Panel settings for the computer's network connections properties. To set this option, complete the following steps:

1. Open Windows Control Panel.
2. Select **Network and Dial-Up Connections** > *some_connection* > **Properties** > **Internet Protocol (TCP/IP)** > **Advanced** > **WINS** > **Enable NetBIOS over TCP/IP**.

See the documentation for your firewall to determine if these ports are not blocked for inbound requests.

To determine if security policies are blocking the connection ports, open Administrative Tools.

Depending on whether your policies are stored locally or in Active Directory, follow these directions:

Locally stored policies

For locally stored policies, complete the following steps:

1. Open Windows Administrative Services.
2. Click **Local Security Policy** > **IP Security Policies on Local Computer**.

Policies stored in Active Directory

For policies stored in Active Directory, examine the IP security policies and edit or remove filters that block the ports:

- Click **Administrative Tools** > **Default Domain Security Settings** > **IP Security Policies on Active Directory**.
- Click **Administrative Tools** > **Default Domain Controller Security Settings** > **IP Security Policies on Active Directory**.

For all Windows systems, the Server service must be running to connect to a Windows system by using the Windows protocol.

The following table lists the ports reserved for NetBIOS. Ensure that these ports are not blocked.

Port	Description
135	NetBIOS Remote procedure call. (Not currently used.)
137	NetBIOS name service.
138	NetBIOS datagram. (Not currently used.)
139	NetBIOS session (for file and print sharing).
445	CIFS (on Windows XP).

For Windows Server 2008, shares must be shared for the Guest or Everyone accounts, and password protected sharing must be disabled. To disable password protected sharing, follow these steps:

1. Click **Control Panel** > **Networking and Sharing Center**.
2. Click the down arrow next to **Password protected sharing**.
3. Click **Turn off password protected sharing**.
4. Click **Apply**.
5. Exit from the Control Panel.

Deployments on Windows 2008 - User Account Control (UAC) remote restrictions

If you are planning to install Storage Resource agents remotely on a Windows 2008 operating system, you must disable the User Account Control (UAC) remote restrictions on the Windows operating system. User Account Control is a security component on Windows operating systems.

Tip: This task contains steps that tell you how to modify the registry. Serious problems might occur if you modify the registry incorrectly. Therefore, make sure that you follow these steps carefully. For added protection, back up the registry before you modify it. Then you can restore the registry if problems occur. For information about how to back up and restore the registry, see http://support.microsoft.com/kb/322756/.

To disable UAC remote restrictions, follow these steps:

1. Open the Windows Run window.
2. Enter **regedit** and click **OK**.
3. Locate and click the following registry subkey:
   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\
       Policies\System
   ```
4. If the `LocalAccountTokenFilterPolicy` registry entry does not exist, follow these steps:
   a. On the **Edit** menu, click **New** > **DWORD Value**.
   b. Enter `LocalAccountTokenFilterPolicy` as the name for the DWORD value in the content pane.
   c. Click **Enter**.
   d. Click **LocalAccountTokenFilterPolicy**, and click **Modify**.
   e. In the Edit DWORD Value window, enter **1**, then click **OK**.
      This value can be 0 or 1:

      0 This value builds a filtered token. This value is the default value. The administrator credentials are removed.

      1 This value builds an elevated token.
   f. Exit the registry editor.

Additional information about Storage Resource agents

See *Planning for Storage Resource agents* in the *IBM Tivoli Storage Productivity Center Installation and Configuration Guide* for more information about Storage Resource agents.

Hypervisors and VMware data sources

Administer the hypervisors and VMware vCenter servers that are monitored by Tivoli Storage Productivity Center. VMware vCenter servers are data sources that can monitor multiple hypervisors. A hypervisor can be an ESX or ESXi host. Each hypervisor can host multiple virtual machines.

Adding hypervisors

Add hypervisors for monitoring. Hypervisors include VMware vSphere data sources such as ESX, ESXi, and vCenter servers.

You can add a hypervisor by specifying connection information for that hypervisor, or you can add multiple hypervisors by specifying connection information for a vCenter Server. For a complete list of hypervisors and their versions that you can add, see the *Agents, Servers, and Browsers* section in the *Tivoli Storage Productivity Center Interoperability Matrix*.

Help tips in the GUI: To view descriptions of the information that you must enter for a hypervisor, hover the mouse pointer over the related help icons in the wizard.

1. In the navigation pane of the web-based GUI, select **Server Resources** > **Hypervisors**.
2. Click **Add Hypervisor**.
3. Enter connection information about the hypervisor. If you enter connection information about a vCenter server, Tivoli Storage Productivity Center automatically discovers and lists the hypervisors that it manages.

   Restriction: The user name that you specify for a hypervisor must have permission to browse through the data stores on VMware. For information about how to check permissions, see "Checking permissions to browse data stores" on page 272.
4. Specify display information and schedule a probe of the hypervisor. When you add a vCenter server, specify display information and schedule a probe for the hypervisors that were discovered. If you select multiple hypervisors at the same time, you cannot specify display information.
5. Complete the wizard.

   Tip: In previous versions of Tivoli Storage Productivity Center, you were required to download the certificate from the hypervisor or vCenter server and add it to the vmware.jks keystore of the Device server. In V5.2 and later, the web-based GUI completes this process automatically.

After a hypervisor is added, a probe collects status and asset information about that hypervisor. You can view detailed information about the hypervisor in the web-based GUI and stand-alone GUI.

Viewing information about hypervisors and VMware VI data sources

View detailed information about hypervisors and VMware vCenter servers that are monitored by Tivoli Storage Productivity Center.

Use the following locations in the Tivoli Storage Productivity Center GUIs to view information about hypervisors and VMware vCenter servers.

Web-based GUI

To view information about hypervisors, complete the following steps:

1. In the navigation pane of the web-based GUI, select **Server Resources** > **Hypervisors**. Information about monitored hypervisors is displayed.
2. On the resource list page, right-click a hypervisor and select **View Properties** to view the key properties of a hypervisor.

Stand-alone GUI

Hypervisors and vCenter servers are displayed in the stand-alone GUI under the following conditions:

- If you add hypervisors in the web-based GUI by discovering them through a vCenter server, only information about the vCenter server is displayed in the stand-alone GUI. Information about the hypervisors is displayed in the web-based GUI.
- If you add a hypervisor as a data source in the web-based GUI, information about that hypervisor is also displayed in the stand-alone GUI.

To view information about hypervisors and vCenter servers, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Administrative Services** > **Data Sources**.
2. Click **VMware VI Data Source**. Information about monitored hypervisors and vCenter servers is displayed. To view more information about a specific hypervisor or vCenter server, click the magnifying glass that is displayed on its row.

Updating the credentials for a hypervisor

Change the user name and password that Tivoli Storage Productivity Center uses to log in to a hypervisor.

1. In the navigation pane of the web-based GUI, select **Server Resources** > **Hypervisors**.
2. Right-click a hypervisor and select **Connections** > **Update Credentials**.
3. Update the user name and password for the hypervisor. The user name and password for a hypervisor must contain valid characters. You can enter the following characters when you update user credentials:
   - A through Z (uppercase characters)
   - a through z (lowercase characters)
   - 0 through 9 (numeric characters)
   - Special characters: ! # % & * + - / = ? ^ _ { } ( ) . ,

   Restrictions:
   - User names and passwords cannot contain spaces and must have at least one character.
   - The maximum length of a user name or password is 128 characters.
   - The user name must have permission to browse the data stores on a hypervisor. For more information about permissions, see "Checking permissions to browse data stores" on page 272.
4. Click **OK** to apply the changes.

Viewing the hypervisors that are managed by a VMware VI data source

View information about the hypervisors that are managed by a vCenter server.

To view the hypervisors that are managed by a vCenter server, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Administrative Services** > **Data Sources** > **VMware VI Data Source**.
2. In the content pane, click **Show Managed Devices**. A list of the hypervisors that are managed by the vCenter server is displayed.

Removing hypervisors and VMware VI data sources

Remove hypervisors and VMware vCenter servers that you no longer want to monitor with Tivoli Storage Productivity Center.

You can use the web-based GUI and stand-alone GUI to remove hypervisors and vCenter servers. The type of resource that you want to delete determines which GUI to use:

- To remove hypervisors, use the web-based GUI.
- To remove vCenter servers, use the stand-alone GUI.

Tips:

- When a vCenter server manages hypervisors, Tivoli Storage Productivity Center considers that vCenter server to be a data source for the hypervisors. In V5.2 and earlier, use the stand-alone GUI when removing data sources.
- When you remove a resource, it is only removed from Tivoli Storage Productivity Center; it is not physically deleted from a storage environment.

Removing hypervisors

To remove a hypervisor in the web-based GUI, complete the following steps:

1. In the navigation pane of the web-based GUI, select **Server Resources** > **Hypervisors**.
2. Right-click a hypervisor and select **Remove**. The hypervisor and all its data are removed from Tivoli Storage Productivity Center immediately. Any data collection jobs and alerts are also removed.

   Restriction: If the data source for a hypervisor is a vCenter server, you must use the stand-alone GUI to remove that server from Tivoli Storage Productivity Center.

Removing vCenter servers

To remove a vCenter server, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Administrative Services** > **Data Sources**.
2. Click **VMware VI Data Source**.
3. Select the row of the vCenter server that you want to remove.
4. Click **Remove VMware VI Data Source**.
5. In the confirmation dialog window, click **OK**.

When you remove a vCenter server, the hypervisors that it manages are also removed from Tivoli Storage Productivity Center. However, information about the hypervisors is not removed immediately, but is retained according to the settings on the Removed Resource Retention panel in the stand-alone GUI. The default setting is 14 days. If the default setting is used, all information about the hypervisors is deleted 14 days after that related vCenter server was removed.

Tips: After a vCenter server is removed, but before its managed hypervisors are removed according to the retention settings, the following conditions occur:

- Any data collection jobs that are scheduled for the hypervisors fail.
- Because data is no longer collected, any alerts that were based on that data are not generated.

Storage systems

Administer the storage systems that are monitored by IBM Tivoli Storage Productivity Center. Administering actions include adding and removing storage systems, updating credentials, and testing connections.

Adding storage systems

Add storage systems for monitoring.

You can add the following types of storage systems for monitoring:

- DS4000
- DS5000

- DS6000
- DS8000
- XIV system
- ESS
- SAN Volume Controller
- Storwize V3500, Storwize V3700, Storwize V7000
- Storwize V7000 Unified
- IBM SONAS
- All others (managed by CIM agents)

Tips:

- For a complete list of the versions of storage systems and CIM agents that you can add, see Tivoli Storage Productivity Center Interoperability Matrix and go to the *Storage* section.
- The storage system that you want to add determines which Tivoli Storage Productivity Center GUI to use. Use the stand-alone GUI to add IBM SONAS storage systems for monitoring. Use the web-based GUI to add all other storage systems and CIM agents.

After a storage system is added, data collection jobs gather information about that storage system. When data collection is complete, you can view detailed information about the storage system in the web-based GUI and stand-alone GUI. Storage systems can be configured for block storage, file storage, or a combination of block and file storage. The method that a storage system uses for managing data determines the internal resources that are monitored.

Related concepts:

"Collecting data" on page 295
Determining the data that you want to gather about storage resources is critical to helping you implement a storage management strategy. Tivoli Storage Productivity Center provides different jobs for collecting data. When you determine the types of data that you want to collect about resources, schedule the corresponding data collection jobs for those resources.

"Collecting asset and status data" on page 296
Use probes to collect asset, status, and storage data about monitored resources in your environment. Probes are also used to discover information about new or removed disks and file systems. You can view information about probes on the resource list and resource details pages, for example, the Storage Systems page, or the Storage Systems details page.

Related reference:

"CIM agents" on page 205
Administer CIM agents that are associated with storage resources that are monitored by Tivoli Storage Productivity Center. Common Information Model (CIM) agents provide a CIM interface for collecting information about certain types of storage systems and switches.

Viewing information about storage systems

View detailed information about storage systems that are monitored by Tivoli Storage Productivity Center.

Use the following locations in the Tivoli Storage Productivity Center GUIs to view information about storage systems.

Web-based GUI

To view information about storage systems, complete the following steps:

1. In the navigation pane of the web-based GUI, select **Server Resources** > **Storage Systems**. Information about monitored storage systems is displayed.
2. Click the **Block Storage** or **File Storage** tab. Depending on how a storage system is configured, information about that storage system might be shown in one or both of these tabs:

   Block Storage tab only
   Storage systems that are configured for storing or retrieving data only in block format include System Storage DS series, SAN Volume Controller, Storwize V7000, and other SAN-based storage systems.

   File Storage tab only
   Storage systems that are configured for storing or retrieving data only in file format include IBM SONAS and other Network Attached Storage (NAS) based storage systems.

   Block Storage and File Storage tabs
   Storage systems that can be configured for both file and block data include Storwize V7000 Unified and NetApp Filers.
3. Right-click a storage system and select **View Properties** to view the key properties for the system.

Stand-alone GUI

To view information about storage systems, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Administrative Services** > **Data Sources**.
2. Click **Storage Systems**. Information about monitored storage systems is displayed. To view more information about a specific storage system, click the magnifying glass that is displayed on its row.

Updating the credentials for storage systems

Change the credentials that Tivoli Storage Productivity Center uses to authenticate to a storage system or the CIM agent that manages a storage system.

If the storage system is managed by multiple data sources, for example multiple CIM agents, the menu is displayed as **Connections** > **Update Credentials** > ***data_sources***. Select the data source for which you want to update the credentials.

The type of storage system determines the credentials that you can update.

Rules for user names and passwords:

- User names and passwords cannot contain spaces and must have at least one character. Some newer versions of virtualizer firmware require a minimum of 8 characters for a password.
- The maximum length of a user name or password is 128 characters.
- You can enter the following characters for a user name and password:
  - A through Z (uppercase characters)
  - a through z (lowercase characters)
  - 0 through 9 (numeric characters)
  - Special characters: ! # % & * + - / = ? ^ _ { } ( ) .,

Updating the credentials for a System Storage DS8000 storage system:

Change the credentials that Tivoli Storage Productivity Center uses to authenticate to a System Storage DS8000 storage system.

To update the credentials for a System Storage DS8000 storage system, complete the following steps. You can update the IP address or host name for the secondary HMC that is used to manage the storage system, the user name, and the password.

1. In the navigation pane of the web-based GUI, click **Storage Resources** > **Storage Systems**. Information about monitored storage systems is displayed.
2. Right-click a storage system and click **Connections** > **Update Credentials**.
3. Change the secondary HMC host name or IP address, the user name, or the password, and then click **OK**.

Updating the credentials for an XIV system:

Change the credentials that Tivoli Storage Productivity Center uses to authenticate to an XIV system.

To update the credentials for an XIV system, complete the following steps. You can update the user name and the password.

1. In the navigation pane of the web-based GUI, click **Storage Resources** > **Storage Systems**. Information about monitored storage systems is displayed.
2. Right-click a storage system and click **Connections** > **Update Credentials**.
3. Change the user name or password, and then click **OK**.

Updating the credentials for a SAN Volume Controller or Storwize family storage system:

Change the credentials that Tivoli Storage Productivity Center uses to authenticate to a SAN Volume Controller or Storwize family storage system.

To update the credentials for a SAN Volume Controller or Storwize family storage system, complete the following steps:

1. In the navigation pane of the web-based GUI, click **Storage Resources** > **Storage Systems**. Information about monitored storage systems is displayed.
2. Right-click a storage system and click **Connections** > **Update Credentials**.
3. Update the following credentials as required and then click **OK**:

   Authentication
   You can use a user name and password or a private Secure Shell (SSH) key to log on to the storage system. The authentication method that you select determines the options that are displayed.

   User Name/Password
   The user name and password for logging on to the storage system.

   Secure Shell (SSH)
   Use an existing SSH key or upload a new key to the storage system. Select one of the following actions:

   Use an existing SSH key
   Use an SSH key that was uploaded to the storage

system by using a method other than through Tivoli Storage Productivity Center, such as the storage system web interface.

SSH key
The location of the SSH key.

The default location is `${device.conf}\tpc_svc.pem`, which represents the Tivoli Storage Productivity Center default key file `tpc_svc.pem`. The `tpc_svc.pem` file is in the `conf` directory where the Device server is installed.

You can enter another location or select **Browse** to search for a key file. If you select **Browse**, the following fields are displayed:

Select file
The location of the SSH key file. You can click **Browse** to search for a file.

Passphrase
The passphrase for the SSH key pair. If you do not have a passphrase, leave this field blank.

The SSH key file is transferred from the computer where the web browser is located to the computer where the Tivoli Storage Productivity Center server is located.

Upload a new SSH key
Provide the following information to upload an SSH key to the storage system:

SSH key
The location of the SSH key. The key must exist on the system where you are running the Tivoli Storage Productivity Center user interface.

The SSH key must be in OpenSSH format or in PuTTY (`.ppk`) format that is not password protected.

Passphrase
The passphrase for the SSH key pair. If you do not have a passphrase, leave this field blank.

User name, Password
The name and password for a user that belongs to the storage system Security Administrator role for the cluster that contains the storage system.

Tivoli Storage Productivity Center uses this value to configure the SSH key for the user that is entered in the **Associate user** field. The user name that is entered in the **User name** field must have privileges to modify other user accounts, otherwise Tivoli Storage Productivity Center cannot configure the SSH key.

Associate user
The user that is associated with the SSH key.

You can click **Get Users** to retrieve all of the existing users from the storage system. You must select a user that belongs to the storage system Administrator role.

The SSH key file is transferred from the computer where the web browser is located to both the computer where the Tivoli Storage Productivity Center server is located and to the storage system.

Updating the credentials for a Storwize V7000 Unified storage system:

Change the credentials that Tivoli Storage Productivity Center uses to authenticate to a Storwize V7000 Unified storage system.

To update the credentials for a Storwize V7000 Unified storage system, complete the following steps:

1. In the navigation pane of the web-based GUI, click **Storage Resources** > **Storage Systems**. Information about monitored storage systems is displayed.
2. Right-click a storage system and click **Connections** > **Update Credentials**.
3. Update the following credentials as required, and then click **OK**:

Authentication
You can use a user name and password or a private Secure Shell (SSH) key to log on to the storage system. The authentication method that you select determines the options that are displayed.

User Name/Password
The user name and password for logging on to the storage system.

Secure Shell (SSH)
Use an existing SSH key or upload a new key to the storage system. Select one of the following actions:

Use an existing SSH key
Use an SSH key that was uploaded to the storage system by using a method other than through Tivoli Storage Productivity Center, such as the storage system web interface.

SSH key
The location of the SSH key.

The default location is `${device.conf}\tpc_svc.pem`, which represents the Tivoli Storage Productivity Center default key file `tpc_svc.pem`. The `tpc_svc.pem` file is in the `conf` directory where the Device server is installed.

You can enter another location or select **Browse** to search for a key file. If you select **Browse**, the following fields are displayed:

Select file
The location of the SSH key file. You can click **Browse** to search for a file.

Passphrase
The passphrase for the SSH key pair. If you do not have a passphrase, leave this field blank.

The SSH key file is transferred from the computer where the web browser is located to the computer where the Tivoli Storage Productivity Center server is located.

Upload a new SSH key
Provide the following information to upload an SSH key to the storage system:

SSH key
The location of the SSH key. The key must exist on the system where you are running the Tivoli Storage Productivity Center user interface.

The SSH key must be in OpenSSH format or in PuTTY (`.ppk`) format that is not password protected.

Passphrase
The passphrase for the SSH key pair. If you do not have a passphrase, leave this field blank.

User name, Password
The name and password for a user that belongs to the storage system Security Administrator role for the cluster that contains the storage system.

Tivoli Storage Productivity Center uses this value to configure the SSH key for the user that is entered in the **Associate user** field. The user name that is entered in the **User name** field must have privileges to modify other user accounts, otherwise Tivoli Storage Productivity Center cannot configure the SSH key.

Associate user
The user that is associated with the SSH key.

You can click **Get Users** to retrieve all of the existing users from the storage system. You must select a user that belongs to the storage system Administrator role.

The SSH key file is transferred from the computer where the web browser is located to both the computer where the Tivoli Storage Productivity Center server is located and to the storage system.

Use different authentication credentials for file storage
Storwize V7000 Unified contains block-level and file-level data. If the credentials are different for block storage and file storage, select this check box to define the credentials for file storage. The options and fields that are displayed are described previously under **Authentication**.

Updating the credentials for an IBM SONAS storage system:

Change the credentials that Tivoli Storage Productivity Center uses to authenticate to an IBM SONAS storage system.

To update the credentials for an IBM SONAS storage system, complete the following steps:

1. In the navigation pane of the web-based GUI, click **Storage Resources** > **Storage Systems**. Information about monitored storage systems is displayed.
2. On the **File Storage** tab, right-click a storage system and click **Connections** > **Update Credentials**.
3. Update the following credentials as required, and then click **OK**.

Authentication
You can use a user name and password or a private Secure Shell (SSH) key to log on to the storage system. The authentication method that you select determines the options that are displayed.

User Name/Password
The user name and password for logging on to the storage system.

Secure Shell (SSH)
Use an existing SSH key or upload a new key to the storage system. Select one of the following actions:

Use an existing SSH key
Use an SSH key that was uploaded to the storage system by using a method other than through Tivoli Storage Productivity Center, such as the storage system web interface.

SSH key
The location of the SSH key.

The default location is `${device.conf}\tpc_svc.pem`, which represents the Tivoli Storage Productivity Center default key file `tpc_svc.pem`. The `tpc_svc.pem` file is in the `conf` directory where the Device server is installed.

You can enter another location or select **Browse** to search for a key file. If you select **Browse**, the following fields are displayed:

Select file
The location of the SSH key file. You can click **Browse** to search for a file.

Passphrase
The passphrase for the SSH key pair. If you do not have a passphrase, leave this field blank.

The SSH key file is transferred from the computer where the web browser is located to the computer where the Tivoli Storage Productivity Center server is located.

Upload a new SSH key

Provide the following information to upload an SSH key to the storage system:

SSH key

The location of the SSH key. The key must exist on the system where you are running the Tivoli Storage Productivity Center user interface.

The SSH key must be in OpenSSH format or in PuTTY (`.ppk`) format that is not password protected.

Passphrase

The passphrase for the SSH key pair. If you do not have a passphrase, leave this field blank.

User name, Password

The name and password for a user that belongs to the storage system Security Administrator role for the cluster that contains the storage system.

Tivoli Storage Productivity Center uses this value to configure the SSH key for the user that is entered in the **Associate user** field. The user name that is entered in the **User name** field must have privileges to modify other user accounts, otherwise Tivoli Storage Productivity Center cannot configure the SSH key.

Associate user

The user that is associated with the SSH key.

You can click **Get Users** to retrieve all of the existing users from the storage system. You must select a user that belongs to the storage system Administrator role.

The SSH key file is transferred from the computer where the web browser is located to both the computer where the Tivoli Storage Productivity Center server is located and to the storage system.

Updating the credentials for a storage system that is managed by a CIM agent:

Change the credentials that Tivoli Storage Productivity Center uses to authenticate to a CIM agent.

Tivoli Storage Productivity Center communicates with CIM agents to collect information about the following resources:

- IBM TotalStorage Enterprise Storage Server
- IBM System Storage DS4000
- IBM System Storage DS5000 Storage Manager
- IBM System Storage DS6000
- Non-IBM storage systems that are managed by SMI-S certified Common Information Model Object Manager (CIMOM), such as EMC, Hitachi, and NetApp
- Switches: Brocade, Cisco, QLogic

To update the credentials for a CIM agent, complete the following steps. You can update the user name and the password.

1. In the navigation pane of the web-based GUI, click **Storage Resources** > **Storage Systems**. Information about monitored storage systems is displayed.
2. Right-click a storage system and click **Connections** > **Update Credentials**.
3. Change the user name or password, and then click **OK**.

Testing the connection to a storage system

Verify that Tivoli Storage Productivity Center can communicate with a monitored storage system. For storage systems that are managed by a CIM agent or Storage Resource agent, the connection to the agent is tested.

To test the connection to a storage system, complete the following steps:

1. In the navigation pane of the web-based GUI, click **Storage Resources** > **Storage Systems**. Information about monitored storage systems is displayed.
2. On the **Block Storage** or **File Storage** tab, right-click a storage system and click **Connections** > **Test Connection**. A message that shows the results of the test is displayed.

Removing a storage system

Remove a storage system that you no longer want to monitor with Tivoli Storage Productivity Center.

To remove a storage system, complete the following steps:

1. From the navigation pane in the web-based GUI, click **Storage Resources** > **Storage Systems**. Information about monitored storage systems is displayed.
2. On the **Block Storage** or **File Storage** tab, right-click a storage system and select **Remove**.
3. Click **Remove** to confirm that you want to remove the storage system.

Switches and fabrics

Administer the switches and fabrics that are monitored by IBM Tivoli Storage Productivity Center. Administering actions include adding and removing switches and fabrics, updating credentials, and testing connections.

Adding fabrics and switches

Add fabrics and switches for monitoring.

In the web-based GUI, you can add switches and fabrics for monitoring at the same time. You can add the following types of switches:

- Brocade
- Cisco
- QLogic

For a complete list of the switches and their versions that you can add, check the switches and directors information in the *Tivoli Storage Productivity Center Interoperability Matrix*.

Help tips in the GUI: To view descriptions of the information that you must enter for a fabric or switch, hover the mouse pointer over the related help icons in the wizard.

1. In the navigation pane, select **Network Resources** > **Switches**.
2. Click **Add Switch**.

3. Determine how to discover the switch and connected fabric.
   - Select **Use CIM agent as data source** to use a CIM agent as the data source. Tivoli Storage Productivity Center automatically discovers the switches that are managed by the agent. After this step, you can also select to use Simple Network Management Protocol (SNMP) agents for the switches.
   - Do not select **Use CIM agent as data source** to skip the use of a CIM agent as a data source. If you skip the use of a CIM agent as a data source, you are prompted to add switches and their fabrics through SNMP agents.

   The type of switch and functions that you want to enable on that switch and connected fabric determine the agent to use for monitoring. For example, to monitor the performance of a switch, you must use a CIM agent. To ensure full monitoring and functions for fabrics, you must also install Storage Resource agent on a server that is connected to the fabric.

   For more information about agent types and the fabric and switch functions that they enable, see "Agent types for switch and fabric functions" on page 294.
4. Specify the display name and location of the switches and fabrics that are discovered. For switches that were discovered through a CIM agent, you can also select to monitor them through SNMP.
5. Schedule data collection for the switches and fabric. The type of resource determines the data collection jobs that you can schedule:
   - Fabrics: probes
   - Switches: probes, performance monitors

   Probes collect status and asset information. Performance monitors collect metrics that measure performance.
6. Complete the wizard.

After a switch or fabric is added, data collection jobs gather information about them. When data collection is complete, you can view detailed information about the switch or fabric in the web-based GUI and stand-alone GUI.

Viewing information about switches and fabrics

View detailed information about switches and fabrics that are monitored by Tivoli Storage Productivity Center.

To view information about switches and fabrics, complete the following steps:

1. In the navigation pane of the web-based GUI, click **Network Resources** > **Switches** or **Network Resources** > **Fabrics**. Information about monitored switches or fabrics is displayed.
2. Right-click a switch or fabric and click **View Properties** to view the key properties for the switch or fabric.

Updating the credentials for switches and fabrics

Change the credentials that Tivoli Storage Productivity Center uses to authenticate to a data source that manages a switch or fabric.

Depending on the functions that you want to enable, you can use a CIM agent, Storage Resource agent, or SNMP agent to manage switches. You can use a CIM agent or Storage Resource agent to manage fabrics.

If the switch or fabric is managed by multiple data sources, for example multiple CIM agents, the menu is displayed as **Connections** > **Update Credentials** > ***data_sources***. Select the data source for which you want to update the credentials.

The type of data source for the fabric or switch determines the credentials that you can update.

Rules for user names and passwords:

- User names and passwords cannot contain spaces and must have at least one character. Some newer versions of virtualizer firmware require a minimum of 8 characters for a password.
- The maximum length of a user name or password is 128 characters.
- You can enter the following characters for a user name and password:
  - A through Z (uppercase characters)
  - a through z (lowercase characters)
  - 0 through 9 (numeric characters)
  - Special characters: ! # % & * + - / = ? ^ _ { } ( ) .,

Updating the credentials for a switch:

Change the credentials that Tivoli Storage Productivity Center uses to authenticate to a data source that manages a switch. The data source can be a CIM agent, Storage Resource agent, or SNMP agent.

To update the credentials for a switch, complete the following steps:

1. In the navigation pane of the web-based GUI, click **Network Resources** > **Switches**. Information about monitored switches is displayed.
2. Right-click a switch and click **Connections** > **Update Credentials**.
3. Update the following credentials as required and then click **OK**. The credentials that are displayed depend on the data source.

   CIM agent
   Change the user name or password for the CIM agent.

   Storage Resource agent
   Change the user name or password for the Storage Resource agent. You can also change the certificate location and the passphrase for the agent.

   Certificate location
   The fully qualified path of the certificate file for the Storage Resource agent. For example, *TPC_install_directory*/data/sra/*operating_system*/certs/sra.pem. This file is on the computer where the Tivoli Storage Productivity Center Data server is located.

   Passphrase
   The passphrase for the certificate file. The passphrase was created when the certificate was generated.

   SNMP agent
   Change the SNMP community string. The default is **public**.

Updating the credentials for a fabric:

Change the credentials that Tivoli Storage Productivity Center uses to authenticate to a data source that manages a fabric. The data source can be a CIM agent or Storage Resource agent.

To update the credentials for a fabric, complete the following steps:

1. In the navigation pane of the web-based GUI, click **Network Resources** > **Fabrics**. Information about monitored fabrics is displayed.
2. Right-click a fabric and click **Connections** > **Update Credentials**.
3. Update the following credentials as required and then click **OK**. The credentials that are displayed depend on the data source.

   CIM agent
   Change the user name or password for the CIM agent.

   Storage Resource agent
   Change the user name or password for the Storage Resource agent. You can also change the certificate location and the passphrase for the agent.

   Certificate location
   The fully qualified path of the certificate file for the Storage Resource agent. For example, *TPC_install_directory*/data/sra/*operating_system*/certs/sra.pem. This file is on the computer where the Tivoli Storage Productivity Center Data server is located.

   Passphrase
   The passphrase for the certificate file. The passphrase was created when the certificate was generated.

Testing the connection to a switch or fabric

Verify that Tivoli Storage Productivity Center can communicate with the data source that manages a switch or fabric.

To test the connection to the data source that manages a switch or fabric, complete the following steps:

1. From the navigation pane in the web-based GUI, click **Network Resources** > **Switches** or **Network Resources** > **Fabrics**. Information about monitored switches or fabrics is displayed.
2. Right-click a switch or fabric and click **Connections** > **Test Connection**. A message that shows the results of the test is displayed.

Removing a switch or fabric

Remove a switch or fabric that you no longer want to monitor with Tivoli Storage Productivity Center.

To remove a switch or fabric, complete the following steps:

1. From the navigation pane in the web-based GUI, click **Network Resources** > **Switches** or **Network Resources** > **Fabrics**. Information about monitored fabrics or switches is displayed.
2. Right-click a fabric or switch and select **Remove**.
3. Follow the directions that are presented in the information message.

Out-of-band fabric agents

Out-of-band agents are switches and directors that communicate with Tivoli Storage Productivity Center through SNMP.

In Tivoli Storage Productivity Center V5.2, the functions for adding and administering out-of-band fabric agents were moved to the web-based GUI. When you add a switch as an SNMP agent in the web-based GUI, it is displayed as

out-of-band fabric agent in the stand-alone GUI. In the stand-alone GUI, you can delete and view information about those agents.

For information about adding switches in the web-based GUI, see "Adding resources" on page 257.

Displaying information about an out-of-band agent

You can view information about an out-of-band agent including the IP address, status, and SNMP community.

To display information about an out-of-band agent, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Click **Out of Band Fabric Agents**.
2. In the content pane, click the icon to the left of the agent for which you want information. In the content pane, a notebook window opens. On the General Information page, the following information is displayed:

   Status The status of the agent.

   Host Name
   The host name of the system on which the agent is installed.

   IP Address
   The IP address of the system on which the agent is installed. This column displays Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6) addresses as appropriate.

   SNMP Community Read
   The name of the Simple Network Management Protocol (SNMP) community to which the agent belongs. The SNMP community name acts as a password that is shared by one or more SNMP hosts. The community name is used to authenticate messages being received by this SNMP host. This field is optional and might be blank if the SNMP community has not been set.

   Note: The default SNMP community is public. If the community name is not the correct community name for your environment, the out-of-band agent might not be able to properly perform scans.

   SNMP Community Write
   The community string for SNMP writes. Tivoli Storage Productivity Center uses the write community string to refresh information for the Out of Band Fabric Agent. The default is private. This value is used for Cisco switches only.

Deleting an out-of-band agent

You can remove an out-of-band agent from the navigation tree and the database. Data discovered by the agent is not removed from the database repository.

To remove an out-of-band agent, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Left-click **Out of Band Fabric**.
2. In the content pane, select an agent and click **Delete**.

Tivoli Storage Productivity Center servers

Use the **Administrative Services** > **Data Sources** > **IBM Tivoli Storage Productivity Center Servers** > function to manage the relationships between a *master* Tivoli Storage Productivity Center server and its *subordinate* Tivoli Storage Productivity Center servers.

By creating master and subordinate relationships between Tivoli Storage Productivity Center servers, you can then use a single interface to generate reports based on data and metrics collected by multiple servers in a storage environment:

- A **master** server is a server that performs normal monitoring and reporting of storage resources like a standard server, but also gathers the storage information (using Tivoli Storage Productivity Center server probes) that was collected by subordinate servers.
- A **subordinate** server is a server that monitors and reports on storage resources like a standard server, but also communicates with the master server during Tivoli Storage Productivity Center server probes. During these probes, the master server collects the storage information gathered by a subordinate server's agents and stores that information in its own database repository.

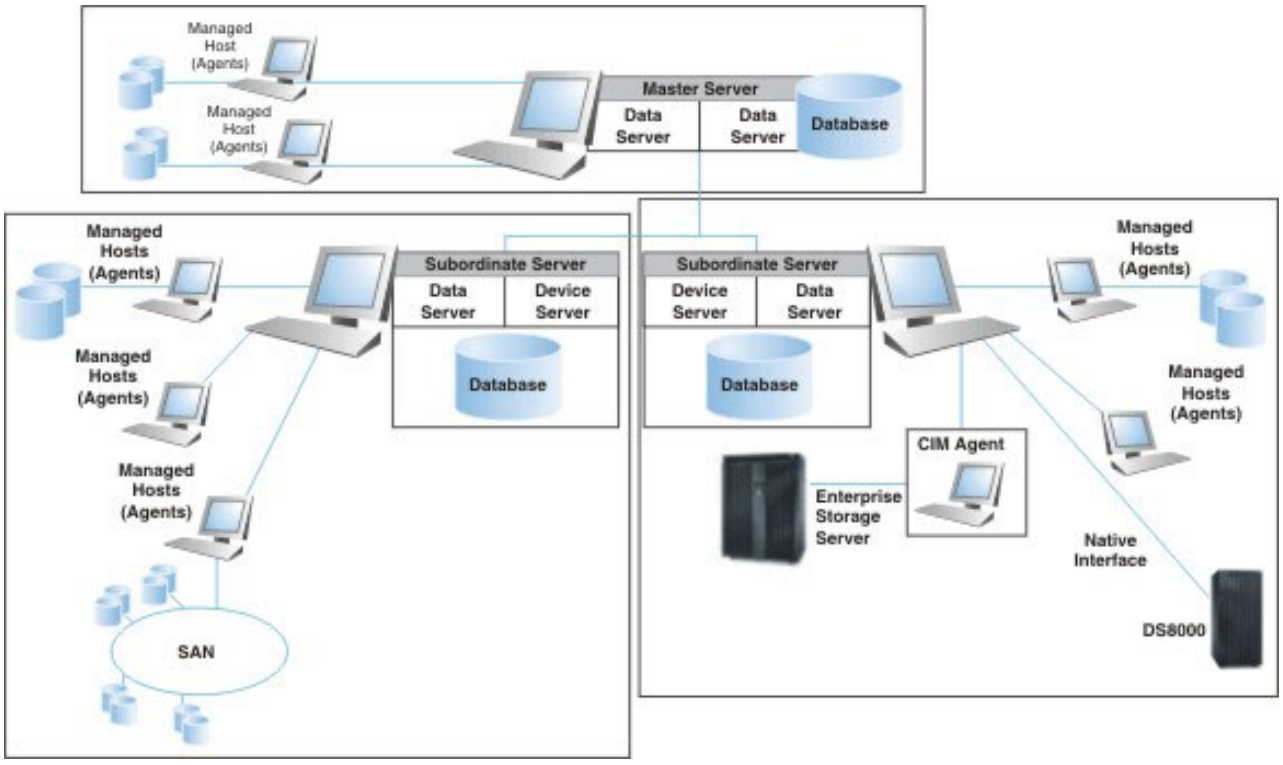

Figure 13. Master and subordinate server architecture

The rollup reports that reflect the storage information collected by the master server from subordinate servers are available in the **Tivoli Storage Productivity Center** > **Rollup Reports** node. This node is available in the navigation tree for the master server.

If the master server is on an IPv6–only server, it can communicate with existing subordinate servers under the following conditions:

- The subordinate servers are upgraded to Tivoli Storage Productivity Center V5.1 or higher and

- The IPv6 protocol is enabled on the machines where they are located.

Before you can configure and manage subordinate servers, keep in mind the following items:

- The master server must be up and running.
- You must be log on to the user interface as an Tivoli Storage Productivity Center administrator or superuser
- We recommend that the master server monitor no more than 500 unique data sources. This number includes subordinate servers, Storage Resource agents, CIMOM agents, and VM servers (VMWare).
- We recommend that each subordinate server monitors no more than 1200 unique data sources. This number includes Storage Resource agents, CIMOM agents, and VM servers (VMWare). After this threshold is met for a server, a new server is deployed and all new agents point to it.
- After upgrading to Tivoli Storage Productivity Center 5.1, you must run probes against the monitored storage assets for the master server to have information about those assets displayed in **Tivoli Storage Productivity Center** > **Reporting** > **Rollup Reports**.
- After you remove a data source (agent) from a subordinate server or a storage entity that it monitors, the master server continues to display information about that agent or storage resource in rollup reports for a specified amount of time. You can specify how quickly that information is removed by accessing the Administrative Services > Configuration > Removed Resource Retention panel and changing the settings for **TPC servers**. Additionally, if a master server is unable to contact a subordinate server for the number of days defined in the **TPC Servers** setting on the Removed Resource Retention panel, that server is removed from the subordinate server as well as information about the storage entities it was monitoring.
- A host might appear more than once in a rollup report and its data is counted multiple times in report totals. Multiple entries for a host might occur if the following conditions are true:
  - An agent is removed from a host, the GUID (globally unique identifier) for the host is changed, and then an agent is redeployed on the host.
  - The amount of time that information about a host is retained is longer than the time it took to remove the original agent and change the GUID for the host.

  For example, if the number of days to retain information about a host is 14 days, and then you 1) remove the agent on the host, 2) change the GUID of the host, and 3) redeploy an agent on that host all in less than 14 days, duplicate entries for the host appears in rollup reports. After 14 days have passed, the entry for the host before the GUID change is removed from the database and only one entry appears for the host in rollup reports.

  You can specify how quickly information is removed by changing the settings for **TPC servers** on the Administrative Services > Configuration > Removed Resource Retention panel.

Adding a Tivoli Storage Productivity Center server as a subordinate server

Learn how to add an Tivoli Storage Productivity Center server as a subordinate sever.

Keep in mind the following when adding a Tivoli Storage Productivity Center server as a subordinate server:

- How subordinate servers are configured to monitor storage entities within your environment can determine if fragmentation occurs within rollup reports. See Rollup reports - Fragmentation for more information.
- Tivoli Storage Productivity Center servers must be at version 3.3 or greater to be added as a subordinate server.
- We recommend that the master server monitor no more than 500 unique data sources. This number includes subordinate servers, Storage Resource agents, CIMOM agents, and VM servers (VMWare).
- We recommend that each subordinate server monitors no more than 1200 unique data sources. This number includes Storage Resource agents, CIMOM agents, and VM servers (VMWare). After this threshold is met for a server, a new server is deployed and all new agents point to it.
- If the master server is located on an IPv6–only server, it can communicate with existing subordinate servers under the following conditions:
  – The subordinate servers are upgraded to Tivoli Storage Productivity Center V4.1 or higher and
  – The IPv6 protocol is enabled on the machines where they are located.

To add a server as a subordinate server, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services** > **Data Sources** > **Tivoli Storage Productivity Center Servers**.
2. In the content pane, click **Add Tivoli Storage Productivity Center Server**.
3. In the Add Tivoli Storage Productivity Center Server window, enter the following information:

   Host Name
   Enter the host name of the machine on which the subordinate server is installed.

   Host Device server Port
   Enter the port with which the subordinate server's Device server is listening for requests. The default port for Tivoli Storage Productivity Center servers is 9550.

   Host Authentication Password
   Enter the password for the subordinate Tivoli Storage Productivity Center server. This host authentication password is entered when installing Tivoli Storage Productivity Center. Contact your system administrator or the person who installed Tivoli Storage Productivity Center to determine this password.

   Display Name (optional)
   Enter a name for the subordinate server.

   Description (optional)
   Enter a description for the subordinate server.

   Test Tivoli Storage Productivity Center Server Connectivity before adding
   Select this check box to have the master server connect and log into the subordinate server using the information provided on this window. This test is performed when you click **Save**.
4. Click **Save** to add the server as a subordinate server.
   - If you selected **Test Tivoli Storage Productivity Center Server Connectivity before adding**, IBM Tivoli Storage Productivity Center attempts to connect to the subordinate server using the information provided on this window.

- If you did not select **Test Tivoli Storage Productivity Center Server Connectivity before adding**, the subordinate server is automatically added as a subordinate server but has a connection status of **Unknown** on the **Administrative Services** > **Data Sources** > **Tivoli Storage Productivity Center Servers** window.

Note: If the login information for a subordinate server is incorrect, an error occurs and an associated error message dialog box is displayed.

Viewing subordinate server details

Learn how to view information about the subordinate servers that have been associated with a master server.

After you add a subordinate server to a master server, basic information about that server appears on the **Administrative Services** > **Data Sources** > **TPC Servers** window. To view additional information about a subordinate server on the TPC Servers window, follow this procedure:

1. In the navigation tree, expand **Administrative Services** > **Data Sources** > **TPC Servers**.
2. In the content pane, double-click a server row or click the icon next to a server for more detailed information about that server. Information about a subordinate server appears on the **Administrative Services** > **Data Sources** > **TPC Servers** > **TPC Server Details** window.

Testing the connection to subordinate servers

Learn how to test the connection from a master server to its associated subordinate servers.

After you add a subordinate server to a master server, basic information about that server appears on the **Administrative Services** > **Data Sources** > **TPC Servers** window. To verify that the master server can use the information provided to connect to subordinate server, follow this procedure:

1. In the navigation tree, expand **Administrative Services** > **Data Sources** > **TPC Servers**.
2. In the content pane, highlight the row for a subordinate server and click **Test TPC Server Connection** to test its connection to the master server. A message window will appear with the results of the tested connection. The information that appears in the **Connection Status** field will be updated to reflect its status.

Removing subordinate servers

Learn how to remove a subordinate server from its association with a master server.

Use the **Administrative Services** > **Data Sources** > **TPC Servers** window to remove a subordinate server from a master server. Removing a subordinate server removes subordinate server information from:

- The master server database repository
- All rollup reports
- The **TPC Servers** list on the **Administrative Services** > **Data Sources** > **TPC Servers** window

To remove a subordinate server, follow this procedure:

1. In the navigation tree, expand **Administrative Services** > **Data Sources** > **TPC Servers**.

2. In the content pane, highlight the row of a subordinate server and click **Remove TPC Server**. A confirmation window appears.
3. Click **Yes** to confirm the removal of the subordinate server.

Modifying subordinate server information

Learn how to modify information about the subordinate servers that have been associated with a master server.

After you add a subordinate server to a master server, basic information about that server appears on the **Administrative Services** > **Data Sources** > **TPC Servers** window. You can edit the information for a subordinate server by modifying its Display Name, Description, and Host Authentication Password. To view modify information about a subordinate server on the TPC Servers window, follow this procedure:

1. In the navigation tree, expand **Administrative Services** > **Data Sources** > **TPC Servers**.
2. In the content pane, double-click a server row or click the icon next to a server to edit information about that server. Information about a subordinate server appears on **Administrative Services** > **Data Sources** > **TPC Servers** > **TPC Server Details**.
3. Edit the information in the following fields:

   Host Authentication Password
   Enter the password for the subordinate IBM Tivoli Storage Productivity Center server.

   Display Name (optional)
   Enter a name for the subordinate server.

   Description (optional)
   Enter a description for the subordinate server.

   Test TPC Server Connectivity before adding
   Select this check box to have the master server connect and log into the subordinate server using the information provided on this window. This test is performed when you click **Save**.
4. Select **File** > **Save** . If successful, the TPC Servers window appears and any new values for **Display Name** and **Description** are shown for a subordinate server.

Rollup reports - fragmentation

Fragmentation occurs when related storage entities are being monitored by different subordinate IBM Tivoli Storage Productivity Center servers.

The following example shows when fragmentation can occur using subordinate servers and rollup reports:

- **Example:** Host A is assigned a volume from storage subsystem 1. Host A is being monitored only by one Tivoli Storage Productivity Center subordinate server (Server 1) and storage subsystem 1 is being monitored by another subordinate server (Server 2). Server 1 and Server 2 are subordinate servers to the master server (Server 3).

  Reason for fragmentation: In this example, the host and subsystem are fragmented because they are not being monitored by the same Tivoli Storage Productivity Center server. This will cause the **Correlated Volume Space** value in the rollup reports for storage subsystems and volumes to be incorrect.

To prevent fragmentation when using master and subordinate server relationships to generate rollup reports, we recommend the following:

- Do not monitor a storage subsystem and the hosts to which its volumes are assigned with different subordinate servers that are associated with the same master server.
- Do not monitor a fabric and its related switches with different subordinate servers that are associated with the same master server.
- If an agent machine is using storage from a subsystem, then both the agent and subsystem should be monitored by the same subordinate server.
- Configure all the subordinate servers to discover all the storage subsystems within your environment. Note that if a subordinate server is monitoring a subsystem with an agent other than the agent that is using storage on that subsystem, the value for the **Correlated Volume Space** column in rollup reports will be incorrect. For example:
  - Host A has storage from Subsystem 1
  - Subordinate server B is Monitoring Host A and Subsystem 1
  - Subordinate server C is Monitoring Subsystem 1 and probed the subsystem more recently than server B.
  - The master server is monitoring subordinate servers B and server C.

  The master server will report the storage subsystem data from server C. Because the master server does not know about Host A, the value for the **Correlated Volume Space** column will be incorrect.

Administering DB2

Administer DB2 by backing up the IBM Tivoli Storage Productivity Center database, starting the IBM Data Studio full client, and starting and stopping DB2.

Backing up the Tivoli Storage Productivity Center database

IBM Tivoli Storage Productivity Center uses circular logs for the DB2 database. Therefore, only an offline database backup is supported. When an offline database backup is run, the database connections are terminated. Therefore, the Tivoli Storage Productivity Center servers must be stopped and then restarted.

1. Close the Tivoli Storage Productivity Center GUI.
2. Stop the Tivoli Storage Productivity Center servers.
3. Log in to the Data server with a user that has DB2 administration privileges.
4. Initialize the DB2 environment with the following command:

   Windows operating systems:
   Open a command window.

   Run the **db2cmd** command.

   Linux and AIX operating systems:
   Run the **db2profile** command
5. In a DB2 command window, run the following commands to force all users or applications off DB2:

   ```
   db2 force application all
   db2 terminate
   ```
6. In a command window, run the following command to ensure that no applications are accessing DB2:

   ```
   db2 list applications
   ```

7. Create a directory to back up the Tivoli Storage Productivity Center database.
8. Run the backup command:

   Windows operating systems:

   ```
   DB2 BACKUP DATABASE TPCDB USER user_name
   USING password TO C:\DB_Backup COMPRESS
   ```

   where *user_name* is the user who owns the DB2 instance where the Tivoli Storage Productivity Center database is located and *password* is the password that is associated with that user name.

   Linux and AIX operating systems:

   ```
   DB2 BACKUP DATABASE TPCDB USER user_name
   USING password TO /tmp/DB_Backup COMPRESS
   ```

   where *user_name* is the user who owns the DB2 instance where the Tivoli Storage Productivity Center database is located and *password* is the password that is associated with that user name.
9. Restart the Tivoli Storage Productivity Center servers.

If you want to do an online backup of the database, you must change the use of circular logs to linear logs. The use of linear logs can fill up the file system quickly if the database is not maintained properly. For information about how to change the use of circular logs to linear logs, see the DB2 Information Center. Click **Designing > Backups**.

Using the command line on UNIX and Linux

This topic describes how to use a command line to perform actions against a DB2 instance under UNIX or Linux.

If the DB2 Control Center is unavailable or you do not have access to a graphical user interface, you can use a command line to execute DB2 commands such as starting and stopping an instance.

Important: If you are using DB2 10.1, you must use the command-line interface.

To use a command line to perform actions against an instance of DB2, complete the following steps:

1. Log in with a user ID or name that has ROOT, SYSADM, SYSCTRL, or SYSMAINT authority on the instance; or log in as the instance owner.
2. Run the startup script:
   - For Bourne or Korn shell, type: **. HOME/sqllib/db2profile**.
   - For C shell, type: **source HOME/sqllib/db2cshrc**.

   where `HOME` is the home directory of the instance you want to use.
3. To start the instance using the command line, type **db2start**.

   Note: When you run commands to start the database manager instance, the DB2 database manager applies the command to the current instance.
4. To stop the instance using the command line, type **db2stop**.

   Note: When you run commands to stop the database manager instance, the DB2 database manager applies the command to the current instance.

Manually starting DB2 on Windows

Start DB2 on Windows operating systems.

To start DB2 manually, complete the following steps:

1. Start the following Windows services:
   - DB2 - DB2-0
   - DB2DAS - DB2DAS00
   - DB2 JDBC Applet Server
   - DB2 License Server
   - DB2 Security Server
2. Open a DB2 command window.
3. From the DB2 Command window, run the **db2start** command.

Manually stopping DB2 on Windows

Stop DB2 manually on Windows.

1. Stop the following Windows services:
   - DB2 Security Server
   - DB2 License Server
   - DB2 JDBC Applet Server
   - DB2DAS - DB2DAS00
   - DB2 - DB2-0

   Note: When you stop the **DB2 Security Server** service, you are prompted to stop the Warehouse logger and Warehouse Serve. Click **Yes**.
2. To open a DB2 command window, click **Start** > **Programs** > **IBM DB2** > **Command Line Tools** > **Command Window**.
3. From the DB2 Command window, issue the **db2stop** command.

To restart DB2, issue the **db2start** command from the DB2 Command window. Before you can issue the command, you must first start these services on the Windows Services panel:

```
DB2 - DB2-0
DB2DAS - DB2DAS00
DB2 JDBC Applet Server
DB2 License Server
DB2 Security Server
```

Starting the IBM Data Studio full client

You can start the Data Studio full client on your workstation from either a menu option or the command line.

Opening IBM Data Studio Administration client on Windows operating systems

1. Choose one of these options:

Option	Description
Windows Server 2012	1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**. 2. On the Start page, right-click, and then click the **All apps** taskbar button.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click **Start** > **All Programs**.

2. Click **IBM Data Studio** > **Data Studio Administration Client**.

Open IBM Data Studio Administration client on Linux and AIX operating systems

On the command line, issue the following command:

```
DS_install_dir/eclipse
```

where `DS_install_dir` is the directory where you installed the full client.

/opt/IBM/DS3.1.1/eclipse

Monitoring DB2

The minimum user authority level needed for monitoring DB2 instances is a user with DB2 system maintenance authority (SYSMAINT).

To check and set SYSMAINT authority, follow these steps:

1. Run this command in the DB2 command prompt window to check to see if there is an operating system user group defined to have SYSMAINT authority:

   ```
   db2 get dbm cfg
   ```

   In the output file, look for this information:

   ```
   SYSADM group name                       (SYSADM_GROUP) =
   SYSCTRL group name                     (SYSCTRL_GROUP) =
   SYSMAINT group name                   (SYSMAINT_GROUP) =
   SYSMON group name                       (SYSMON_GROUP) =
   ```

   If the setup for the operating system group has not been done, you do not see a value set.

   If the setup has been done, this example shows what you can expect to see:

   ```
   SYSADM group name                       (SYSADM_GROUP) =
   SYSCTRL group name                     (SYSCTRL_GROUP) =
   SYSMAINT group name                   (SYSMAINT_GROUP) = ADMINISTRATORS
   SYSMON group name                       (SYSMON_GROUP) =
   ```

   In this example, the "ADMINISTRATORS" group has SYSMAINT_GROUP authority.
2. If the setup has been done, add the user you want to use to the ADMINISTRATORS group using the operating system utilities or use a user that already belongs to the ADMINISTRATORS group.

If you want to give a user group "SYSMAINT_GROUP" authority, follow these steps:

a. If a user (for example `userA`) belongs to an operating system group called `db2monitor`, here is an example of setting the db2monitor group with SYSMAINT authority. From the DB2 command prompt window, run the following command:

   ```
   db2 update dbm cfg using SYSMAINT_GROUP  db2monitor
   ```

b. After issuing the **db2 update** command, restart DB2 by running the following command from the DB2 command prompt window or restarting the system:

   ```
   db2 force application all
   ```

   This command might need to be issued a few times to stop all the database connections.

c. Run the following commands from the DB2 command prompt window:

   ```
   db2stop
   db2start
   ```

d. `UserA` can now monitor the DB2 database.

Chapter 5. Managing resources

Use Tivoli Storage Productivity Center to monitor and manage the resources and storage infrastructure in an enterprise environment.

Tivoli Storage Productivity Center provides a web-based GUI and stand-alone GUI for monitoring the condition, capacity, performance, and relationships of storage resources. Each GUI provides a different set of storage-management functions. To view a comparison of the functions that are available in each GUI, see "Available functions in the interfaces" on page 34.

Tivoli Storage Productivity Center also includes Tivoli Storage Productivity Center for Replication. Tivoli Storage Productivity Center for Replication manages copy services in storage environments. *Copy services* are features that are used by storage systems such as System Storage DS8000 to configure, manage, and monitor data-copy functions. Copy services include IBMFlashCopy, Metro Mirror, Global Mirror, and Metro Global Mirror.

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to manage a vSphere virtual storage environment. You can complete provisioning tasks and view information about storage resources in the vSphere Web Client extension.

Related concepts:

"Starting Tivoli Storage Productivity Center" on page 41
Tivoli Storage Productivity Center provides two graphical user interface (GUI) applications for managing and monitoring the resources in a storage environment: a web-based GUI and a stand-alone GUI. You can start these GUIs on the Tivoli Storage Productivity Center server or on a remote system.

Related reference:

"Available functions in the interfaces" on page 34
Tivoli Storage Productivity Center provides the web-based GUI and stand-alone GUI for managing storage resources. In V5.2, many of the functions for managing storage were moved to the web-based GUI. Determine which GUI to use based on the tasks that you must complete.

Adding resources

You must add resources to Tivoli Storage Productivity Center before you can collect data, generate reports, and manage storage that is related to those resources.

The type of resource that you want to add determines which Tivoli Storage Productivity Center GUI to use:

Web-based GUI

The web-based GUI guides you through the steps for discovering the following resources in your environment, adding the resources as data sources, and scheduling data collection:

- Storage systems
- Servers (without Storage Resource agents)
- Hypervisors
- Fabrics

© Copyright IBM Corp. 2001, 2013

- Switches

Stand-alone GUI

You can add the following resources in the stand-alone GUI:

- IBM SONAS systems
- Servers (with Storage Resource agents)
- NAS filers
- Tivoli Storage Productivity Center servers as subordinate servers

Tips:

- For a complete list of the resources and their versions that you can add, see

  Tivoli Storage Productivity Center Interoperability Matrix.
- You cannot add resources that are already being monitored by Tivoli Storage Productivity Center. A resource is considered to be monitored if it is included in a data collection job such as a probe, performance monitor, or scan.
- Some types of storage systems and switches are managed by CIM agents. Tivoli Storage Productivity Center communicates with the CIM agent to discover those resources. During the discovery process, you can click **Close** to exit the wizard. When you click **Close**, the discovery process continues in the background, but any discovered resources are not configured for data collection. Before you can manage and view information about the discovered resources, you must schedule data collection.

Related concepts:

"Alerting" on page 341
Determining when and how you are alerted to conditions and violations of performance thresholds within your storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines the data about your resources for the conditions and performance thresholds that you specify.

Related reference:

"Triggering conditions for alerts" on page 1665
You can define alerts so that IBM Tivoli Storage Productivity Center notifies you when a specified event occurs. Such events are the *triggering conditions* for the alert. The specific triggering conditions that you can use to define an alert depend on the sort of storage resource that you are monitoring.

"Performance thresholds" on page 1690
Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

"Adding servers with Storage Resource agents" on page 210
Use the stand-alone GUI to add Storage Resource agents to your system. You can add Storage Resource agents by scheduling a job or by adding an agent immediately.

"Adding a Tivoli Storage Productivity Center server as a subordinate server" on page 247
Learn how to add an Tivoli Storage Productivity Center server as a subordinate sever.

Resources that you can monitor in the web-based GUI

You can use the web-based GUI to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources

are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

Fabrics

A *fabric* is a network of hubs, switches, adapter endpoints, and connecting cables that support a communication protocol between devices. The following table shows the internal resources and related resources that you can monitor for fabrics.

Table 45. Fabric resources that you can monitor in the web-based GUI

Top-level resource icon	Internal resources	Related resources
	Blades Ports Switches Zone sets	Hypervisors Servers Storage systems Discovered ports

Hypervisors

A *hypervisor* is a layer of software or a physical device that manages multiple, virtual instances of operating systems on the same hardware. Hypervisors that you can monitor in the web-based GUI include VMWare vSphere data sources such as ESX, ESXi, and vCenter servers. If a hypervisor is part of a cluster, information about that cluster is also available.

The following table shows the internal resources and related resources that you can monitor for hypervisors.

Table 46. Hypervisor resources that you can monitor in the web-based GUI

Top-level resource icon	Internal resources	Related resources
	Virtual machines Disks Controllers Data stores	Storage systems Fabrics Switches Cluster

Servers

A *server* is a computer or host that is connected to a network and provides an access point to that network. If a server is part of a cluster, information about that cluster is also available. The following table shows the internal resources and related resources that you can monitor for servers.

Table 47. Server resources that you can monitor in the web-based GUI

Top-level resource icon	Internal resources	Related resources
	Disks Controllers Volume groups File systems and logical volumes Hypervisors Shares and exports Monitored directories	Storage systems Fabrics Switches

Storage systems

A *storage system* is a device, such as a Redundant Array of Independent Disks (RAID) controller that creates and manages other storage devices. A storage virtualizer is a storage system that virtualizes storage space from internal storage or from another storage system.

Storage systems can be configured for block storage, file storage, or a combination of block and file storage. The method that a storage system uses for managing data determines the internal resources that are monitored.

The following table shows the internal resources and related resources that you can monitor for storage systems.

Table 48. Storage system resources that you can monitor in the web-based GUI

Top-level resource icon	Internal resources	Related resources
	Block storage resources Controllers Disks External disks Host connections I/O groups Managed disks Modules Nodes Pools Ports RAID arrays Volumes **File storage resources** Controllers Clusters Disks Shares File sets File systems Network shared disks Nodes Pools	Back-end storage systems Fabrics Hypervisors Servers Switches Virtualizer storage systems

Switches

A *switch* is a network infrastructure component to which multiple nodes attach. In Tivoli Storage Productivity Center, you can monitor Fibre Channel (FC) switches. The following table shows the internal resources and related resources that you can monitor for switches.

Table 49. Switch resources that you can monitor in the web-based GUI

Top-level resource icon	Internal resources	Related resources
	Blades Fabrics Logical switches Parent switches Ports	Fabrics Hypervisors Servers Storage systems Switches Discovered ports

Related tasks:

"Adding hypervisors" on page 229
Add hypervisors for monitoring. Hypervisors include VMware vSphere data sources such as ESX, ESXi, and vCenter servers.

"Adding storage systems" on page 232
Add storage systems for monitoring.

"Adding fabrics and switches" on page 241
Add fabrics and switches for monitoring.

Adding servers

Add servers for monitoring. You can add a server by deploying a Storage Resource agent in the stand-alone GUI, or you can add an agentless server in the web-based GUI.

In Tivoli Storage Productivity Center, you can view the connectivity of all your top-level resources in a SAN environment. These resources include servers, hypervisors, fabrics, switches, and storage systems. In previous releases, to get a complete view of the connectivity between servers and other top-level resources, you were required to deploy Storage Resource agents to all of those servers. However, in some environments it was not always possible to deploy Storage Resource agents, and thus a complete view of all resource relationships was not available. In V5.2 and later, you can add servers without deploying Storage Resource agents and still view the overall connectivity, capacity, and performance of those servers.

To determine which type of server to add, follow these guidelines:

Table 50.

Condition	Add an agentless server	Deploy a Storage Resource agent to a server
You want to view the monitored resources to which a server is connected.	X	X
You want to view the amount of storage that the server is assigned from other monitored resources.	X	X

Table 50. (continued)

Condition	Add an agentless server	Deploy a Storage Resource agent to a server
You want to view the performance of the SAN-related resources that are connected to the server.	X	X
The restrictions within your organization prohibit the deployment of agents to a server.	X	
You want to view information about the resources that are internal to a server. Internal resources include controllers, disks, file systems, logical volumes, shares, and directories.		X
You want to collect information about the fabrics that are visible to the server.		X

Example: Server-centric monitoring of SAN resources without an agent

In Tivoli Storage Productivity Center, you can view the connectivity of all your top-level resources in a SAN environment. These resources include servers, hypervisors, fabrics, switches, and storage systems. In previous releases, to get a complete view of the connectivity between servers and other top-level resources, you were required to deploy Storage Resource agents on all of those servers. However, in some environments it was not always possible to deploy Storage Resource agents, and thus a complete view of all resource relationships was not available. In Tivoli Storage Productivity Center V5.2 and later, you can add servers without deploying Storage Resource agents and still view the overall connectivity and performance of SAN storage that is assigned to those servers.

When you add an agentless server for monitoring and specify its world wide port names, Tivoli Storage Productivity Center automatically correlates that server with the ports on known host connections. If matches are found between the server and host connections on monitored storage systems, the following information can be determined:

- The amount of SAN storage that is assigned to the server.
- The storage systems that are mapping volumes to the server.
- The relationships between the server and other monitored resources.
- If performance monitors collect data about the related storage systems, you can view the performance of the internal resources that are directly associated with storage that is assigned to the server. For example, if a SAN Volume Controller maps two volumes to the server, you can view the performance of those volumes and the related managed disks.

Adding servers without Storage Resource agents

Use the web-based GUI to add agentless servers for monitoring. An agentless server is a server or virtual machine that is monitored by Tivoli Storage Productivity Center but does not have a Storage Resource agent deployed to it.

When you add an agentless server, enter its IP address or host name and other optional information such as its OS type, location, and ports. Tivoli Storage Productivity Center uses this information to correlate the agentless server with the storage systems, fabrics, and switches that are being monitored. For example, the ports that you enter for an agentless server are compared to the ports on host connections for monitored storage systems. If there is a match between the server and a host connection, you can view the connectivity between the server and storage system and the capacity of volumes that are assigned to the server.

Use any of the following methods to add agentless servers for monitoring:

- Add a server by manually entering its information.
- Add a server from a host connection on a monitored storage system.
- Add a server from a discovered port that is connected to a monitored switch.
- Add a server from a discovered port that is connected to a monitored fabric.

Add an agentless server by manually entering its information.

1. In the navigation pane, select **Server Resources** > **Servers**.
2. Click **Add Server**.
3. Select a method for adding a server.
4. Complete the wizard.

Add a server from a host connection on a monitored storage system.

Use this method when you want to configure a host connection to represent a physical server. When you select a host connection, you are specifying the internal resources (volumes, ports) that belong to that server from the perspective of the related storage system.

To add an agentless server that is based on a known host connection, complete the following steps:

1. In the navigation pane, select **Storage Resources** > **Storage Systems**.
2. Right-click a storage system and select **View Details**.
3. Click **Host Connections** to view a list of host connections that are mapped to a storage system.
4. Right-click a host connection and select **Add Server**. This action is not available under the following conditions:
   - A Storage Resource agent or agentless server is already associated with the host connection.
   - The host type of the host connection is configured for SAN Volume Controller.
5. Enter information about the server and assign ports.
6. Complete the wizard.

Add a server from a discovered port that is connected to a monitored switch.

Discovered ports are ports on resources that are not monitored by Tivoli Storage Productivity Center but are visible to a monitored switch. To add an agentless server that is based on a discovered port, complete the following steps:

1. In the navigation pane, select **Network Resources** > **Switches**.
2. Right-click a switch and select **View Details**.
3. Click **Discovered Ports** to view a list of unmanaged ports that are visible to the switch.
4. Right-click an unknown port and select **Add Server**.
5. Enter information about the server and assign ports.

6. Complete the wizard.

   After the wizard is complete, the port is removed from the list of discovered ports and added as a server to the list of related resources for the switch.

Add a server from a discovered port that is connected to a monitored fabric.
Discovered ports are ports on resources that are not monitored by Tivoli Storage Productivity Center but are connected to ports on a monitored fabric. To add an agentless server that is based on a discovered port, complete the following steps:

1. In the navigation pane, select **Network Resources** > **Fabrics**.
2. Right-click a switch and select **View Details**.
3. Click **Discovered Ports** to view a list of unknown ports that are connected to ports in the fabric.
4. Right-click an unknown port and select **Add Server**.
5. Enter information about the server or virtual machine and assign ports.
6. Complete the wizard.

   After the wizard is complete, the port is removed from the list of discovered ports and added as a server to the list of related resources for the fabric.

Help tips in the GUI: To view descriptions of the information that you must enter for a server, hover the mouse pointer over the related help icons in the wizard.

Restrictions:

- The host name or IP address of an agentless server can be up to 24 characters in length and must contain valid characters. You can enter the following characters when you enter a host name or IP address:
  - A through Z (uppercase characters)
  - a through z (lowercase characters)
  - 0 through 9 (numeric characters)
  - Special characters: - (minus) . (period)
- If the location of an agentless server exceeds 64 characters, it is truncated when the server is added. To avoid the truncation, ensure that the location value for an agentless server is equal to or less than 64 characters in length.
- If you add an agentless server in the web-based GUI, you cannot use the stand-alone GUI to deploy a Storage Resource agent to that server for full monitoring. Before you can deploy a Storage Resource agent in the stand-alone GUI, you must first remove the agentless server in the web-based GUI.

  To remove an agentless server in the web-based GUI, go to the Servers page, right-click the agentless server, and select **Remove**. When the agentless server is removed, you can use the stand-alone GUI to deploy an Storage Resource agent to that server.

Methods:

Determine how to add agentless servers or virtual machines.

There are several methods that you can use to add one or more servers for monitoring. Use the following information to determine the method that best meets the needs of your environment:

Manually
Add one server by manually entering configuration information.

Import List from a File
Add one or more servers by importing configuration information from a comma-delimited file.

Discovered Servers
Add one or more Windows servers that are members of a Windows domain. If you select this method, the Add Server wizard automatically searches for servers that are in the Windows domain where Tivoli Storage Productivity Center is installed. When you select to add one server from the list of discovered servers, you can configure more information about that server. When you select to add multiple servers, the servers are automatically created without additional configuration information.

Import List:

Add one or more agentless servers by importing configuration information from a comma-delimited file. Tivoli Storage Productivity Center reads the entries in the file and automatically adds the corresponding servers.

The input file for importing server information can contain entries for one or more servers. Each line in the file represents a server that you want to add. Information about each server must be organized in the following format:

```
host name or IP address,OS type,machine type,location,WWPN
```

where:

- `host name or IP address` is required for each server entry. An IP address can be in an IPv4 or IPv6 format. You must include a host name or IP address for each server that you want to add. The host name or IP address can be up to 24 characters in length and must contain valid characters. You can enter the following characters when you enter a host name or IP address:
  - A through Z (uppercase characters)
  - a through z (lowercase characters)
  - 0 through 9 (numeric characters)
  - Special characters: - (minus) . (period)
- `OS type` is optional and represents the operating system of the server. The OS type for a server must be one of the following values:
  - Windows
  - Linux
  - AIX
  - Solaris
  - HPUX
- `Machine type` is optional and determines if the server is a physical device or a virtual machine. The machine type for a server must be one of the following values:
  - Physical
  - Virtual
- `Location` is optional and represents the physical location of the server. If the location of an agentless server exceeds 64 characters, it is truncated when the server is added. To avoid the truncation, ensure that the location value for an agentless server is equal to or less than 64 characters in length.

- `WWPN` is optional and represents the worldwide port name that is associated with the server. You can include one or more WWPNs for each server. The value for World Wide Port Name (WWPN) must be 16 hexadecimal characters in length. If the WWPN for a server entry is already associated with a server that is being monitored, that WWPN is not created for the server.

Example

```
host1,Windows,Virtual,San Jose,2100000056AE45B2,2100000056AE45B3 # Comment
host5,Linux,Physical,London,2100000056EC45A4,21000000565E4FB6
systemxyz.storage.widgets.com,Solaris,,Chicago,21000000567E5FB7
198.51.100.22,HPUX,Virtual,,21000000566E6FB8
2001:DB8:0:0:0:0:0:0,Windows,Physical,Tokyo,21000000569E7FB9
```

Tip: If you add many servers in the same input file, the adding process might require some time to complete. To run the process in the background, click **Close** on the Add Server window. If you remain on the Servers page, a message is displayed when the servers are added. If you leave the Servers page during processing, a message is not displayed. To confirm that the servers were added, you must go back to the Servers page and refresh the list.

Restrictions

Keep in mind the following restrictions when you import server information from an input file:

- Each line in the file represents one server.
- Enter a "#" at the beginning of a line to comment out a line. The server on that line is not added when the list is imported in the add server wizard.
- If there are syntax problems in the file, then none of the servers in the file are added.

Adding servers with Storage Resource agents

Use the stand-alone GUI to add Storage Resource agents to your system. You can add Storage Resource agents by scheduling a job or by adding an agent immediately.

To add a Storage Resource agent, complete the following steps:

1. In the navigation tree, click **Administrative Services** > **Configuration**. Right-click **Storage Resource Agent Deployments**. On the menu, click **Create Storage Resource Agent Deployments**.
2. The Create Storage Resource agent Deployments window opens displaying the Computers page. Enter the following information:

   Header fields

   Creator
   Predefined creator name.

   Name Predefined name.

   Description (optional)
   Enter a user-defined description name for the job.

   Enabled
   Check this box to enable Storage Resource agent deployments.

   Add Host List
   Click this button to enter names and logon credentials for the computers on which you want to deploy Storage Resource agents.

Remember: When you use an IPv6 address to add a server, the TPC Server must also have an IPv6 address. Both of these servers must be in the same domain name system (DNS).

You can enter information in one of the following ways:

- Click **Add Agents from MS Directory** to install Storage Resource agents on one or more Windows computers that are members of a Windows domain. Enter the domain controller name and login information about the displayed window. On the next window, select the domain and click **Get List of Domain Computers**. The host names of computers that you add by using this method are displayed in the **Remote Agent Machines** section of the window. In that section, you can also specify an installation location for each computer and indicate whether you are reinstalling a Storage Resource agent on a computer.
- Click this button to retrieve a list of computer names from a flat text file. The computers listed in the file are automatically added to the deployment schedule and displayed in the **Remote Agent Machines** section. In that section, you can also specify an installation location for each computer and indicate whether you are reinstalling a Storage Resource agent on a computer.

  The file that you retrieve must contain one host name or IP address per line. For example:

  ```
  tpcserver1.tpc.example.com
  9.79.179.179
  tpcserver1.tpc.example.com
  9.89.189.189
  ```

  If Tivoli Storage Productivity Center discovers any syntax problems in the file, none of the host entries in the file are added to the deployment job.
- In the **Remote Agent Machines** section, enter the names or IP addresses of the computers where you want to deploy the Storage Resource agents. Specify the installation location for the Storage Resource agent runtime files. If you do not enter an installation location, the default location is used. The default installation directories are:
  - For Windows: `C:\Program Files\IBM\TPC\agent`
  - For Linux and UNIX: `/opt/IBM/TPC/agent`
- Select **Force** under the following circumstances:
  - If an earlier Storage Resource agent installation failed and there are damaged agent files on the computer that cause further installations to fail. If you select this option, Tivoli Storage Productivity Center attempts to overwrite the previous failed deployment on the computer with a new Storage Resource agent.
  - If you want an existing Storage Resource agent to communicate with an additional Tivoli Storage Productivity Center server. To do this, you must create the deployment job from the additional Tivoli Storage Productivity Center server to which you want the Storage Resource agent to communicate.

    To install the Storage Resource agent to communicate with two different servers using the same location, follow the example in these steps:

a. From server 1 (S1), install the Storage Resource agent to the location `C:\Program Files\IBM\TPC\agent` using port 9510.
b. From server 2 (S2), install the Storage Resource agent in the same location using the same port 9510 and specify the **Force** option. This option is required when you install the Storage Resource agent in the same location.

If you want to install two Storage Resource agents on the same system, each with a different port, install the agents using different port numbers and different locations.

Note:

– You cannot change the communication type for a Storage Resource agent (daemon or non-daemon) when you select the **Force** option. Make sure to select the same communication type as the existing Storage Resource agent when you create a deployment job.
– If a Storage Resource agent exists on a target computer and you do not select **Force**, an error occurs during validation and the Storage Resource agent is not installed.

- Select **Initial Scan** to perform a file system scan on the computer where a Storage Resource agent is deployed. If this box is checked, the scan is run immediately after a Storage Resource agent is deployed.
- Click **Windows Services Information** to enter login information that Tivoli Storage Productivity Center uses to create a Windows service for the Storage Resource agent. Use the Windows Service Information window to select the local system account, enter information about an existing user account, or create a new local account. If you do not specify a Windows service account, the default itcauser account is created. This action is only available for agents that are deployed to Windows computers as daemon services. This button is disabled if Use Daemon Service for Runtime Operation is not selected.

Enter the following information:

User Enter the user ID that Tivoli Storage Productivity Center uses when logging in to the host computer to install a Storage Resource agent. The value in the field is applied to all the computers that appear in the **Remote Agent Machines** table.

Attention: This ID must have administrative or root privileges on the target computer.

Password
Enter the password for the user ID.

Certificate Location
Enter the fully qualified path of the certificate on the computer where the Data server is located. This file is used for certificate-based authentication. If you do not enter a value in this field the default location is used.

Passphrase
Enter the passphrase for the certificate file. The passphrase was created when the certificate was generated.

Port Enter the port number on which a Storage Resource agent listens for requests. The default is 9510.

Tip: This value is required if you run the Storage Resource agent as a daemon service.

Use Daemon Service for Runtime Operation
Select this option to run a Storage Resource agent on the monitored computer as a daemon service.

See Storage Resource agent Deployment Considerations for more information about deploying agents as daemon and non-daemon services.

Restriction: Do not select this option to run the Storage Resource agent as a stand-alone executable file on the monitored computer.

Validate before Save
Check this box to indicate that Tivoli Storage Productivity Center should attempt to communicate with the computers that are added to the deployment job when you click **Save** on this page. If the validation fails for a computer, the agent is not installed on that computer.

Edit Selected Entries
Click this button to edit settings for the selected Storage Resource agents. This information includes user ID, password, certificate location, passphrase, and port. You must submit the deployment job for the changes to take effect.

Remove
Click this button to remove a Storage Resource agent from the deployment job.

Tip: This button does not remove the agent from Tivoli Storage Productivity Center if it is already installed; it only removes it from the deployment job.

3. Click the **When to run** tab to enter the following information:

   How often to run
   Specify a time to run:
   - **Run now**
   - **Run once at** (specify a date and time to run)

   How to handle time zones
   Specify a time zone to use:
   - **Use the time zone that the server runs in**
   - **Use this time zone** (select a time zone)

4. Click the **Alert** tab to specify the following information:

   Triggering-Condition
   The triggering condition you can specify is: Storage Resource agent Deployment Failed

   Triggered-Actions
   Choose from the following check boxes:
   - **SNMP Trap**
   - **TEC / OMNIbus Event**
   - **Login Notification**
   - **Windows Event Log**

- **Email**

  Depending on what action you select, you might have other choices to make. For example, if you select the **Windows Event Log** check box, the **Event Type** field becomes active so that you can specify a severity for the event in the Windows event log.

5. Click **File > Save**.

For related tasks, see the *Planning > General Planning > Planning for Storage Resource agents* section in the *Tivoli Storage Productivity Center Installation and Configuration Guide* for additional configuration information about certificates used after deployment.

Check the TPCD_#####.log file in the Tivoli Storage Productivity Center installation directory to view the error messages related to failed validations. The following are the default locations for this log file:

- Windows: `C:\Program Files\IBM\TPC\data\log`
- UNIX/Linux: `/opt/IBM/TPC/data/log`

If multiple Storage Resource agents are set up to probe or scan the same storage resources, the Storage Resource agent that was added to Tivoli Storage Productivity Center first is used for the probe or scan. Therefore, only data that is gathered by the first Storage Resource agent is shown.

Adding hypervisors

Add hypervisors for monitoring. Hypervisors include VMware vSphere data sources such as ESX, ESXi, and vCenter servers.

You can add a hypervisor by specifying connection information for that hypervisor, or you can add multiple hypervisors by specifying connection information for a vCenter Server. For a complete list of hypervisors and their versions that you can add, see the *Agents, Servers, and Browsers* section in the *Tivoli Storage Productivity Center Interoperability Matrix*.

Help tips in the GUI: To view descriptions of the information that you must enter for a hypervisor, hover the mouse pointer over the related help icons in the wizard.

1. In the navigation pane of the web-based GUI, select **Server Resources > Hypervisors**.
2. Click **Add Hypervisor**.
3. Enter connection information about the hypervisor. If you enter connection information about a vCenter server, Tivoli Storage Productivity Center automatically discovers and lists the hypervisors that it manages.

   Restriction: The user name that you specify for a hypervisor must have permission to browse through the data stores on VMware. For information about how to check permissions, see "Checking permissions to browse data stores" on page 272.
4. Specify display information and schedule a probe of the hypervisor. When you add a vCenter server, specify display information and schedule a probe for the hypervisors that were discovered. If you select multiple hypervisors at the same time, you cannot specify display information.
5. Complete the wizard.

Tip: In previous versions of Tivoli Storage Productivity Center, you were required to download the certificate from the hypervisor or vCenter server and add it to the vmware.jks keystore of the Device server. In V5.2 and later, the web-based GUI completes this process automatically.

After a hypervisor is added, a probe collects status and asset information about that hypervisor. You can view detailed information about the hypervisor in the web-based GUI and stand-alone GUI.

Checking permissions to browse data stores

Determine if the user name that you specified for a VMware data source has permission to browse through the data stores on a hypervisor.

When you add a VMware data source in Tivoli Storage Productivity Center, the user name that you specify must have permission to browse through the data stores on VMware. IBM Tivoli Storage Productivity Center must browse through the data stores to collect information from the hypervisors. However, the "Read Only" role as defined by VMware does not allow IBM Tivoli Storage Productivity Center to browse the data stores. You can use the "Virtual Machine Power User" role if you do not want to use the Administrator role, or you can create a custom role with the required permissions.

To check the permissions for a VMware user, complete the following steps:

1. Connect a VI client to the VMware data source that you added for monitoring. The data source can be an ESX server or vCenter server.
2. Click the server and go to the **Summary** tab.
3. Right-click the data source and select **Browse** to confirm that the browse permission is working correctly.
4. Connect the VI client to the data source with the Administrator role, go to the **Permission** tab, and determine if the role name is used for the user.
5. Open the **Permissions** view, right-click that role, and select the **Edit** role.
6. Verify that the permissions are checked. Expand the check groups to find the specific permissions.

For more information about VMware user roles, see Managing VMware VirtualCenter Roles and Permissions.

Adding storage systems

Add storage systems for monitoring.

You can add the following types of storage systems for monitoring:

- DS4000
- DS5000
- DS6000
- DS8000
- XIV system
- ESS
- SAN Volume Controller
- Storwize V3500, Storwize V3700, Storwize V7000
- Storwize V7000 Unified
- IBM SONAS

- All others (managed by CIM agents)

Tips:

- For a complete list of the versions of storage systems and CIM agents that you can add, see Tivoli Storage Productivity Center Interoperability Matrix and go to the *Storage* section.
- The storage system that you want to add determines which Tivoli Storage Productivity Center GUI to use. Use the stand-alone GUI to add IBM SONAS storage systems for monitoring. Use the web-based GUI to add all other storage systems and CIM agents.

After a storage system is added, data collection jobs gather information about that storage system. When data collection is complete, you can view detailed information about the storage system in the web-based GUI and stand-alone GUI. Storage systems can be configured for block storage, file storage, or a combination of block and file storage. The method that a storage system uses for managing data determines the internal resources that are monitored.

Related concepts:

"Collecting data" on page 295
Determining the data that you want to gather about storage resources is critical to helping you implement a storage management strategy. Tivoli Storage Productivity Center provides different jobs for collecting data. When you determine the types of data that you want to collect about resources, schedule the corresponding data collection jobs for those resources.

"Collecting asset and status data" on page 296
Use probes to collect asset, status, and storage data about monitored resources in your environment. Probes are also used to discover information about new or removed disks and file systems. You can view information about probes on the resource list and resource details pages, for example, the Storage Systems page, or the Storage Systems details page.

Related reference:

"CIM agents" on page 205
Administer CIM agents that are associated with storage resources that are monitored by Tivoli Storage Productivity Center. Common Information Model (CIM) agents provide a CIM interface for collecting information about certain types of storage systems and switches.

DS4000 or DS5000

Add a DS4000 or DS5000 storage system by connecting to the CIM agent that manages the storage system.

Adding a CIM agent that manages multiple storage systems: If a CIM agent manages multiple storage systems, all of the storage systems that are managed by the agent are added to Tivoli Storage Productivity Center.

To add a DS4000 or DS5000 storage system, complete the following steps:

1. In the navigation pane, select **Storage Resources** > **Storage Systems**.
2. Click **Add Storage System**.
3. Click the **DS4000/DS5000** icon.
4. On the Discover page, enter the following information about a CIM agent:

   CIM agent host name or IP address
   The host name or IP address for the computer on which the CIM agent

that manages the storage system is installed. Depending on what is supported in your environment, you can enter an Internet Protocol version 4 (IPv4) or IPv6 address. If you enter an IPv6 address, the preferred representation is written as eight groups of four hexadecimal digits. Example: 2001:DB95:0000:1234:0000:0000:5678:ABCD.

User name, Password
The user name and password that are used to authenticate to the CIM agent.

CIM agent display name
The name of the CIM agent, as specified by the agent provider. This name is displayed in the Tivoli Storage Productivity Center interface.

CIM agent description
A description of the CIM agent. This description can be any combination of alphanumeric characters.

Advanced
Expand **Advanced** to set the optional protocol, port, and namespace for the CIM agent.

Protocol
The version of the cim-xml protocol, http or https.

Port The port on which the CIM agent is listening. By default this port is 5989 for a secure connection and 5988 for an unsecured connection. See the documentation for the appropriate SMI-S CIM Provider for more information about available ports.

Namespace
This namespace within the CIM agent allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how IBM Tivoli Storage Productivity Center interacts with the CIM agent when retrieving information. For information about the namespace to use, see http://www-01.ibm.com/support/docview.wss?uid=swg21386446. Click the appropriate version for **Storage**.

To view a list of default namespaces for CIM agents for storage systems that is included with Tivoli Storage Productivity Center, open the following file in a text editor: *install_directory*/data/config/namespace.config, where *install_directory* represents the directory where the product is installed.

Check the documentation of the appropriate SMI-S CIM Provider or contact the storage system vendor to ensure that you use the most current namespaces.

5. On the Configure page, enter the following information. If multiple storage systems are managed by the CIM agent, select the storage system or systems that you want to configure. To select multiple storage systems, press Ctrl and click or Shift and click. The default configuration settings are applied to the remaining storage systems.

Display name
The name of the storage system that is displayed in the Tivoli Storage Productivity Center interface. If you do not enter a value, the default name is provided. This is the name that is provided with the storage system.

This field is inactive if you selected multiple storage systems.

Location
The physical location of the storage system. You can accept the default, **No Location**; enter a new location; or select a location that is already defined.

Data Collection
Schedule the probe and performance monitor for the storage system. The probe collects status, asset, and storage information about the storage system. The performance monitor collects metrics that measure the performance of the storage system.

Probe Enter the time and schedule for the storage system probe to run.

Enter the time in *hh:mm* format, where *hh* equals the hour and *mm* equals the minute. The time zone that is shown is the time zone where the Tivoli Storage Productivity Center Data server is located.

Performance monitor
Enable or disable the performance monitor for the storage system. If you enable the performance monitor, select the schedule for how often you want performance data to be collected.

DS6000

Add a DS6000 storage system by connecting to the CIM agent that manages the storage system.

Adding a CIM agent that manages multiple storage systems: If a CIM agent manages multiple storage systems, all of the storage systems that are managed by the agent are added to Tivoli Storage Productivity Center.

To add a DS6000 storage system, complete the following steps:

1. In the navigation pane, select **Storage Resources** > **Storage Systems**.
2. Click **Add Storage System**.
3. Click the **DS6000** icon.
4. On the Discover page, enter the following information about a CIM agent:

CIM agent host name or IP address
The host name or IP address for the computer on which the CIM agent that manages the storage system is installed. Depending on what is supported in your environment, you can enter an Internet Protocol version 4 (IPv4) or IPv6 address. If you enter an IPv6 address, the preferred representation is written as eight groups of four hexadecimal digits. Example: 2001:DB95:0000:1234:0000:0000:5678:ABCD.

User name, Password
The user name and password that are used to authenticate to the CIM agent.

CIM agent display name
The name of the CIM agent, as specified by the agent provider. This name is displayed in the Tivoli Storage Productivity Center interface.

CIM agent description
A description of the CIM agent. This description can be any combination of alphanumeric characters.

Advanced

Expand **Advanced** to set the optional protocol, port, and namespace for the CIM agent.

Protocol

The version of the cim-xml protocol, http or https.

Port The port on which the CIM agent is listening. By default this port is 5989 for a secure connection and 5988 for an unsecured connection. See the documentation for the appropriate SMI-S CIM Provider for more information about available ports.

Namespace

This namespace within the CIM agent allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how IBM Tivoli Storage Productivity Center interacts with the CIM agent when retrieving information. For information about the namespace to use, see http://www-01.ibm.com/support/docview.wss?uid=swg21386446. Click the appropriate version for **Storage**.

To view a list of default namespaces for CIM agents for storage systems that is included with Tivoli Storage Productivity Center, open the following file in a text editor: *install_directory*/data/config/namespace.config, where *install_directory* represents the directory where the product is installed.

Check the documentation of the appropriate SMI-S CIM Provider or contact the storage system vendor to ensure that you use the most current namespaces.

5. On the Configure page, enter the following information. If multiple storage systems are managed by the CIM agent, select the storage system or systems that you want to configure. To select multiple storage systems, press Ctrl and click or Shift and click. The default configuration settings are applied to the remaining storage systems.

Display name

The name of the storage system that is displayed in the Tivoli Storage Productivity Center interface. If you do not enter a value, the default name is provided. This is the name that is provided with the storage system.

This field is inactive if you selected multiple storage systems.

Location

The physical location of the storage system. You can accept the default, **No Location**; enter a new location; or select a location that is already defined.

Data Collection

Schedule the probe and performance monitor for the storage system. The probe collects status, asset, and storage information about the storage system. The performance monitor collects metrics that measure the performance of the storage system.

Probe Enter the time and schedule for the storage system probe to run.

Enter the time in *hh:mm* format, where *hh* equals the hour and *mm* equals the minute. The time zone that is shown is the time zone where the Tivoli Storage Productivity Center Data server is located.

Performance monitor
Enable or disable the performance monitor for the storage system. If you enable the performance monitor, select the schedule for how often you want performance data to be collected.

DS8000

Add a DS8000 storage system by connecting to the Hardware Management Console (HMC) that is used to manage the storage system.

Adding an HMC that manages multiple storage systems: If an HMC manages multiple storage systems, all of the storage systems that are managed by the HMC are added to Tivoli Storage Productivity Center.

To add a DS8000 storage system, complete the following steps:

1. In the navigation pane, select **Storage Resources** > **Storage Systems**.
2. Click **Add Storage System**.
3. Click the **DS8000** icon.
4. On the Discover page of the Add Storage System wizard, enter the following information:

   Primary HMC host name or IP address
   The host name or IP address for the Hardware Management Console (HMC) that is used to manage the storage system.

   Secondary HMC host name or IP address
   The host name or IP address for a second HMC that is used to manage the storage system. This value is optional.

   User name, Password
   The user name and password for logging on to the storage system. This user name is the same as the ESSNI user name.

   Display name for HMC
   The name of the HMC that is displayed in the Tivoli Storage Productivity Center interface.

   Description of HMC
   A description of the HMC. This description can be any combination of alphanumeric characters.
5. On the Configure page, enter the following information. If multiple storage systems are managed by the HMC, select the storage system or systems that you want to configure. To select multiple storage systems, press Ctrl and click or Shift and click. The default configuration settings are applied to the remaining storage systems.

   Display name
   The name of the storage system that is displayed in the Tivoli Storage Productivity Center interface. If you do not enter a value, the default name is provided. This is the name that is provided with the storage system.

   This field is inactive if you selected multiple storage systems.

Location
The physical location of the storage system. You can accept the default, **No Location**; enter a new location; or select a location that is already defined.

Data Collection
Schedule the probe and performance monitor for the storage system. The probe collects status, asset, and storage information about the storage system. The performance monitor collects metrics that measure the performance of the storage system.

Probe Enter the time and schedule for the storage system probe to run.

Enter the time in *hh:mm* format, where *hh* equals the hour and *mm* equals the minute. The time zone that is shown is the time zone where the Tivoli Storage Productivity Center Data server is located.

Performance monitor
Enable or disable the performance monitor for the storage system. If you enable the performance monitor, select the schedule for how often you want performance data to be collected.

XIV

Add an XIV system by using a direct connection to the system.

To add an XIV system, complete the following steps:

1. In the navigation pane, select **Storage Resources** > **Storage Systems**.
2. Click **Add Storage System**.
3. Click the **XIV** icon.
4. On the Discover page of the Add Storage System wizard, enter the following information:

   Host name or IP address
   The host name or IP address for the storage system.

   User name, Password
   The user name and password for logging on to the storage system. The user name must have administrative rights on the XIV system.
5. On the Configure page, enter the following information:

   Display name
   The name of the storage system that is displayed in the Tivoli Storage Productivity Center interface. If you do not enter a value, the default name is provided. This is the name that is provided with the storage system.

   Location
   The physical location of the storage system. You can accept the default, **No Location**; enter a new location; or select a location that is already defined.

   Data Collection
   Schedule the probe and performance monitor for the storage system. The probe collects status, asset, and storage information about the storage system. The performance monitor collects metrics that measure the performance of the storage system.

Probe Enter the time and schedule for the storage system probe to run.

Enter the time in *hh:mm* format, where *hh* equals the hour and *mm* equals the minute. The time zone that is shown is the time zone where the Tivoli Storage Productivity Center Data server is located.

Performance monitor
Enable or disable the performance monitor for the storage system. If you enable the performance monitor, select the schedule for how often you want performance data to be collected.

ESS

Add an ESS storage system by connecting to the CIM agent that manages the storage system.

Adding a CIM agent that manages multiple storage systems: If a CIM agent manages multiple storage systems, all of the storage systems that are managed by the agent are added to Tivoli Storage Productivity Center.

To add an ESS storage system, complete the following steps:

1. In the navigation pane, select **Storage Resources** > **Storage Systems**.
2. Click **Add Storage System**.
3. Click the **ESS800** icon.
4. On the Discover page, enter the following information about a CIM agent:

   CIM agent host name or IP address
   The host name or IP address for the computer on which the CIM agent that manages the storage system is installed. Depending on what is supported in your environment, you can enter an Internet Protocol version 4 (IPv4) or IPv6 address. If you enter an IPv6 address, the preferred representation is written as eight groups of four hexadecimal digits. Example: 2001:DB95:0000:1234:0000:0000:5678:ABCD.

   User name, Password
   The user name and password that are used to authenticate to the CIM agent.

   CIM agent display name
   The name of the CIM agent, as specified by the agent provider. This name is displayed in the Tivoli Storage Productivity Center interface.

   CIM agent description
   A description of the CIM agent. This description can be any combination of alphanumeric characters.

   Advanced
   Expand **Advanced** to set the optional protocol, port, and namespace for the CIM agent.

   Protocol
   The version of the cim-xml protocol, http or https.

   Port The port on which the CIM agent is listening. By default this port is 5989 for a secure connection and 5988 for an unsecured connection. See the documentation for the appropriate SMI-S CIM Provider for more information about available ports.

Namespace

This namespace within the CIM agent allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how IBM Tivoli Storage Productivity Center interacts with the CIM agent when retrieving information. For information about the namespace to use, see http://www-01.ibm.com/support/docview.wss?uid=swg21386446. Click the appropriate version for **Storage**.

To view a list of default namespaces for CIM agents for storage systems that is included with Tivoli Storage Productivity Center, open the following file in a text editor: *install_directory*/data/config/namespace.config, where *install_directory* represents the directory where the product is installed.

Check the documentation of the appropriate SMI-S CIM Provider or contact the storage system vendor to ensure that you use the most current namespaces.

5. On the Configure page, enter the following information. If multiple storage systems are managed by the CIM agent, select the storage system or systems that you want to configure. To select multiple storage systems, press Ctrl and click or Shift and click. The default configuration settings are applied to the remaining storage systems.

   Display name

   The name of the storage system that is displayed in the Tivoli Storage Productivity Center interface. If you do not enter a value, the default name is provided. This is the name that is provided with the storage system.

   This field is inactive if you selected multiple storage systems.

   Location

   The physical location of the storage system. You can accept the default, **No Location**; enter a new location; or select a location that is already defined.

   Data Collection

   Schedule the probe and performance monitor for the storage system. The probe collects status, asset, and storage information about the storage system. The performance monitor collects metrics that measure the performance of the storage system.

   Probe Enter the time and schedule for the storage system probe to run.

   Enter the time in *hh:mm* format, where *hh* equals the hour and *mm* equals the minute. The time zone that is shown is the time zone where the Tivoli Storage Productivity Center Data server is located.

   Performance monitor

   Enable or disable the performance monitor for the storage system. If you enable the performance monitor, select the schedule for how often you want performance data to be collected.

SAN Volume Controller

Add a SAN Volume Controller storage system by using a direct connection to the system.

To add a SAN Volume Controller storage system, complete the following steps:

1. In the navigation pane, select **Storage Resources** > **Storage Systems**.
2. Click **Add Storage System**.
3. Click the **SVC** icon.
4. On the Discover page of the Add Storage System wizard, enter the following information:

SAN Volume Controller or Storwize V7000 family

Host name or IP address
The host name or IP address for the storage system.

Authentication
You can use a user name and password or a private Secure Shell (SSH) key to log on to the storage system. The authentication method that you select determines the options that are displayed.

User Name/Password
The user name and password for logging on to the storage system.

Secure Shell (SSH)
Use an existing SSH key or upload a new key to the storage system. Select one of the following actions:

Use an existing SSH key
Use an SSH key that was uploaded to the storage system by using a method other than through Tivoli Storage Productivity Center, such as the storage system web interface.

SSH key
The location of the SSH key.

The default location is `${device.conf}\tpc_svc.pem`, which represents the Tivoli Storage Productivity Center default key file `tpc_svc.pem`. The `tpc_svc.pem` file is in the `conf` directory where the Device server is installed.

You can enter another location or select **Browse** to search for a key file. If you select **Browse**, the following fields are displayed:

Select file
The location of the SSH key file. You can click **Browse** to search for a file.

Passphrase
The passphrase for the SSH key

pair. If you do not have a passphrase, leave this field blank.

The SSH key file is transferred from the computer where the web browser is located to the computer where the Tivoli Storage Productivity Center server is located.

Upload a new SSH key

Provide the following information to upload an SSH key to the storage system:

SSH key

The location of the SSH key. The key must exist on the system where you are running the Tivoli Storage Productivity Center user interface.

The SSH key must be in OpenSSH format or in PuTTY (`.ppk`) format that is not password protected.

Passphrase

The passphrase for the SSH key pair. If you do not have a passphrase, leave this field blank.

User name, Password

The name and password for a user that belongs to the storage system Security Administrator role for the cluster that contains the storage system.

Tivoli Storage Productivity Center uses this value to configure the SSH key for the user that is entered in the **Associate user** field. The user name that is entered in the **User name** field must have privileges to modify other user accounts, otherwise Tivoli Storage Productivity Center cannot configure the SSH key.

Associate user

The user that is associated with the SSH key.

You can click **Get Users** to retrieve all of the existing users from the storage system. You must select a user that belongs to the storage system Administrator role.

The SSH key file is transferred from the computer where the web browser is located to both the computer where the Tivoli Storage Productivity Center server is located and to the storage system.

5. On the Configure page, enter the following information:

Display name
The name of the storage system that is displayed in the Tivoli Storage Productivity Center interface. If you do not enter a value, the default name is provided. This is the name that is provided with the storage system.

Location
The physical location of the storage system. You can accept the default, **No Location**; enter a new location; or select a location that is already defined.

Data Collection
Schedule the probe and performance monitor for the storage system. The probe collects status, asset, and storage information about the storage system. The performance monitor collects metrics that measure the performance of the storage system.

Probe Enter the time and schedule for the storage system probe to run.

Enter the time in *hh:mm* format, where *hh* equals the hour and *mm* equals the minute. The time zone that is shown is the time zone where the Tivoli Storage Productivity Center Data server is located.

Performance monitor
Enable or disable the performance monitor for the storage system. If you enable the performance monitor, select the schedule for how often you want performance data to be collected.

Storwize family

Add a Storwize V3500, Storwize V3700, or Storwize V7000 storage system by using a direct connection to the system.

To add a Storwize V3500, Storwize V3700, or Storwize V7000 storage system, complete the following steps:

1. In the navigation pane, select **Storage Resources** > **Storage Systems**.
2. Click **Add Storage System**.
3. Click the **Storwize Family** icon.
4. On the Discover page of the Add Storage System wizard, enter the following information:

SAN Volume Controller or Storwize V7000 family

Host name or IP address
The host name or IP address for the storage system.

Authentication
You can use a user name and password or a private Secure Shell (SSH) key to log on to the storage system. The authentication method that you select determines the options that are displayed.

User Name/Password
The user name and password for logging on to the storage system.

Secure Shell (SSH)

Use an existing SSH key or upload a new key to the storage system. Select one of the following actions:

Use an existing SSH key

Use an SSH key that was uploaded to the storage system by using a method other than through Tivoli Storage Productivity Center, such as the storage system web interface.

SSH key

The location of the SSH key.

The default location is `${device.conf}\tpc_svc.pem`, which represents the Tivoli Storage Productivity Center default key file `tpc_svc.pem`. The `tpc_svc.pem` file is in the `conf` directory where the Device server is installed.

You can enter another location or select **Browse** to search for a key file. If you select **Browse**, the following fields are displayed:

Select file

The location of the SSH key file. You can click **Browse** to search for a file.

Passphrase

The passphrase for the SSH key pair. If you do not have a passphrase, leave this field blank.

The SSH key file is transferred from the computer where the web browser is located to the computer where the Tivoli Storage Productivity Center server is located.

Upload a new SSH key

Provide the following information to upload an SSH key to the storage system:

SSH key

The location of the SSH key. The key must exist on the system where you are running the Tivoli Storage Productivity Center user interface.

The SSH key must be in OpenSSH format or in PuTTY (`.ppk`) format that is not password protected.

Passphrase

The passphrase for the SSH key pair. If you do not have a passphrase, leave this field blank.

User name, Password
The name and password for a user that belongs to the storage system Security Administrator role for the cluster that contains the storage system.

Tivoli Storage Productivity Center uses this value to configure the SSH key for the user that is entered in the **Associate user** field. The user name that is entered in the **User name** field must have privileges to modify other user accounts, otherwise Tivoli Storage Productivity Center cannot configure the SSH key.

Associate user
The user that is associated with the SSH key.

You can click **Get Users** to retrieve all of the existing users from the storage system. You must select a user that belongs to the storage system Administrator role.

The SSH key file is transferred from the computer where the web browser is located to both the computer where the Tivoli Storage Productivity Center server is located and to the storage system.

5. On the Configure page, enter the following information:

Display name
The name of the storage system that is displayed in the Tivoli Storage Productivity Center interface. If you do not enter a value, the default name is provided. This is the name that is provided with the storage system.

Location
The physical location of the storage system. You can accept the default, **No Location**; enter a new location; or select a location that is already defined.

Data Collection
Schedule the probe and performance monitor for the storage system. The probe collects status, asset, and storage information about the storage system. The performance monitor collects metrics that measure the performance of the storage system.

Probe Enter the time and schedule for the storage system probe to run.

Enter the time in *hh:mm* format, where *hh* equals the hour and *mm* equals the minute. The time zone that is shown is the time zone where the Tivoli Storage Productivity Center Data server is located.

Performance monitor
Enable or disable the performance monitor for the storage

system. If you enable the performance monitor, select the schedule for how often you want performance data to be collected.

Storwize V7000 Unified

Add a Storwize V7000 Unified storage system by using a direct connection to the system.

To add a Storwize V7000 Unified storage system, complete the following steps:

1. In the navigation pane, select **Storage Resources** > **Storage Systems**.
2. Click **Add Storage System**.
3. Click the **Storwize V7000U** icon.
4. On the Discover page of the Add Storage System wizard, enter the following information:

   Storwize V7000 Unified

   Host name or IP address
   The host name or IP address for the storage system.

   Authentication
   You can use a user name and password or a private Secure Shell (SSH) key to log on to the storage system. The authentication method that you select determines the options that are displayed.

   User Name/Password
   The user name and password for logging on to the storage system.

   Secure Shell (SSH)
   Use an existing SSH key or upload a new key to the storage system. Select one of the following actions:

   Use an existing SSH key
   Use an SSH key that was uploaded to the storage system by using a method other than through Tivoli Storage Productivity Center, such as the storage system web interface.

   SSH key
   The location of the SSH key.

   The default location is `${device.conf}\tpc_svc.pem`, which represents the Tivoli Storage Productivity Center default key file `tpc_svc.pem`. The `tpc_svc.pem` file is in the `conf` directory where the Device server is installed.

   You can enter another location or select **Browse** to search for a key file. If you select **Browse**, the following fields are displayed:

   Select file
   The location of the SSH key file. You can click **Browse** to search for a file.

Passphrase
The passphrase for the SSH key pair. If you do not have a passphrase, leave this field blank.

The SSH key file is transferred from the computer where the web browser is located to the computer where the Tivoli Storage Productivity Center server is located.

Upload a new SSH key
Provide the following information to upload an SSH key to the storage system:

SSH key
The location of the SSH key. The key must exist on the system where you are running the Tivoli Storage Productivity Center user interface.

The SSH key must be in OpenSSH format or in PuTTY (`.ppk`) format that is not password protected.

Passphrase
The passphrase for the SSH key pair. If you do not have a passphrase, leave this field blank.

User name, Password
The name and password for a user that belongs to the storage system Security Administrator role for the cluster that contains the storage system.

Tivoli Storage Productivity Center uses this value to configure the SSH key for the user that is entered in the **Associate user** field. The user name that is entered in the **User name** field must have privileges to modify other user accounts, otherwise Tivoli Storage Productivity Center cannot configure the SSH key.

Associate user
The user that is associated with the SSH key.

You can click **Get Users** to retrieve all of the existing users from the storage system. You must select a user that belongs to the storage system Administrator role.

The SSH key file is transferred from the computer where the web browser is located to

both the computer where the Tivoli Storage Productivity Center server is located and to the storage system.

Use different authentication credentials for file storage
Storwize V7000 Unified contains block-level and file-level data. If the credentials are different for block storage and file storage, select this check box to define the credentials for file storage. The options and fields that are displayed are described previously under **Authentication**.

5. On the Configure page, enter the following information:

Display name
The name of the storage system that is displayed in the Tivoli Storage Productivity Center interface. If you do not enter a value, the default name is provided. This is the name that is provided with the storage system.

Location
The physical location of the storage system. You can accept the default, **No Location**; enter a new location; or select a location that is already defined.

Data Collection
Schedule the probe and performance monitor for the storage system. The probe collects status, asset, and storage information about the storage system. The performance monitor collects metrics that measure the performance of the storage system.

Probe Enter the time and schedule for the storage system probe to run.

Enter the time in *hh:mm* format, where *hh* equals the hour and *mm* equals the minute. The time zone that is shown is the time zone where the Tivoli Storage Productivity Center Data server is located.

Performance monitor
Enable or disable the performance monitor for the storage system. If you enable the performance monitor, select the schedule for how often you want performance data to be collected.

IBM SONAS

Use the Configure Devices wizard in the stand-alone GUI to add IBM SONAS systems for monitoring by Tivoli Storage Productivity Center.

Tip: You can configure more than one IBM SONAS system at a time by using the Configure Devices wizard.

To add IBM SONAS systems for monitoring, complete the following steps:

1. In the navigation tree of the stand-alone GUI, go to **Disk Manager** > **Storage Subsystems**.
2. Click **Add Storage Subsystem**.
3. Specify connection properties for the IBM SONAS system that you want to add. For information about the connection properties, see "Connection properties for IBM SONAS systems" on page 289.

4. Click **Add** to verify the connection information. When verified, the storage system is displayed at the bottom of the page.
5. Repeat steps 2 and 3 to add more IBM SONAS systems at the same time.
6. Click **Next** after you define the connection properties for all the storage systems that you want to add.
7. On the summary page of the wizard, review the selections that you made.

After you add an IBM SONAS system for monitoring, you can complete the following tasks:

- In the stand-alone GUI, add the storage system to a monitoring group. Use monitoring groups to manage a collection of storage systems in the same manner. Tivoli Storage Productivity Center provides predefined monitoring groups that are associated with an existing set of alerts. When you add a storage system to a predefined group, the set of alerts is automatically applied to the storage system.
- In the web-based GUI, schedule probes and performance monitors to collect data about the storage system. Probes collect status, asset, and storage information about storage systems. Performance monitors collect metrics that measure the performance of storage systems.
- In the stand-alone GUI or web-based GUI, view detailed information about the storage system.

"Starting the Tivoli Storage Productivity Center stand-alone GUI" on page 41
The stand-alone GUI contains functions for monitoring the condition of storage resources, and all the tools for managing data, disk, fabric, and replication.

"Alerts and monitoring groups" on page 1685
Use alerts to notify you when conditions on monitored resources are detected by Tivoli Storage Productivity Center. Include resources in predefined monitoring groups to automatically associate a set of alerts with those resources.

"Collecting data" on page 295
Schedule probes and performance monitors to collect data about a storage system.

Connection properties for IBM SONAS systems:

Use the Configure storage subsystem connections page in the Configure Devices wizard to enter connection properties for the IBM SONAS systems that you are configuring. Tivoli Storage Productivity Center uses this information to collect data about the storage systems.

The following information is required when you add IBM SONAS systems:

IP Address

Enter the IP address or host name of the IBM SONAS system. You can enter an IPv4 or IPv6 address, for example:

- 127.0.0.1
- 2001:DB8:0:0:0:0:0:0/32

Select Key

You must use a Secure Shell (SSH) key to access the storage system. Select one of the following actions for the SSH key:

- **Upload new key:** Select this option to upload a new key. Public key content is extracted from the private key that is provided in the **Private**

SSH Key field and is uploaded to the IBM SONAS system. If you select this action, you must specify a value for the **Administrator User Name** and **Administrator Password** fields.

- **Use existing uploaded key:** Select this option if you want to use a public key that was already uploaded by using a method outside of Tivoli Storage Productivity Center, such as the IBM SONAS system web interface. Enter the location of the private key file that corresponds to the public key in the **Private SSH Key** field. You do not have to specify a value in the **Administrator User Name** and **Administrator Password** fields.
- **Use user password:** Select this option if you want to use a user name and password to connect to the file module. If you select this option, the **User Password** field is enabled and you can use a user name and password to connect to the file module.

  If you select **Use user password**, the **User Password** field is enabled and the **Administrator User Name** and **Administrator Password** fields are optional.

Administrator User Name
If you selected **Upload new key** in the **Select Key** list, enter the name of a user that has the Security Administrator role for the IBM SONAS system. Tivoli Storage Productivity Center uses this value during the setup process to contact the system and configure the SSH keys for the user. This user name must have privileges to modify other user accounts (Security Administrator role); otherwise, Tivoli Storage Productivity Center cannot configure the SSH keys.

If you selected **Use existing uploaded key**, this field is optional.

Administrator Password
Enter the password for the user name that you entered in the **Administrator User Name** field.

User Name
Enter an authorized user name for the storage system. Tivoli Storage Productivity Center uses this value when it configures and monitors the storage system. If the user name does not exist, it is created. This user name must have the Administrator role.

This field is active if you selected **Upload new key** in the **Select Key** list and entered a user name and password in the **Administrator User Name** and **Administrator Password** fields.

Select User
Select from a list of existing user names that are defined on the IBM SONAS system.

Private SSH Key
By default, the private key entry includes `${device.conf}\tpc_svc.pem`, which represents the Tivoli Storage Productivity Center default key `tpc_svc.pem`. The `tpc_svc.pem` key is in the `conf` directory where the Device server is installed. If you want to use your own key, the key must exist on the system where you are running the Tivoli Storage Productivity Center user interface. The key is uploaded to the Tivoli Storage Productivity Center server system. To identify a new key, click **Browse**.

The private SSH key must be in OpenSSH format or in PuTTY (`.ppk`) format that is not password protected. If you have a password

protected PuTTY .ppk key, use the PuTTYgen utility on a Windows operating system to convert that key to OpenSSH format by following these steps:

1. Start the PuTTYgen utility.
2. Load the `.ppk` file by clicking **File** > **Load private key**.
3. From the menu bar, click **Conversion** > **Export OpenSSH key conversion**.

This action saves the key in OpenSSH format.

Passphrase (Optional)
Enter the passphrase for the private SSH key. The passphrase protects the private SSH key file. If you do not have a passphrase, leave this field blank.

User Password
Enter the password for the user name that you entered in the User Name field.

All others

Add other storage systems by connecting to the CIM agents that manage the storage systems.

Adding a CIM agent that manages multiple storage systems: If a CIM agent manages multiple storage systems, all of the storage systems that are managed by the agent are added to Tivoli Storage Productivity Center.

The CIM agents for the storage systems must comply with SMI-S standards. Examples of other storage systems are EMC, Hitachi, and NetApp.

To add a storage system, complete the following steps:

1. In the navigation pane, select **Storage Resources** > **Storage Systems**.
2. Click **Add Storage System**.
3. Click the **All others** icon.
4. On the Discover page, enter the following information about a CIM agent:

   CIM agent host name or IP address
   The host name or IP address for the computer on which the CIM agent that manages the storage system is installed. Depending on what is supported in your environment, you can enter an Internet Protocol version 4 (IPv4) or IPv6 address. If you enter an IPv6 address, the preferred representation is written as eight groups of four hexadecimal digits. Example: 2001:DB95:0000:1234:0000:0000:5678:ABCD.

   User name, Password
   The user name and password that are used to authenticate to the CIM agent.

   CIM agent display name
   The name of the CIM agent, as specified by the agent provider. This name is displayed in the Tivoli Storage Productivity Center interface.

   CIM agent description
   A description of the CIM agent. This description can be any combination of alphanumeric characters.

Advanced

Expand **Advanced** to set the optional protocol, port, and namespace for the CIM agent.

Protocol

The version of the cim-xml protocol, http or https.

Port The port on which the CIM agent is listening. By default this port is 5989 for a secure connection and 5988 for an unsecured connection. See the documentation for the appropriate SMI-S CIM Provider for more information about available ports.

Namespace

This namespace within the CIM agent allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how IBM Tivoli Storage Productivity Center interacts with the CIM agent when retrieving information. For information about the namespace to use, see http://www-01.ibm.com/support/docview.wss?uid=swg21386446. Click the appropriate version for **Storage**.

To view a list of default namespaces for CIM agents for storage systems that is included with Tivoli Storage Productivity Center, open the following file in a text editor: *install_directory*/data/config/namespace.config, where *install_directory* represents the directory where the product is installed.

Check the documentation of the appropriate SMI-S CIM Provider or contact the storage system vendor to ensure that you use the most current namespaces.

5. On the Configure page, enter the following information. If multiple storage systems are managed by the CIM agent, select the storage system or systems that you want to configure. To select multiple storage systems, press Ctrl and click or Shift and click. The default configuration settings are applied to the remaining storage systems.

Display name

The name of the storage system that is displayed in the Tivoli Storage Productivity Center interface. If you do not enter a value, the default name is provided. This is the name that is provided with the storage system.

This field is inactive if you selected multiple storage systems.

Location

The physical location of the storage system. You can accept the default, **No Location**; enter a new location; or select a location that is already defined.

Data Collection

Schedule the probe and performance monitor for the storage system. The probe collects status, asset, and storage information about the storage system. The performance monitor collects metrics that measure the performance of the storage system.

Probe Enter the time and schedule for the storage system probe to run.

Enter the time in *hh:mm* format, where *hh* equals the hour and *mm* equals the minute. The time zone that is shown is the time zone where the Tivoli Storage Productivity Center Data server is located.

Performance monitor
Enable or disable the performance monitor for the storage system. If you enable the performance monitor, select the schedule for how often you want performance data to be collected.

Adding fabrics and switches

Add fabrics and switches for monitoring.

In the web-based GUI, you can add switches and fabrics for monitoring at the same time. You can add the following types of switches:

- Brocade
- Cisco
- QLogic

For a complete list of the switches and their versions that you can add, check the switches and directors information in the *Tivoli Storage Productivity Center Interoperability Matrix*.

Help tips in the GUI: To view descriptions of the information that you must enter for a fabric or switch, hover the mouse pointer over the related help icons in the wizard.

1. In the navigation pane, select **Network Resources** > **Switches**.
2. Click **Add Switch**.
3. Determine how to discover the switch and connected fabric.
   - Select **Use CIM agent as data source** to use a CIM agent as the data source. Tivoli Storage Productivity Center automatically discovers the switches that are managed by the agent. After this step, you can also select to use Simple Network Management Protocol (SNMP) agents for the switches.
   - Do not select **Use CIM agent as data source** to skip the use of a CIM agent as a data source. If you skip the use of a CIM agent as a data source, you are prompted to add switches and their fabrics through SNMP agents.

   The type of switch and functions that you want to enable on that switch and connected fabric determine the agent to use for monitoring. For example, to monitor the performance of a switch, you must use a CIM agent. To ensure full monitoring and functions for fabrics, you must also install Storage Resource agent on a server that is connected to the fabric.

   For more information about agent types and the fabric and switch functions that they enable, see "Agent types for switch and fabric functions" on page 294.
4. Specify the display name and location of the switches and fabrics that are discovered. For switches that were discovered through a CIM agent, you can also select to monitor them through SNMP.
5. Schedule data collection for the switches and fabric. The type of resource determines the data collection jobs that you can schedule:
   - Fabrics: probes
   - Switches: probes, performance monitors

Probes collect status and asset information. Performance monitors collect metrics that measure performance.

6. Complete the wizard.

After a switch or fabric is added, data collection jobs gather information about them. When data collection is complete, you can view detailed information about the switch or fabric in the web-based GUI and stand-alone GUI.

Agent types for switch and fabric functions

Depending on the functions that you want to enable, you can use a CIM agent, SNMP agent, or Storage Resource agent to monitor switches and fabrics.

The vendors of switches can help you determine which agents to use.

Table 51. Agent types for switch and fabric functions

Function > Switch	Brocade	Cisco [1]	QLogic
Switch Performance Monitoring	CIM agent	CIM agent	CIM agent
Switch and switch port information collected	Recommended: CIM agent Also supported: SNMP agent or Storage Resource agent	Recommended: CIM agent Also supported: SNMP agent or Storage Resource agent	SNMP agent or Storage Resource agent
Topology connectivity information collected	Recommended: CIM agent Also supported: SNMP agent or Storage Resource agent	SNMP agent or Storage Resource agent	SNMP agent or Storage Resource agent
Zoning information collected [2]	Recommended: CIM agent	A Storage Resource agent is required in each VSAN [3] [4] [5]	Storage Resource agent
Tivoli Storage Productivity Center Alerts	Recommended: CIM agent Also supported: SNMP agent or Storage Resource agent	SMP agent or Storage Resource agent	SNMP agent or Storage Resource agent
Hosts, endpoint devices, device-centric and host-centric information collected	Storage Resource agent	Storage Resource agent	Storage Resource agent
Switch “Sensors and Events”	SNMP agent	SNMP agent	SNMP agent

Table 51. Agent types for switch and fabric functions (continued)

Function > Switch	Brocade	Cisco [1]	QLogic
Tips: [1] The Storage Resource agent collects the information about a VSAN basis. Each VSAN is viewed as an individual SAN. The SNMP agent and CIM agent get the physical fabric information and can correlate the VSAN information to a physical infrastructure. The SNMP agent also collects some VSAN information. [2] "Zone information collected" indicates that information about the active zone set and the inactive zoning library is collected during a fabric probe. [3] For many Cisco proprietary zone members, the values are displayed as zeros. [4] If a zoning configuration has Cisco proprietary zone members, then you cannot modify the zoning for that configuration through Tivoli Storage Productivity Center. You must use the element manager to remove proprietary zone members or do zoning. [5] Orphan Cisco zones are not reported. [6] Tivoli Storage Productivity Center does not get a transaction lock on the switch while zone changes are made on the fabric.			

Collecting data

Determining the data that you want to gather about storage resources is critical to helping you implement a storage management strategy. Tivoli Storage Productivity Center provides different jobs for collecting data. When you determine the types of data that you want to collect about resources, schedule the corresponding data collection jobs for those resources.

Data collection jobs include probes, performance monitors, scans, and pings.

Tips:

- The type of data collection job that you want to schedule determines which Tivoli Storage Productivity Center GUI to use.
- In the web-based GUI, data collection jobs are automatically scheduled when resources are added for monitoring. In the stand-alone GUI, you must schedule data collection jobs after resources are added.

Table 52. Data collection jobs, resources, and the GUI where you define the jobs

Data collection job	Resources	GUI
Probe Use probes to collect status, asset, and capacity information about resources.	Storage systems Servers Hypervisors Fabrics Switches	Web-based GUI: • Schedule a probe when you add a resource for monitoring. • Modify the schedule for a probe at any time on the list and details pages for a resource. For example, to modify the schedule for a storage system, go to the Storage Systems page or the Storage System details page. Learn more...

Table 52. Data collection jobs, resources, and the GUI where you define the jobs (continued)

Data collection job	Resources	GUI
Performance monitor Use performance monitors to collect metrics that measure the performance of storage systems and switches.	Storage systems Switches	Web-based GUI: • Schedule a performance monitor when you add a resource for monitoring. • Modify the schedule for a performance monitor at any time on the Performance page, and on the list and details pages for the resource. Learn more...
Scan Use scans to collect file usage and trending information about servers or hypervisors.	Servers Hypervisors	Stand-alone GUI: Schedule a scan after a resource is added for monitoring. You can define a scan in the **Data Manager** > **Monitoring** > **Scans** node in the navigation tree of the stand-alone GUI. Learn more...
Ping Use pings to collect information about the availability (uptime) of servers.	Servers	Stand-alone GUI: Schedule a scan after a resource is added for monitoring. You can define a ping in the **Data Manager** > **Monitoring** > **Pings** node in the navigation tree of the stand-alone GUI. Learn more...
TPC Server probe Use TPC server probes to gather the information that was collected by multiple Tivoli Storage Productivity Center servers.	Tivoli Storage Productivity Center servers	Stand-alone GUI: You can define a TPC Server probe in the **IBM® Tivoli Storage Productivity Center** > **Monitoring** > **TPC Server Probes** node in the navigation tree of the stand-alone GUI. Learn more...

Collecting asset and status data

Use probes to collect asset, status, and storage data about monitored resources in your environment. Probes are also used to discover information about new or

removed disks and file systems. You can view information about probes on the resource list and resource details pages, for example, the Storage Systems page, or the Storage Systems details page.

Creating probes

You can create and schedule probes to gather information about different storage resources at regular intervals. *Probes* are data collection jobs that provide detailed asset and status information about your storage resources.

You can set up a data collection schedule from a resource list page, or from the details page of a particular storage resource. For example, to create a data collection schedule for a storage system, use the Storage Systems page.

1. From the navigation pane in the web-based GUI, go to the resource list page for a resource type. For example, to create a probe for a storage system, click **Storage Resources** > **Storage Systems**.
2. Locate the resource that you want to monitor.
3. Right-click the resource row, and select **Data Collection** > **Schedule**.
4. On the Data Collection Schedule page, set the probe to **Enabled**.
5. Enter the time of the data collection, select the frequency, and click **Save**. A message is displayed that confirms that a probe schedule is defined for the storage resource.
6. Click **Close**.

The data collection job runs according to the defined schedule.

Verifying that a probe is running for a resource

Verify that a probe is running for a resource to ensure that Tivoli Storage Productivity Center is collecting asset and status information about the resource.

To verify that a probe is running for a resource, use the resource list page in the web-based GUI.

1. From the navigation pane, go to a resource list page for a resource type. For example, to verify that a probe is running for a storage system, click **Storage Resources** > **Storage Systems**.
2. Locate the resource that you want to check.
3. View the Probe Status column. A value of Successful or Running indicates that the probe is collecting asset and status information about the resource. The Last Successful Probe column shows the most recent date and time when the probe successfully collected data.
4. Optional: If the Probe Status column shows a status of Failed or Warning, to view the error or warning messages, right-click the resource row and click **Data Collection** > **Open Probe Logs**.
5. Optional: If the Probe Status column shows a status of Never probed, to schedule a probe for the resource, right-click the resource row and click **Data Collection** > **Schedule**.

Modifying probes

Detailed information about monitored resources is obtained from data collection jobs called *probes*. To modify a probe in Tivoli Storage Productivity Center, use the resource list page for the monitored resource in the web-based GUI.

1. From the navigation pane, go to the resource list page for a resource type. For example, to modify a data collection schedule for a storage system, click **Storage Resources** > **Storage Systems**.

2. Locate the resource that you want to monitor.
3. Right-click the resource row, and select **Data Collection** > **Schedule**.
4. On the Data Collection Schedule page, modify the data collection details, and click **Save**. A message is displayed to indicate that the probe is updated.
5. Click **Close**.

The data collection job runs according to the modified schedule. If you disabled the probe, Tivoli Storage Productivity Center stops collecting asset and status data for the monitored resource.

Configuring alerts for probes

Configure the alert notifications that are generated if a probe fails to run.

When a probe is scheduled for a resource in the web-based GUI, an entry for that probe is created in the stand-alone GUI. The entry is displayed in the navigation tree, under **IBM Tivoli Storage Productivity Center** > **Monitoring** > **Probe Alert Configurations**. The name of the entry is *userID*.Probe_*resourcename*, where *userID* represents the user ID who created the probe and *resourcename* is the name of the resource that is being probed. For example, "user1.Probe_DS8000a" represents a node for a probe that was created by user1 for a storage system named DS8000a.

You can edit the alert configuration for a resource if you want to modify how you are notified when a probe fails and which actions are taken. By default, if a probe fails, an error is recorded in the alert log. If you want to specify more actions if a probe fails for a resource, such as receiving an email and automatically running a script, you can modify its alert configuration.

Tip: You can view the alert log on the **Home** > **Alerts** page in the web-based GUI.

To modify the alert configuration for a probe, complete the following steps:

1. Start the stand-alone GUI.
2. In the navigation tree, expand **IBM Tivoli Storage Productivity Center** > **Monitoring** > **Probe Alert Configurations**.
3. Click the node for the alert configuration that you want to modify. For example, if you want to modify the alert configuration for the probe of a storage system named "Storwize_IBM02", click the node **user1.Probe_Storwize_IBM02**.
4. In the **Triggered actions** section, determine the actions that occur as a result of a probe that fails. For information about triggered actions, see “Triggered actions for alerts” on page 1682.
5. To save the alert configuration, click **File** > **Save**.

If a probe for the resource fails, Tivoli Storage Productivity Center takes the actions that are defined in the alert configuration for that resource.

Starting probes

Data is collected according to the probe schedule that is defined when a data collection job is created. To immediately collect data about a resource, you can use the **Start Probe** action in the web-based GUI.

If a data collection job is not defined for the resource, use the **Schedule** action to create a probe.

1. From the navigation pane, go to the resource list page for a resource type. For example, to start a probe for a storage system, click **Storage Resources** > **Storage Systems**.
2. Locate the resource that you want to probe.
3. Right-click the resource row and select **Data Collection** > **Start Probe**.

The probe begins, and data is collected from the monitored resource. The status of the probe changes to Running.

Viewing probes for a specific resource

You can view information for specific resources on the resource details pages. For example, use the Storage Systems details page to view the date and time when data was last collected.

1. From the navigation pane in the web-based GUI, select a type of top-level resource. For example, if you want to view the resource information for a storage system, click **Storage Resources** > **Storage Systems**.
2. Locate the resource that you want to view information for.
3. Right-click the resource row and select **View Details**.
4. In the General section of the resource details page, click **Data Collection**.

The Data Collection pane is displayed, which contains information about the data collection jobs that are defined for the resource.

Viewing probe logs

A probe log file is created for each run of a probe on a storage resource. Probe logs include detailed information about the status, actions, and progress of a probe.

1. From the navigation pane in the web-based GUI, go to a resource list page for a resource type. For example, to view the probe logs for a storage system, click **Storage Resources** > **Storage Systems**.
2. Locate the resource that you want to view log files for.
3. Right-click the resource and select **Data Collection** > **Open Probe Logs**.

The Logs page opens in a separate window and shows log entries for each of the previous runs of the probe.

You can view the following information on the Logs page:

- The overall status of a probe. The icon that is shown in the Select a log list represents the most critical status that was generated by an action in the job run.
- The status for each probe action.
- The date and time when an action was completed. The date, time, and time zone of an action is shown in the Date and Time column.
- The ID of the message that is associated with an action. You can click the value in the ID column to view more information about a message.
- The description of an action.

Collecting performance data

Before you can monitor the performance of the resources in your environment, you must collect data about those resources. Tivoli Storage Productivity Center uses performance monitors to collect metrics for measuring the performance of storage systems and switches. Use the Performance Monitors page to view and manage performance monitors and to track threshold violations.

You must schedule and run performance monitors before you can complete other tasks, such as optimizing storage tiering, balancing pools, and running performance reports.

Before you can view performance information for resources, you must complete the following tasks:

- Add the storage system or switch for monitoring by Tivoli Storage Productivity Center.
- Schedule a performance monitor to collect performance data about the resource.
- Define a threshold in an alert so that you are notified when the performance of a resource might represent a potential problem. For example, you can define an alert threshold that notifies you when the total I/O port rate for a storage system falls outside a specified range. You can define alerts in the stand-alone GUI.

The performance information that is gathered by Tivoli Storage Productivity Center includes metrics that measure the performance of volumes, ports, and disks. You can view performance metrics in the web-based GUI and stand-alone GUI at the following locations:

- Resource list and detail pages for switches and storage systems. For example, you can view performance metrics that are related to switches on the Switches page in the web-based GUI.
- Performance reports in the web-based GUI and stand-alone GUI. For example, you can access reports from the web-based GUI that show data rates, I/O rates, response times, and many other performance metrics.

Creating performance monitors

Performance monitors are data collection jobs that collect metrics for measuring the performance of storage systems and switches. You can create performance monitors in the web-based GUI when you add resources for monitoring or at a later time.

To create performance monitors, you must have Administrator privileges.

To create a performance monitor for a resource in your environment, you can use the list and detail pages for the resource. For example, to create a performance monitor for a switch, use the Switches page.

Restrictions:

- You can define only one performance monitor for a resource.
- You can create performance monitors for the following resources only:
  - Storage systems
  - Switches that are managed by a CIM agent
- You cannot schedule a performance monitor for multiple resources at the same time. The **Schedule** action for performance monitors is available only if you click a single resource row.

1. From the navigation pane, go to a resource list page for a resource type. For example, to create a performance monitor for a storage system, click **Storage Resources** > **Storage Systems**.
2. Locate the resource that you want to monitor.
3. Right-click the resource row, and select **Data Collection** > **Schedule**.
4. On the Data Collection Schedule page, set the performance monitor to **Enabled**. You can set the performance monitor to **Enabled** only if the probe is set to **Enabled**.

5. Select the performance monitor interval, and click **Save**. The interval represents the number of minutes over which samples of performance data are averaged.

If a successful probe run is completed for the resource, the performance monitor runs according to the defined interval.

To check the progress of the performance monitor, you can complete the following actions:

- From the resource list page, right-click the resource row, and select **Data Collection** > **Open Performance Monitor Logs**. You can view detailed informational, warning, and error messages that are related to the performance monitor. Use this information to troubleshoot any errors that might occur when the performance monitor runs.
- View the Performance Monitor Status column on the resource list page.

Verifying that a performance monitor is running for a resource

Verify that a performance monitor is running for a resource to ensure that Tivoli Storage Productivity Center is collecting performance metrics about the resource. Use performance monitors to collect metrics about the performance of storage systems and switches.

To verify that a performance monitor is running for a resource, use the resource list page in the web-based GUI.

1. From the navigation pane, go to a resource list page for a resource type. For example, to verify that a performance monitor is running for a storage system, click **Storage Resources** > **Storage Systems**.
2. Locate the resource that you want to check.
3. View the Performance Monitor Status column. A value of Completed or Running indicates that the performance monitor is collecting performance metrics about the resource. The Last Successful Monitor column shows the most recent date and time when the performance monitor successfully collected data.
4. Optional: If the Performance Monitor Status column shows a status of Running with problems, Completed with warnings, Failed, or Canceled, you can view the logs to troubleshoot the error or warning messages. To view the performance monitor logs, right-click the resource row and click **Data Collection** > **Open Performance Monitor Logs**.
5. Optional: If the Performance Monitor Status column shows a status of Disabled, to enable the performance monitor for the resource, right-click the resource row and click **Data Collection** > **Schedule**. On the Data Collection Schedule page, select **Enabled**.

   Tip: You can set the performance monitor to **Enabled** only if the probe is set to **Enabled**. If a probe is set to **Disabled**, the performance monitor cannot run.

Modifying performance monitors

To modify all of the performance monitors that are used by Tivoli Storage Productivity Center, use the Performance Monitors page in the web-based GUI. To modify performance monitors for specific resources, use the resource list and resource detail pages.

To modify performance monitors, you must have Administrator privileges.

You can enable or disable performance monitors and modify the interval. The interval represents the number of minutes over which samples of performance data are averaged.

1. From the navigation pane in the web-based GUI, click **Home** > **Performance Monitors**. Detailed information about all performance monitors is shown in the **Performance Monitors** tab.
2. Locate the performance monitor that you want to modify.
3. Right-click the performance monitor row, and select **Schedule**.
4. On the Data Collection Schedule page, modify the performance monitor schedule. You can enable or disable the performance monitor and change the performance monitor interval. If the performance monitor is disabled, you cannot modify the interval value.
5. Click **Save**.

The performance monitor runs according to the modified schedule. If you disabled the performance monitor, Tivoli Storage Productivity Center stops collecting performance metrics for the storage system or switch.

To check the progress of the performance monitor, you can complete the following actions on the Performance Monitors page:

- View information such as the status of the performance monitor and the most recent date and time when performance data was collected about a resource.
- Right-click the performance monitor row, and select **Open Logs**. You can view detailed informational, warning, and error messages that are related to the performance monitor job. Use this information to troubleshoot any errors that might occur when the job runs.

Configuring alerts for performance monitors

Configure the alert notifications that are generated if a performance monitor fails to run.

When a performance monitor is scheduled for a switch or storage system in the web-based GUI, an entry for that performance monitor is created in the stand-alone GUI. The entry is displayed in the navigation tree. For a storage system, an entry is displayed under **Disk Manager** > **Monitoring** > **Subsystem Performance Alert Configurations**. For a switch, an entry is displayed under **Fabric Manager** > **Monitoring** > **Switch Performance Alert Configurations**.

The name of the entry is *userID*.Performance_Monitor_*resourcename*, where *userID* represents the user ID who created the performance monitor and *resourcename* is the name of the associated resource. For example, "user1.Performance_Monitor_brocade1" represents a node for a performance monitor that was created by user1 for a switch named brocade1.

You can edit the alert configuration for a resource if you want to modify how you are notified when a performance monitor fails and which actions are taken. By default, if a performance monitor fails, an error is recorded in the alert log. If you want to specify more actions if a performance monitor fails for a resource, such as receiving an email and automatically running a script, you can modify its alert configuration.

Tip: You can view the alert log on the **Home** > **Alerts** page in the web-based GUI.

To modify the alert configuration for a performance monitor, complete the following steps:

1. Start the stand-alone GUI.
2. Choose one of the following options:

Option	Description
Storage systems	In the navigation tree, expand **Disk Manager** > **Monitoring** > **Subsystem Performance Alert Configurations**.
Switches	In the navigation tree, expand **Fabric Manager** > **Monitoring** > **Switch Performance Alert Configurations**.

3. Click the node for the alert configuration that you want to modify. For example, if you want to modify the alert configuration for the performance monitor of a switch named "switch_brocade02", click the node **user1.Performance_Monitor_switch_brocade02**.
4. In the **Triggered actions** section, determine the actions that occur as a result of a performance monitor that fails. For information about triggered actions, see "Triggered actions for alerts" on page 1682.
5. To save the alert configuration, click **File** > **Save**.

If a performance monitor for the resource fails, Tivoli Storage Productivity Center takes the actions that are defined in the alert configuration for that resource.

Starting and stopping performance monitors

To start or stop performance monitors for storage systems and switches, use the Performance Monitors page in the web-based GUI.

To start or stop performance monitors, you must have Administrator privileges.

Restrictions:

- A successful probe must be completed on the resource before a performance monitor can start. If a successful probe is not run, you cannot start the performance monitor for the resource.
- The performance monitor actions are available only if you click a single performance monitor row. You cannot start or stop multiple performance monitors at the same time.

1. From the navigation pane in the web-based GUI, click **Home** > **Performance Monitors**.
2. Click the **Performance Monitors** tab.
3. Locate the performance monitor that you want to start or stop.
4. Right-click the performance monitor row, and select **Start** or **Stop**.

Starting or stopping a performance monitor has the following effects:

Start Starting a performance monitor starts the immediate collection of data from the monitored resource. The status of the performance monitor changes to Running.

Stop Stopping a performance monitor stops the collection of performance data from the resource. If a performance monitor is running, it stops.

Performance data is not collected from the resource until you restart the performance monitor. The status of the performance monitor changes to Canceled.

To check the progress of the performance monitor, you can complete the following actions on the Performance Monitors page:

- View information such as the status of the performance monitor and the total number of threshold violations that were detected during the last sampled time of the performance monitor.
- Right-click the performance monitor row, and select **Open Logs**. You can view detailed informational, warning, and error messages that are related to the performance monitor. Use this information to troubleshoot any errors that might occur when the performance monitor runs.

Viewing performance monitors

To view all of the performance monitors that Tivoli Storage Productivity Center uses to collect metrics for measuring the performance of resources, use the Performance Monitors page. To view the performance monitors for specific resources, use resource list and resource details pages.

Viewing all performance monitors:

To view all of the performance monitors that Tivoli Storage Productivity Center uses to collect metrics for measuring the performance of storage systems and switches, use the Performance Monitors page in the web-based GUI.

The Performance Monitors page displays one row for each storage system or switch that Tivoli Storage Productivity Center is monitoring. If a performance monitor is not scheduled for a resource, the status column is set to Disabled for the performance monitor row.

Before Tivoli Storage Productivity Center can collect performance metrics for a resource, the resource must fulfill the following conditions:

- A probe must be run for the resource.
- For switches, the switch must be managed by a CIM agent.

1. From the navigation pane in the web-based GUI, click **Home** > **Performance Monitors**. On the **Performance Monitors** tab, the information about performance monitors is organized into columns. These columns include the status of performance monitors, the names of the monitored resources, and the most recent dates and times when data was successfully collected.
2. Optional: View the status icons at the top of the page to view a summary of performance monitor statuses. This summary includes the number of performance monitors that have a Running, Running with problems, Failed, or Not Running status.

Viewing performance monitors for specific resources:

To view the performance monitors for specific storage systems and switches, use the resource detail pages in the web-based GUI. For example, use the Storage Systems detail page to view performance monitor information for a specific storage system.

If performance monitor information is not displayed on the resource detail page, this indicates that the resource does not support the collection of performance metrics by Tivoli Storage Productivity Center.

1. From the navigation pane in the web-based GUI, select a type of top-level resource. For example, if you want to view the performance monitor for a storage system, click **Storage Resources** > **Storage Systems**.
2. Right-click a resource in the list, and select **View Details**. The resource detail page is displayed.

   In the **General** section, the number next to **Data Collection** represents the number of data collection jobs that are associated with the resource. For example, **Data Collection (2)** indicates that you can view a performance monitor and a probe for the resource.

   The icon next to the **Data Collection** link represents the most critical status of the data collection jobs that are associated with the resource. For example, if a job failed, the following icon is shown:

   Data Collection (2)
3. In the **General** section, click **Data Collection** to view information about the performance monitor for the switch or storage system. On the **Data Collection** pane, you can view information such as the status of the most recent run of a performance monitor, the performance monitor interval, and the most recent data and time that data was successfully collected.

   Tip: You can use the columns on the resource list pages to view information about the performance monitors that are defined for storage systems or switches. For example, use the Storage Systems page to view performance monitor information for storage systems.

Viewing performance monitors logs:

Use performance monitor logs to view detailed information about the status, actions, and progress of a performance monitor. You can use this information to troubleshoot any errors that might occur when a performance monitor is running.

Performance monitor logs contain informational, warning, and error messages that are related to each action that is taken during the processing of a performance monitor.

1. From the navigation pane in the web-based GUI, click **Home** > **Performance Monitors**.
2. In the **Performance Monitors** tab, locate the performance monitor that you want to view.
3. Right-click the performance monitor row, and select **Open Logs**. The Logs page opens and shows log entries for the most recent performance monitor log. A new log file is created when you stop and restart a performance monitor.
4. Optional: On the Logs page, to view a previous performance monitor log, select a log from the **Select a log** list.
5. Optional: To view only the log entries that have a Warning or Error status, select an option from the **Show all** list. You can choose to view only entries that have the following statuses:
   - Only error entries
   - Only warning entries
   - Error and warning entries
6. Optional: To view an explanation of the message that is associated with a log entry, click the link in the **ID** column.

The information on the Logs page is automatically updated every 30 seconds. New entries are added to the end of a log. You can view the following information on the Logs page:

- The overall status of a performance monitor. The icon that is shown in the **Select a log** list represents the most critical status that was generated by an action in the job run.
- The status for each performance monitor action.
- The date and time when an action was completed. The date, time, and time zone of an action is shown in the Date and Time column.
- The ID of the message that is associated with an action. You can click the value in the ID column to view more information about a message.
- The description of an action.

Tip: You can use the navigation and search functions of the web browser to locate information on the page.

Collecting information about shares on storage systems

To collect information about shares in file systems, you must modify the `TPCD.config` file. When you provision shares, you can then view such information as the name of the share, the name of the file system, and the protocols that are used to share files.

Add **`saveNonRoot=1`** to the `TPCD.config` file, to view the information that is collected about shares in the web-based GUI.

1. Stop the Data server.
2. Open the `TPCD.config` file.

   For Windows operating systems:
   By default, the `TPCD.config` file is in the `C:\Program Files\IBM\TPC\data\config\` directory.

   For Linux or UNIX operating systems:
   By default, the `TPCD.config` file is in either the `usr/IBM/TPC/data/config/` directory or the `opt/IBM/TPC/data/config` directory.
3. In the `[server]` section of the `TPCD.config` file, add the following line of code:
   ```
   saveNonRoot=1
   ```
4. Save the `TPCD.config` file.
5. Restart the Data server.

When you provision new shares, you can view information about the shares on the Provisioning task page in the web-based GUI. On the Shares page, you can view more information about shares, such as the servers that have access to the share.

Collecting storage usage and trending information

Use scans to collect statistics about the usage and trending of storage consumption of storage resources in your environment. To create and manage scans, use the stand-alone GUI. You can view the information that is collected by scans in reports and use it as part of a quota analysis.

Use a Storage Resource agent to collect information for a scan. The Storage Resource agent must be deployed on the storage resources that you want to monitor. By running scans and generating the results of those scans in reports, you can complete the following tasks:

- View when files were created, accessed, and modified and by what group or user
- Map storage resources to the consumers of those resources
- Generate historical views of storage consumption and utilization that show usage trends over time. System administrators can use these views to project and plan for future storage consumption.
- Collect detailed RDBMS space usage reports at the table, index, and user levels
- Gather individual user space usage at the table space, database, instance, computer, and network wide levels
- Collect detailed space usage information about monitored tables and their associated indexes. For example, when you scan UDB tables, this information includes many of the calculated REORGCHK values. These include Used Page %, Used Space %, and Cluster Ratio.
- Historically track database, table space, and user usage and use that data to forecast and display future usage
- Gather information for quotas that can be created to trigger alerts (and create reports) when users exceed a specified amount of space usage
- (UDB only) Perform (optionally) runstats on all tables involved in the scan process
- Assign costs to space used at the user level when combined with Data Manager for Chargeback

Note:

1. A licensed agent must be installed on a computer that contains instances against which you want to run a scan. You must run a probe against an instance before you can select any table spaces against which you want to run a scan.
2. The login ID that Data Manager uses to log in to Microsoft SQL Server instances that you want to scan must have "Permit" access.

Creating scans

Create and schedule a scan to gather usage information about your storage resources. To create a scan, use the stand-alone GUI. You can specify which file systems, computers, clusters, directories, databases, and other storage resources that you want to scan. You can define any number of scans to gather information about different storage resources at different times.

To create a scan, complete the following steps:

1. In the stand-alone GUI, expand **Data Manager** > **Monitoring** > **Scans**. For Data Manager for Databases, expand **Data Manager for Databases** > **Monitoring** > **Scans**.
2. Right-click the **Scans** node and select **Create Scan**.
3. For Data Manager, define the following scan properties:
   - The computers and file systems whose storage usage and consumption you want to monitor. For more information, see “Including storage resources in a scan” on page 308.
   - The directory groups to scan. For more information, see “Selecting directory groups to scan” on page 309.
   - The profiles to use during a scan. For more information, see “Using profiles in a scan” on page 311.
   - The schedule when you want a scan to run. For more information, see “When to run the scan” on page 312.

- An alert to notify you if the scan fails. For more information, see "Defining alert values for a scan" on page 313.

4. For Data Manager for Databases, define the following scan properties:
   - The databases, table spaces, instances, and computers whose storage usage and consumption you want to scan. For more information, see "Selecting databases, tablespaces to scan" on page 310.
   - The table groups that you want to scan. For more information, see "Select the table groups to scan" on page 310.
   - The profiles to use during a scan. For more information, see "Using profiles in a scan" on page 311.
   - The schedule when you want a scan to run. For more information, see "When to run the scan" on page 312.
   - The options for Oracle and UDB. For more information, see "Specify options for Oracle and UDB" on page 312.
   - An alert to notify you if the scan fails. For more information, see "Defining alert values for a scan" on page 313.
5. Click to save the scan. You must save the scan before Tivoli Storage Productivity Center can collect data about storage resources.
6. Enter a name for the scan and click **OK**. An entry for the scan is displayed under the **Scans** node in the navigation tree. The scan is run at the date and time that is specified in the schedule definition.
7. Use the Job Management panel to manage view information about the runs and jobs that are related to the scan. For example, use the Job Management panel to view the status of a run for a schedule, the date and time when it was run, and the log files of the jobs that are associated with the run.

Related reference:

"Managing schedules, jobs, and runs" on page 605
Use the Job Management panel in the stand-alone GUI to view and manage the schedules, runs, and jobs that are related to the resources that are monitored by Tivoli Storage Productivity Center.

Including storage resources in a scan:

In a scan, include the computers, computer groups, file systems, file system groups, and clusters whose storage usage and consumption that you want to monitor.

In the stand-alone GUI, use the **Filesystems** page, to select the storage resources that you want to scan. Include clusters, individual file systems and computers, or file system and computer groups (that have been predefined using the "Working with groups of resources" on page 337 facility).

The **Current Selections** list box shows the storage resources that are selected for the scan. The **Available** list box shows which storage resources are still available for selection in the scan.

Note: A cluster name is displayed next to individual computers that are members of a PowerHA SystemMirror for AIX or Microsoft Cluster Server cluster. Cluster resource groups are identified with a special icon. Clusters displayed in the **Available** list box cannot be expanded to show member nodes. You can move a cluster to the **Current Selections** list box to view and, if necessary, exclude member nodes.

To include storage resources in a scan:

1. Highlight the file systems, file system groups, computer groups, computers, and clusters you want to include in the scan in the **Available** list box.

   Note: Before you can include file systems in a scan, you must probe the computers that contain those file systems. If you do not probe a computer before you define a scan, the file systems on that computer are unavailable for the scan.

2. Click >>. The storage resources that you selected are displayed in the **Current Selections** list box.

To remove storage resources from a scan:

- Highlight the file systems, file system groups, computer groups, computers, and clusters that you want to remove and click <<. This action removes the storage resources from the **Current Selections** list box and moves them to the **Available** list box.

To exclude one or more file systems or computers from a selected cluster or group:

Perform the following steps if you want to exclude specific file systems or computers from a cluster or group that you included in a scan:

1. Double click the cluster or group name in the **Current Selections** list box.
2. Right-click the file systems or computers that you want to exclude. A pop-up menu is displayed.
3. Select **Exclude**. This action creates an **Excludes** section in the **Current Selections** list box that displays the names of the file systems or computers that you excluded from a scan. This action is useful when you want to select a group or cluster for scanning, but you want to exclude a few file systems or computers from that monitoring.

Selecting directory groups to scan:

This topic describes how to select the directory groups that you want to scan.

In the stand-alone GUI, use the Directory Groups tab to select the groups of directories that you want to scan. Including directory groups within a scan enables you to monitor and gather information about specific directories across a number of machines.

The **Selected Directory Groups** list box shows what directory groups have been included in the scan. The **Available Directory Groups** list box shows the groups that are still available for selection. You can select directory groups that have been predefined using the groups facility.

To select a directory group:

1. Highlight the desired directory group in the **Available Directory Groups** list box.
2. Click >>. The group you selected appears in the **Selected Directory Groups** list box.

To remove a directory group from a scan:

- Highlight the groups you want to remove and click << . This will remove the groups from the **Selected Directory Groups** list box and move them to the **Available Directory Groups** list box.

To create a new directory group:

- Click **New Directory Group**. The **New Directory Group** window opens, enabling you to create a new group. See "Creating directory groups" on page 338 for more information on how to create a directory group.

Selecting databases, tablespaces to scan:

This topic describes how to select the databases and tablespaces that you want to scan.

In the stand-alone GUI, use the Database-Tablespaces page to select what databases, tablespaces, instances, and computers whose storage usage and consumption you want to monitor. You can include single or groups of databases, tablespaces, and computers against which to run the scan.

The **Current Selections** list box shows the databases, tablespaces, tablespace groups, computers, and computer groups that have been selected for the job. The **Available** list box shows what items are still available for selection.

To include a database, tablespace, computer, instance, or group in a scan:

1. Highlight the desired databases, tablespaces, computers, instances, or groups in the **Available** list box.
2. Click >> . The databases, tablespaces, computers, instances, or groups you selected appear in the **Current Selections** list box.

   For example, if you select a computer to include in a scan, that job will scan all the databases and/or tablespaces on all the instances that reside on that computer.

To remove a database, tablespace, computer, instance, or group from a scan:

- Highlight the items you want to remove and click << . This will remove the databases, tablespaces, computers, instances, or groups from the **Current Selections** list box and move them to the **Available** list box.

Select the table groups to scan:

Select the table groups that you want to scan.

In the stand-alone GUI, use the Table Groups page to select the table groups you want to scan. A table group represents a group of defined tables. By choosing a table group on this window, you can perform more detailed scans of tables that might not appear within other Data Manager for Databases reports (such as the Largest Table or Monitored Tables reports).

Note: Before you can view detailed information about tables in an RDBMS, you must have 1) defined a table group that contains the tables about which you want to view storage information and 2) scheduled and run a scan against that table group.

The **Selected Table Groups** list box shows the table groups that have been selected for the job. The **Available Table Groups** list box shows the table groups that are still available for selection.

To include a table group in a scan:

1. Highlight the desired group in the **Available Table Groups** list box.
2. Click . The group you selected appears in the **Selected Table Groups** list box.

To remove a table group from a scan:

- Highlight the groups you want to remove and click . This will remove the groups from the **Selected Table Groups** list box and move them to the **Available Table Groups** list box.

To create a new table group:

- Click **New Table Group**. The **New Table Group** window opens, enabling you to create a new group.

Using profiles in a scan:

Select the profile to use as part of a scan. Profiles enable you to control what files are scanned and what statistics are gathered.

In the stand-alone GUI, use the **Profiles** page to select which profiles to use during a scan. Using profiles is a powerful method for determining the statistics that you want to gather during scans. Keep in mind that the more profiles you include within a scan, the longer the scan will take to gather the statistics that are specified by those profiles.

To include profiles in a scan:

1. Highlight the desired profiles in the **Available Profiles** list box.
2. Click the for the **Profiles to apply to Filesystems** list box to apply the selected profiles to file systems; click the for the **Profiles to apply to Directories** list box to apply the selected profile(s) to directories. The profiles that you selected appear in either the **Profiles to apply to Filesystems** or **Profiles to apply to Directories** list box

   Note: To collect statistics according to Data Manager default settings, select all of the default profiles.

To remove profiles from a scan:

- Highlight the profiles you want to remove and click the corresponding . This will move the selected profiles to the **Available Profiles** list box.

To create a new profile:

- Click **New Profile**. The **Statistics** tab on the **Create Profile** window is displayed, enabling you to create a new profile. See “Use profiles” on page 315 for more information.

Related concepts:

"Use profiles" on page 315
In the stand-alone GUI, use profiles to specify what statistical information is gathered and to fine-tune and control what files are scanned during a scan. You can also use profiles with some reports to determine what statistics are displayed.

When to run the scan:

In the stand-alone GUI, schedule when and how often to run a scan.

You have the following options for determining when to run a scan:

- Run the job immediately
- Run the job once at a specified time and date. Use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM.
- Run the job repeatedly according to a schedule you specify. You can:
  - use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM when the job should begin running.
  - indicate how often the job should run. You can use the list boxes to specify the number of minute(s), hour(s), day(s), week(s), or month(s).
  - select which days you wish to run the job.

Use the fields in the **How to handle timezones** section to indicate the time zone that Data Manager should use when scheduling and running the scan. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:

- 9:00 AM of the time zone where the server resides
- 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

Use the timezone that the server runs in. Select this option to use the time zone of the location where the Data server resides.

Use this timezone. Select this option to indicate a specific time zone for when to run the job. When you select this option, a list box is activated. Select a timezone from the list box.

Specify options for Oracle and UDB:

Specify options for Oracle and UDB when defining a scan.

In the stand-alone GUI, use the Options page to do the following:

- Determine if the job should scan the freelist chains that are associated with the segments in your Oracle environment. Free list chains represent the pages allocated to a segment that have enough free space for additional rows. Note that checking this option will increase the time required by the scan to complete.
- Set the number of pages Data Manager for Databases should scan on a chain before stopping.
- Determine if the job should gather segment extent data for locally-managed tablespaces.
- Run the DB2 utility runstats against the tables being scanned. Checking this box will enable Data Manager for Databases to collect the most up-to-date storage statistics about tables from the DB2 catalog. Note that checking this option will increase the time required by the scan to complete, but will ensure that table and

user statistics are up-to-date in UDB/DB2 databases. Alerts based on table and user stats will not function until runstats is run.

Defining alert values for a scan:

This topic describes how to set an alert that will be triggered if the scan job fails.

In the stand-alone GUI, use the Alert page to define an alert to notify you if a scan fails. You can choose to be notified by SNMP trap, IBM Tivoli Enterprise Console event, login notification, Windows event log, UNIX Syslog, running a script, and e-mail. If you select SNMP trap, Tivoli Enterprise Console event, or e-mail, you must also define a destination (for example, an e-mail address). Define this destination by expanding **Administrative Services** > **Configuration** > **Alert Dispositions**.

To define an alert for a scan, complete the following steps:

1. Click the **Alert** tab; the Alert page opens.
2. In the **Triggering Condition** area, select Scan Failed. This indicates that an alert will be triggered if the scan fails to run.
3. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

   SNMP trap
   An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

   TEC / OMNIbus Event
   A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the `tivoliSRM.baroc` file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, or running programs.

   Login notification
   The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

   Login ID
   The ID of the user who receives alert notification upon logging in to the system.

   Windows Event Log, UNIX Syslog
   The ID of the user who receives alert notification upon logging in to the system.

   Event Type (Windows Event Log only)
   The type of event that is recorded to the OS log.

   Facility (UNIX Syslog only)
   The location where the UNIX Syslog is sent. You can select User or Local.

   Level (UNIX Syslog only)
   The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters that are listed in the **Specify Script** window are passed to the script when the alert is triggered.

Email An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit email** to change the actual email message.

4. Click **File** > **Save** to save the scan.
5. Enter a name for the scan and click **OK**. An entry for the scan appears under the **Scans** node.

Scanning a Storwize V7000 Unified file system:

In the stand-alone GUI, create and schedule a scan to gather usage information about the file storage in a IBM Storwize V7000 Unified file system.

To scan a Storwize V7000 Unified file system, complete the following steps:

1. Export the Storwize V7000 Unified file system to be scanned as a CIFS or NFS export.
2. Mount the export on a host that has a Storage Resource agent file system that is deployed on it. Exports are accessible to the TCP/IP data network through the Storwize V7000 Unified interface node **public IPs**. Exports must be mounted as a CIFS export on Windows agents or an NFS export on UNIX agents. On Windows operating systems, the export must be mounted as the local system user or from a service that is running under the local system account so that it is visible to all logon sessions. One way of mapping a drive as the local system user is to map the drive from a **psexec** command window. The **psexec** command window is available as part of the **Sysinternals** command-line tools for Windows operating systems.
3. Probe the Storwize V7000 Unified and the Storage Resource agent. The default behavior of the probe is to discard export paths that are not at the root of the file system. The export paths are discarded to prevent the product from counting information twice on a file system. The reason for this behavior is because the probe is treating two exports to the same file system as separate file systems. This behavior can be overridden by setting the **saveNonRoot** parameter. The probe picks up all exports (shares) of a specific NAS device as file systems. To set the **saveNonRoot** parameter, insert the line **saveNonRoot=1** in the server section of the `TPCD.config` file and restart the Data server to take effect. The `TPCD.config` file is in the *TPC_installation_directory*`/data/config` directory where *TPC_installation_directory* is where Tivoli Storage Productivity Center is installed.
4. Assign a scan agent to Storwize V7000 Unified file system.
5. Open the stand-alone GUI and click **Administrative Services** > **Configuration** > **Scan/Probe Agent Administration**.
6. Save the page.
7. A scan job for the Storwize V7000 Unified can be created and run from **Data Manager** > **Monitoring** > **Scans**.

Editing scans

Change the settings in an existing scan job.

To edit an existing scan, complete the following steps:

1. In the stand-alone GUI, expand **Data Manager** > **Monitoring** > **Scans**. For Data Manager for Databases, expand **Data Manager for Databases** > **Monitoring** > **Scans**.
2. Click the name of the scan you want to edit. The content pane displays information about that job.
3. Edit the scan definition by changing the values that appear on the following pages:
   - Filesystems. Use this page to edit the computers and file systems whose storage usage and consumption you want to monitor as part of a scan.
   - Directory Groups. Use this page to edit the groups of directories that you want to scan.
   - Databases-Tablespaces. Use this page to edit the databases, tablespaces, instances, and computers whose storage usage and consumption you want to monitor as part of a scan.
   - Table Groups. Use this page to edit the table groups that you want to scan. A table group represents a group of defined tables.
   - Profiles. Use this page to edit which profiles are used during a scan. Using profiles is a powerful method for determining the statistics you want to gather during scans.
   - When to Run. Use this page to edit when a scan is run.
   - Options (Oracle and UDB only). Use this page to edit the scan options for Oracle or UDB.
   - Alert. Use this page to edit how you are notified if a scan fails.
4. Click to save the updated scan definition.

Use profiles

In the stand-alone GUI, use profiles to specify what statistical information is gathered and to fine-tune and control what files are scanned during a scan. You can also use profiles with some reports to determine what statistics are displayed.

You can use the default profiles that are supplied with IBM Tivoli Storage Productivity Center until you begin creating your own profiles. With each of these default profiles, you can select a specific statistic that you want to gather. You might want to select ALL of the default profiles when you create your scan job. This selection scans a computer or file system and gathers all of the statistical information that is available.

For more information about default profiles, see “Default profiles”.

Default profiles:

Default Profiles are provided for your use with IBM Tivoli Storage Productivity Center. Each default profile represents a specific statistic that you can gather (which is included in the name of the profile for convenient identification).

To scan a computer or file system and gather all of the statistical information available, select all the default profiles when you create your scan in the stand-alone GUI. TPCUser is shown as the creator of the default profiles, and TPCUser is the prefix name for each of the profiles. The following list identifies the default Data Manager profiles that are available and describes their functions.

By Access
: Gathers statistics by length of time since last access of files.

By Creation
Gather storage statistics by the length of time since the creation of files. This profile applies to scans run against Windows computers only. Scans with this profile do not gather any creation time data from computers that are not a Windows system.

By Mod Not Backed Up
Gathers statistics by length of time since last modification, only for files that are not backed up since modification (Windows operating systems only).

By Modification
Gathers statistics by length of time since last modification of files.

Size Distribution
Gathers information on the size distribution of files.

Largest Directories
Gathers statistics on the *n* largest directories (20 is the default amount).

Largest Files
Gathers statistics on the *n* largest files (20 is the default amount).

Largest Orphans
Gathers statistics on the *n* largest orphan files (20 is the default amount).

Most At Risk
Windows operating systems only: Gathers statistics on the *n* files that are modified the longest time and are not yet backed up since they were modified (20 is the default amount).

Oldest Orphans
Gathers statistics on the *n* oldest orphan files (20 is the default amount).

Most Obsolete Files
Gathers statistics on the *n* most obsolete files. These files are not yet accessed or modified for the longest time (20 is the default amount).

Summary By File Type
Summarizes space usage by file types. Examples of these file types are .exe, .dll, .doc, .mp3, and so forth.

Summary By File System/Directory
Summarizes space usage by file system or directory.

Summary By Group
Summarizes space usage by OS Group.

Summary By Owner
Summarizes space usage by Owner.

Temporary Files
Gathers statistics on the non-OS files that are not accessed in the last year and orphan files.

Wasted Space
Gathers statistics on the non-OS files that are not accessed in the last year and orphaned files.

The following list identifies the default Data Manager for Databases profiles and describes their functions.

DB User Space
Gathers statistics about the user space usage within a database or table space.

Largest Tables
Gathers statistics on the *n* largest tables (20 is the default amount).

Segment Most Extents
Gathers statistics on the *n* segments with the most extents (20 is the default amount).

Segment Most Unused Space
Gathers statistics on the *n* segments with the most empty used space (20 is the default amount). Empty used space represents the space that is allocated to a segment that is empty/not being used. Gathering this information can help you discover space that can be reclaimed and allocated to other objects. Specifically, this Profile is useful:

- To find objects that are over-allocated (that does not use all their associated free space)
- To find the free space that exists in excess of the high-water mark for a table. This mark is the highest block that was written for a table.
- To find segments that have the largest number of blocks that are not filled

Summary
Summarizes space usage by database or table space. The summary includes how many segments are in a table space (Oracle only), how many of those segments are tables (Oracle only), indexes, and so forth.

Creating a profile:

Create a profile for use within a scan.

To create a profile perform the following steps:

1. In the stand-alone GUI, expand **Data Manager** > **Monitoring**. For Data Manager for Databases, expand **Data Manager for Databases** > **Monitoring**.
2. Perform one of the following actions:
   - Select **Create Profile** from the pop-up menu.
   - Click **New Profile** from the **Profiles** tab of the Create scan window.

   Use the Create Profile window to specify exactly what level of detail to collect when running a scan. The more information that you specify, the more reports that you will be able to run against the information that you collect during the scan. Statistics are tied to a specific profile.
3. The **Statistics** tab on the **Create Profile** window is displayed.
4. Select statistics you want the scan to collect. Statistics are tied to a specific profile.

   See “Default profiles” on page 315 for more information about statistics.
5. If you want to further limit specific files to be scanned, select the **File Filters** tab. You can define exclusive conditions that filter what files are analyzed during the scan.

   See the “Creating file filters” on page 320 section for more information on using this tab.
6. The first time you view a new profile, this window is blank except for the text **All files selected**. That is the default and indicates that information will be

collected for all files. To create your first group or condition, simply right-click **All files selected** and choose to create a new group or condition.

7. In the **File Filters** page, the default condition **All files selected** is displayed. Right-click **All files selected** and select **New Condition** from the pop-up menu. You can also specify groups for these conditions, by selecting **New Group**.

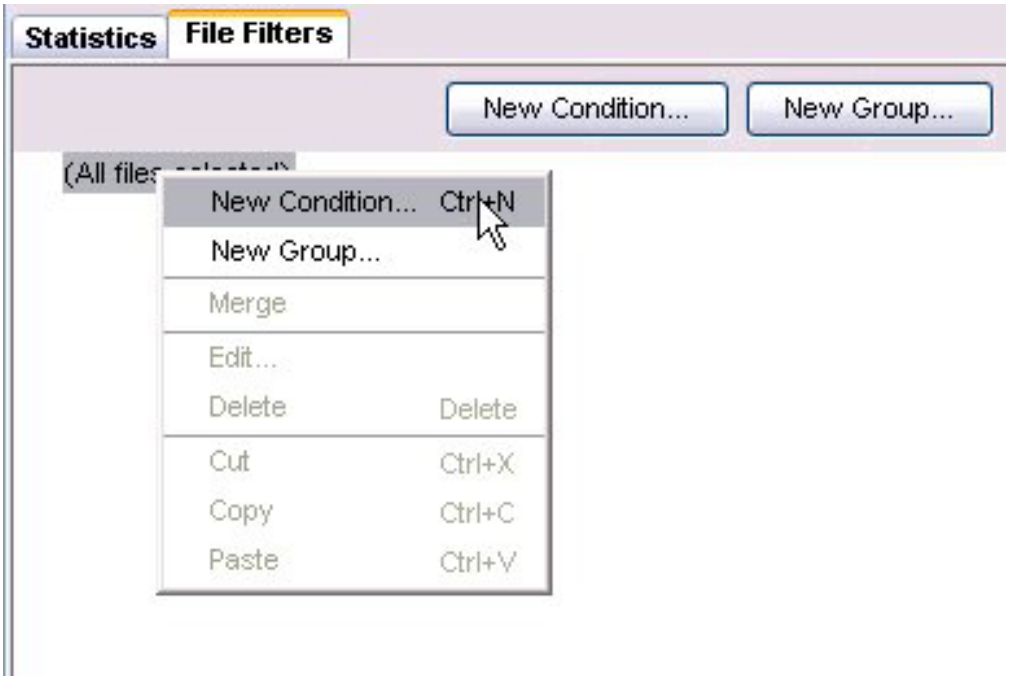

Figure 14. File Filters tab, pop-up menu

8. In the preceding figure, **New Condition** is selected. The **Create Condition** window opens, enabling you to define your file conditions. You can create conditions based on access time, creation time, file name or group, modification time, file attributes, size, type, owner, or length.

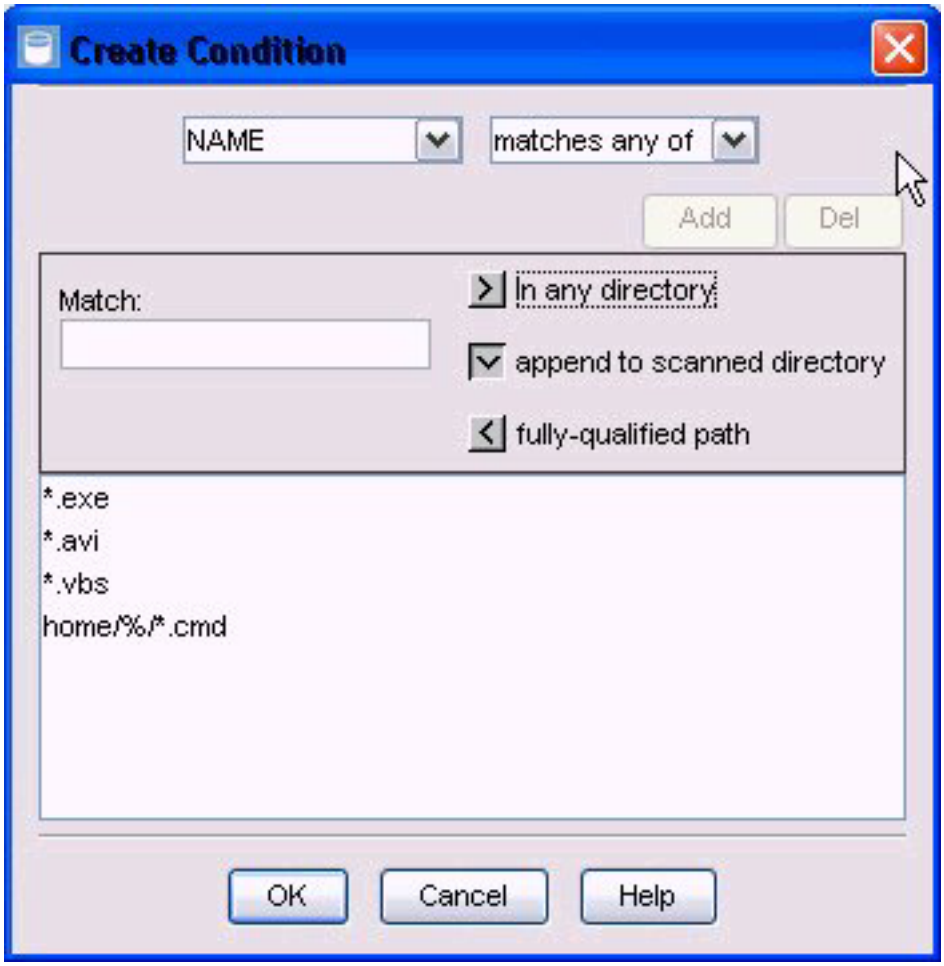

Figure 15. Create Condition window

9. Press **OK** after defining each condition.
10. If you are defining more than one condition, the **Create Group** dialog box displays to help you associate the conditions you created:

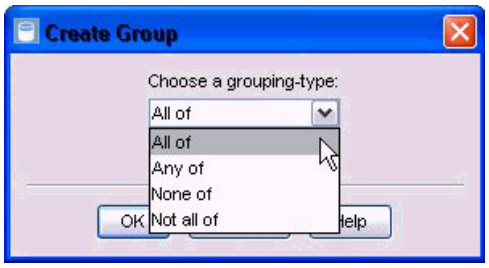

Figure 16. Create Group Window

11. Select how you want to associate the conditions you have defined. See "Associating conditions with groups" on page 325 for more information on creating groups.
12. The condition is displayed in the **File Filter** window.

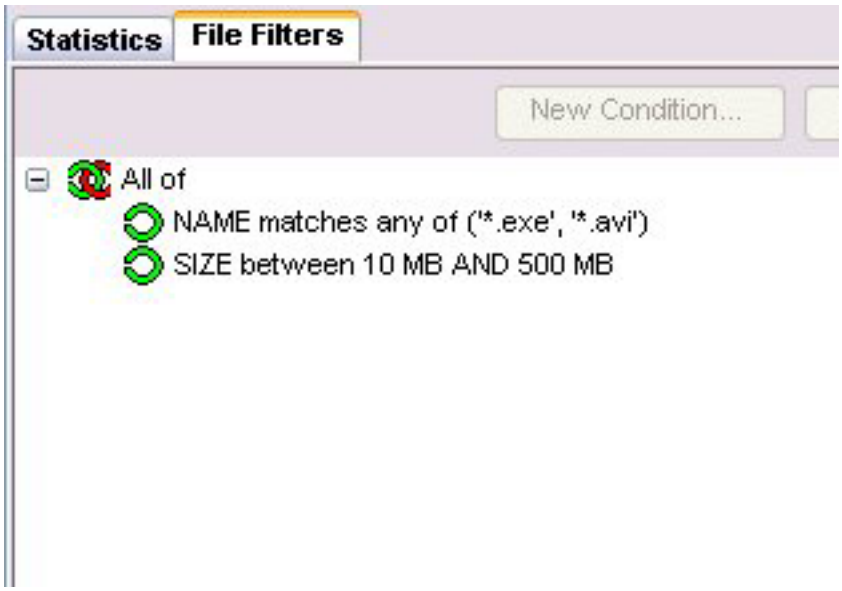

Figure 17. File Filter window

13. Click to save the profile.
14. Enter a name for the profile and click **OK**. An entry for the profile appears as a node under the **Profiles** node.

Creating file filters:

Use file filters as part of a profile to further define the files that you want to include or exclude in a scan. Filters limit the files on which you are gathering data during the scan based on criteria that you define.

In the following example, a group (`All of`) has been created with three conditions:

- The first controls the type of files to collect information on (MP3, GIF, AVI, and JPG files).
- The second specifies that the files must have been created within the last year.
- The third limits the profile to analyzing only files greater than 500 KB.

To create file filters within profile, complete the following steps in the stand-alone GUI:

1. To edit the filer filters in an existing profile, expand **Data Manager** > **Monitoring** > **Profiles** and select an existing profile in the navigation tree. To create a new profile and add file filters, expand **Data Manager** > **Monitoring**, right-click **Profiles**, and select **Create Profile**.
2. Click the **File Filters** tab.
3. The first time you view a new profile, this window is blank except for **All files selected**. Selecting all files is the default and indicates that information is collected for all files. To create your first group or condition, right-click **All files selected** and choose to create a new group or condition.
4. In the File Filters page, the default condition **All files selected** is displayed. Right-click **All files selected** and select **New Condition**. You can also specify groups for these conditions by selecting **New Group**.

The **Create Condition** window opens, enabling you to define your file conditions. You can create conditions based on access time, creation time, file name or group, modification time, file attributes, size, type, owner, or length.

5. Press **OK** after defining each condition. See "Defining a filter condition" for more information on creating groups.

   Each condition is displayed in the **File Filter** window.

   If you are defining more than one condition, the **Create Group** dialog box helps you associate the conditions you create.

6. Associate the conditions you have defined.

   See "Associating conditions with groups" on page 325 for more information on creating groups.

Defining a filter condition:

Define a condition for a filter.

Use the stand-alone GUI to define a filter condition. On the File Filters page, right-click **All files selected** and select **New Condition**. The Create Condition window opens.

The first drop down list displays file criteria that you can select to define your condition. You can define file criteria based on access time, creation time, file name or group, modification time, file attributes, size, type, owner, or length. Depending on the criteria you choose to define, the second drop down list and window fields change dynamically to help you set parameters for the selected criteria. The following table describes the file criteria you can select for your file filter.

Condition	Description	Method for filtering the files	Condition Modifiers
Last Accessed	Use this condition to limit files based on the last time they were accessed. For example, you can define a filter that includes files that were last accessed between 12:01 on January 1, 2005 and 11:59 on January 30, 2005.	• later than • earlier than • between • not between	• Length of time before scan: enter the years, weeks, days, hours, or minutes before a scan • Timestamp - select files based on the timestamp
Created	Use this condition to limit files based on when they were created. For example, you can define a filter that includes files that were created between 12:01 on January 1, 2005 and 11:59 on January 30, 2005.	• later than • earlier than • between • not between	• Length of time before scan: enter the years, weeks, days, hours, or minutes before a scan • Timestamp: select files based on the timestamp

Condition	Description	Method for filtering the files	Condition Modifiers
Name	Use this condition to limit files based on their name. For example, you can define a filters that include files that end in *.html, or include files that end in .cmd and appear within a specified subdirectory.	• matches any of • matches none of • matches • does not match	Enter the file name upon which you want the condition to be based. For example, you can specify, *.exe, *.mp3, *.avi. in the **Match** field. See "Using wildcards in file filters" on page 326 for information about how to use wildcards (*, %, etc.) in this field. Select the location where the file name will be searched during the scan: • **In any directory**. During a scan, Data Manager will look for the file name in any directory. In this condition, the search string will occur at the end of the full directory path. • **Append to scanned directory**. This option enables you to restrict what directory to scan. Data Manager will look for the file type in the specified file system root directory or in a monitored directory in a directory group tied to a scan. In this option, you must first specify what directory groups you want scanned in the **Directory Groups** tab of the **Scan** window. See "Selecting directory groups to scan" on page 309 for more information. • **Fully-qualified path**. Enter the fully-qualified path when you enter the file type in the **Match** field. During a scan, Data Manager will look for the file type in the exact path you specify here. This is useful when you want to exclude system directories and files in a Most Obsolete Files report.
Group	Use this condition to limit files based on the OS group to which those files belong.	• matches any of • matches none of • between • not between • matches • does not match • >= • < • > • <=	Name of the OS group
Last Modified	Use this condition to limit files based on when they were last modified.	• later than • earlier than • between • not between	Specify the time as: • Length of time before scan: enter the years, weeks, days, hours, or minutes before a scan • Timestamp: select files based on the timestamp

Condition	Description	Method for filtering the files	Condition Modifiers
Attributes	Use this condition to limit files based on file attributes such as read only, system, temporary, etc.	• include any of • include none of • include all of • don't include all of	Select the file attribute upon which you want the condition to be based. You can select: • Read only: Users cannot write to the file. • Hidden: A file that cannot be seen by users during a normal directory listing. • System: A file used by the operating system that cannot normally be viewed by users. • Archive: A file that has been backed-up and the archive bit is turned off. An Archive bit signifies that the file has been changed since it was last archived. • Normal: A file with no special attributes set to it. • Temporary: (Windows only) A temporary or backup copy of a file. • Sparse: A file with a special kind of compression that does not take up a lot of CPU time. The logical size of the file looks bigger than is actually allocated. It contains some blank pages which are not on disk. • Reparse Point: New NTFS file system objects that have a definable attribute containing user-controlled data and are used to extend functionality in the input/output (I/O) Subsystem. • Compressed: A file to which a compression algorithm was applied. This file takes up less space on disk, but it will be slower to read. • Offline: A file that was migrated to tape. • Not Content Indexed: A file that is not indexed--the Windows indexing service was told that this is not a file that should be searched through for text strings. • Encrypted: A file that was saved with a password (is password-protected). • Orphaned: (for Windows, UNIX, and Linux) The user who owns the file is no longer present in the system database.

Condition	Description	Method for filtering the files	Condition Modifiers
Attributes (continued)			• Block Device: (UNIX/Linux only) A disk device file (not a real file). Executable: A file that you can run or execute. Sticky: A special type of executable file that is read into memory when it is first run and never released. Use Scans to find these types of files and set them to "unsticky". SETGID: (UNIX/Linux only, for executable files) When a file is run, its group changes to the group of the file. For example, if a user belongs to User group and runs a file from a different group (e.g., Root), than the user's group will be changed to Root while running the file (for security purposes). SETUID: (UNIX/Linux only, for executable files) When a file is run, the user's effective user will be changed when the file is run. For example, if a file is SETUID =Root, a users effective user ID becomes Root when running the file. This is considered a security hole by UNIX Administrators, so you can define a Profile to search for files that have a SETUID = Root.
Size	Use this condition to limit files based on their physical file size.	• between • not between • >= • < • > • <=	• KB (kilobytes) • MB (megabytes) • GB (gigabytes) • TB (terabytes)
Type	Use this condition to limit files based on pre-defined file types.	• in • not in • = • < • >	• Link • Directory • Fifo • Socket • Device
Owner	Use this condition to limit files based on their owner.	• matches any of • matches none of • between • not between • matches • does not match • >= • < • > • <=	Enter the name of the owner upon which you want the condition to be based in the **Entry** field.

Condition	Description	Method for filtering the files	Condition Modifiers
Length	Use this condition to limit files based on the logical size of the files contents.	• between • not between • >= • < • > • <=	Enter the file length range. You can qualify both values by selecting from one of the following: • KB (kilobytes) • MB (megabytes) • GB (gigabytes) • TB (terabytes)

Associating conditions with groups:

Associate multiple conditions in a group or groups.

Groups can further define your file filter by including or excluding certain files based on specific criteria.

1. In the stand-alone GUI, create a new group using one of the following options.
   - In the **File Filters** tab, right-click a condition and select **New Group** from the pop-up menu.

     The **Create Group** dialog box is displayed.
   - Create each file condition for your file filter.

     As you create a new file condition, the **Create Group** dialog box is displayed.
2. Select the grouping for the file conditions you created. The grouped condition is displayed in the **File Filters** tab.

Select one of the following groups.

Group	Description
All of	All inclusive. All the conditions must be met. For example: • a *and* b *and* c
Any of	Any one or more of the following conditions will be met. For example: • a *or* b *or* c
None of	None of the conditions are met. For example: • *not* (a *or* b *or* c)
Not all of	All of the following conditions are not met. For example: • *not* (a *and* b *and* c)

For example, say you want to create a file filter with the following conditions:

- Name in (*.gif, *.jpg)
- Size > 1 MB
- Attribute is Read Only

How you group the conditions produces different results. The following table shows the different results depending on the group you select:

File	Group	Result
A.gif 500KB read only	• All of • Any of • None of • Not all of	• False • True • False • True
B.jpg 2 MB read only	• All of • Any of • None of • Not all of	• True • True • False • False
C.gif 2 MB read/write	• All of • Any of • None of • Not all of	• False • True • False • True
D.txt 500KB read/write	• All of • Any of • None of • Not all of	• False • False • True • True

Using wildcards in file filters:

In the stand-alone GUI, use wildcard characters within the conditions of file filters to perform pattern matching of file names. A wildcard is a special symbol that represents one or more characters.

You can use wildcards in File Filter conditions to identify multiple files and directories with a single specification. The following table describes the wildcards that are available in Data Manager:

Wildcard Character	Description	Examples
*	Use the asterisk (*) to represent unknown characters in file names. It can act as a wildcard for zero or more characters of any value within a directory or file name.	• **s*** refers to all files that begin with **s** • ***.exe** refers to all files that have **.exe** as an extension • **/user/*** refers to all files under the **/user** directory
%	Use the percent (%) character to specify a wildcard for a directories within the path of a file. It can act as a wildcard for zero or more directories within a path.	• **/home/%/*.cmd** will match all files with the extension **.cmd** and located in any of the subdirectories under the **/home** directory • **/usr/%/*** will match all files and subdirectories under the **/usr** directory (for example, /usr/files/c, /usr/c, /usr/files/documents/c)

Wildcard Character	Description	Examples
?	Use the question mark (?) to represent a character within a file name or directory name. It can as a wildcard for a character within a directory or file name.	• **/home/usr?/** will match directories like: /home/usr1/, /home/usr2/, etc. • **/usr/temp/temp??.tmp** will match files like /usr/temp/temp01.tmp, /usr/temp/temp02.tmp, etc.

Data Manager supports standard UNIX shell path name expansion wildcarding within file filter conditions. See the appropriate UNIX man pages for more information about the wildcards that you can use.

Determining the availability of storage resources

Use pings to collect information about the availability of the storage resources in your environment. To create pings, use the stand-alone GUI. You can view the information that is collected by pings in Availability and System reports.

Pings enable you to monitor and report on the availability of your storage from a network point of view or from a computer uptime perspective. See the percentage of off-network time due to network problems or system downtime. Define any number of pings that each check the availability of a unique storage resource.

Creating pings

To create pings, use the stand-alone GUI. Create pings to collect information about the availability of the storage resources in your environment. You can define any number of pings to gather information about different storage resources at different times.

To create a ping, complete the following steps:

1. In the stand-alone GUI, expand **Data Manager** > **Monitoring** > **Pings**.
2. Right-click **Pings** A menu is displayed.
3. Select **Create Ping**.
4. Select the computers whose availability you want to check. “Select computers to ping” on page 328.
5. Schedule when you want a ping to run. “When to run the ping” on page 328.
6. Specify how often you want the server to write out the ping results to the database repository for viewing in reports. “Specify ping options” on page 329.
7. Click to save the ping. You must save the ping before Tivoli Storage Productivity Center can collect data about storage resources.
8. Enter a name for the ping and click **OK**. An entry for the ping is displayed under the **Pings** node in the navigation tree. The ping is run at the date and time specified in the schedule definition.
9. Use the Job Management panel in the stand-alone GUI to view information about the runs and jobs that are related to the ping. For example, use the Job Management panel to view the status of a run for a schedule, the date and time when it was run, and the log files of the jobs associated with the run.

Related reference:

“Managing schedules, jobs, and runs” on page 605
Use the Job Management panel in the stand-alone GUI to view and manage the schedules, runs, and jobs that are related to the resources that are monitored by

Tivoli Storage Productivity Center.

Select computers to ping:

In the stand-alone GUI, select the computers and clusters that you want to ping.

Use the **Computers** page in the stand-alone GUI to select the computers and clusters whose availability statistics you want to collect. The **Available** list box shows the computers, computer groups, or clusters that are currently being monitored and available for inclusion in the ping. After you include storage resources in the ping, it is removed from this list and moved to the **Current Selections** list box. The **Current Selections** list box shows the computers, computer groups, or clusters that have been selected for the ping.

Note: A cluster name is displayed next to individual computers that are members of a IBM PowerHA SystemMirror for AIX or Microsoft Cluster Server cluster. Cluster resource groups are identified with a special icon. Clusters displayed in the **Available** list box cannot be expanded to show member nodes. You can move a cluster to the **Current Selections** list box to view and, if necessary, exclude member nodes.

You can include individual computers, computer groups, or clusters against which to run a ping. See "Working with groups of resources" on page 337 for information about groups.

- **To include storage resources in a ping:**
  1. In the **Available** list box, highlight the computers, computer groups, and clusters that you want to include in the ping.
  2. Click . The storage resource that you selected is moved to the **Current Selections** list box.
- **To remove storage resources from a ping:**
  - Highlight the computers, computer groups, and clusters that you want to remove and click . The storage resource is removed from the ping and moved from the **Current Selections** list box to the **Available** list box.
- **To exclude one or more computers from a selected computer group or cluster:**
  1. Double click the group or name of the cluster in the **Current Selections** list box.
  2. Right-click the computers that you want to exclude. A menu is displayed.
  3. Select **Exclude**. An **Excludes** section is created in the **Current Selections** list box that displays the names of the computers that you excluded from the ping.

When to run the ping:

Use the stand-alone GUI to schedule when and how often to run a ping.

You have the following options for determining when to run a ping:

- Run the job immediately
- Run the job once at a specified time and date. Use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM.
- Run the job repeatedly according to a schedule you specify. You can do the following:

- Use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM when the job should begin running.
- Indicate how often the job should run. Use the list boxes to specify the number of minutes, hours, days, weeks, or months.
- Select which days you wish to run the job.

Use the fields in the **How to handle timezones** section to indicate the time zone that Data Manager should use when scheduling and running the ping. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:

- 9:00 AM of the time zone where the server resides
- 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent that is located in the PST time zone.

Use the timezone that the server runs in. Select this option to use the time zone of the location where the Data server resides.

Use this timezone. Select this option to indicate a specific time zone for when to run the job. Select the timezone from the list provided.

Specify ping options:

To specify how often the results of pings are written from memory and stored in the database repository, use the stand-alone GUI .

For greater efficiency, the server temporarily stores ping results in memory. The unit of time you specify in the **Options** tab determines how often these results are externalized to the repository.

You can specify a unit of time in a number of pings, minutes, hours, or days.

Define alert values for a ping:

To set an alert that will be triggered if the ping job fails, use the stand-alone GUI.

Use the Alert page in the stand-alone GUI to define an alert to notify you if a computer that a ping is trying to check is unreachable for a certain amount of time. You can choose to be notified by SNMP trap, IBM Tivoli Enterprise Console event, login notification, Windows event log, UNIX Syslog, running a script, and e-mail. If you select SNMP trap, Tivoli Enterprise Console event, or e-mail, you must also define a destination (for example, an e-mail address). Define this destination by expanding **Administrative Services** > **Configuration** > **Alert Dispositions**.

To define an alert for a ping, complete the following steps:

1. Click the **Alerts** tab; the Alert page opens.
2. Define the conditions that will trigger the alert. The default setting is **Computer Is Unreachable More Than**. No other alerting conditions are currently available. Enter the **Value** and **Value Units**. These values specify how many pings occur before an error is generated. By default, **Value Units** is set to **Times**.
3. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

SNMP trap
An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

TEC / OMNIbus Event
A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the `tivoliSRM.baroc` file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, or running programs.

Login notification
The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who receives alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who receives alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that is recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog is sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters that are listed in the **Specify Script** window are passed to the script when the alert is triggered.

Email An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit email** to change the actual email message.

4. Click **File** > **Save** to save the ping.
5. Enter a name for the ping and click **OK**. An entry for the ping appears under the **Pings** node.

Editing pings

To change the settings in an existing ping job, use the stand-alone GUI.

To edit an existing ping, complete the following steps:

1. In the stand-alone GUI, expand **Data Manager** > **Monitoring** > **Pings**.
2. Click the name of the ping you want to edit. The content pane displays information about that job.
3. Edit the ping definition by changing the values that appear on the following pages:

- Computers. Use this page to edit the computers, computer groups, and clusters whose availability you want to monitor as part of a ping.
- When to Run. Use this page to edit when a ping is run.
- Options Use this page to edit how often the results of pings are written from memory and stored in the database repository.
- Alert. Use this page to edit how long a storage entity is unreachable before you are notified that the ping has failed.

4. Click to save the updated ping definition.

Collect storage statistics from Tivoli Storage Productivity Center servers

Schedule probes from the master server for Tivoli Storage Productivity Center to collect storage information that is gathered by subordinate servers. The information that is gathered by probes of subordinate servers is displayed in the **Tivoli Storage Productivity Center** > **Rollup Reports** node in the navigation tree of the master server. Create and manage server probes by selecting **TPC Server Probe** in the navigation tree of the stand-alone GUI.

A subordinate server monitors and reports on storage resources like a standard server. It also communicates with the master server during TPC server probes. During these probes, the master server collects the storage information that is gathered by a subordinate server 's agents and stores that information in its own database repository.

Note: If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

Before you run TPC server probes:

1. Associate subordinate servers with the master server on the **Administrative Services** > **Data Sources** > **TPC Servers** window, in the stand-alone GUI. See "Adding a Tivoli Storage Productivity Center server as a subordinate server" on page 247 for information on how to associate subordinate servers with a master server.
2. Run discoveries, probes, and scans on the subordinate servers to gather information about their locally managed entities.

The information that is collected by TPC server probes from subordinate servers includes information:

- Asset information: detailed statistics about agents (Tivoli Storage Productivity Center), computers, storage subsystems, disk and volume groups, disks, file systems, logical volumes, volumes, and fabrics.
- Database asset information: detailed statistics about the RDBMSs in your environment, including Oracle, SQL Server, Sybase, and UDB.
- Capacity information: storage metrics that are related to the disk capacity, file system capacity, file system used space, and file system free space of the storage entities in your environment.
- Database capacity information: storage metrics that are related to the storage capacity of the RDBMSs in your environment, including Oracle, SQL Server, Sybase, and UDB.

Extra storage information that is collected by a subordinate server and is available for only that server. The storage information is used within Tivoli Storage Productivity Center functions such as the topology viewer, data path explorer, and volume performance. That information is not collected by the master server, and is not rolled up into the reports available through the master server.

Also, the Tivoli Storage Productivity Center functions within the master server use the storage information that is collected by only the master server.

Creating Tivoli Storage Productivity Center server probes

Use Tivoli Storage Productivity Center server probes from the master server to collect storage information that has been gathered by subordinate servers.

To create a Tivoli Storage Productivity Center server probe, complete the following steps:

1. In the stand-alone GUI, expand **IBM Tivoli Storage Productivity Center** > **Monitoring** > **TPC Server Probes**.
2. Right-click **TPC Server Probes** and select **Create TPC Server Probes** from the menu.
3. Enter a brief description of the probe in the **Description** field.
4. In the **What to Probe** page, complete the following steps:
   a. In the **Available** list box, highlight the Tivoli Storage Productivity Center servers you want to probe. You can do any of the following: 1. Select the **All TPC Servers** item to automatically include all subordinate servers in the probe. 2. Select the name of a specific subordinate server and all its related storage entities that you want to probe. 3. Select the specific entities on a subordinate server that you want to probe. For example, you can expand the node for a subordinate server and select the specific entities for which you want to collect data: All Resources, Clusters, Computer Groups, Computers (including Hypervisors), Database Instances, Fabrics, RDBMS Computer Groups, Storage Subsystem Groups, and Storage Subsystems.
   b. Click

to include highlighted Tivoli Storage Productivity Center servers or storage entities in the probe. The entities that you selected are moved to **Current Selections** list box.
5. In the **When to Run** page, schedule when you want the probe to run. You have the following options:
   - Run the job immediately
   - Run the job once at a specified time and date. Use the list boxes to specify the month, day year, and the hour, minute, and AM/PM.
   - Run the job repeatedly according to a schedule you specify. You can perform the following tasks:
     – Specify the month, day, year, and the hour, minute, and AM or PM when the job is to run.
     – Indicate how often the job runs. Use the list boxes to specify the number of minutes, hours, days, weeks, or months.
     – Indicate which days on which run the job. A check mark is displayed next to the days you select.

   Use the fields in the **How to handle timezones** section to indicate the time zone that Data Manager for Databases uses when scheduling and running the action. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
   - 9:00 AM of the time zone where the server resides

- 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on a Storage Resource agent located in the PST time zone.

Use the timezone that the server runs in
Select this option to use the time zone of the location where the Data server resides.

Use this timezone
Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.

6. In the **Alert** page, set the following options:
   - The condition that triggers an alert. The default setting is **Probe Failed**. No other conditions are currently available.
   - The notification method of an alert: SNMP Trap, IBM Tivoli Enterprise Console Event, Login Notification, Windows Event Log, UNIX Syslog, running a script, and email. If you select SNMP Trap, Tivoli Enterprise Console Event, or email for the method by which you are notified of an alert, you must define a destination for the method on the **Administrative Services → Configuration → Alert Dispositions** window.
7. Click to save the probe. You must save the probe before Tivoli Storage Productivity Center can collect data about storage resources.
8. Enter a name for the probe and click **OK**. An entry for the probe is displayed under the **TPC Server Probes** node in the navigation tree. The probe is run at the date and time specified in the schedule definition.
9. Use the Job Management panel to manage view information about the runs and jobs that are related to the probe. For example, use the Job Management panel to view the status of a run for a schedule, the date and time when it was run, and the log files of the jobs associated with the run.

Related reference:

"Managing schedules, jobs, and runs" on page 605
Use the Job Management panel in the stand-alone GUI to view and manage the schedules, runs, and jobs that are related to the resources that are monitored by Tivoli Storage Productivity Center.

Editing TPC Server probes

Learn how to edit the definition of a TPC server probe.

To edit an existing TPC server probe, complete the following steps:

1. In the stand-alone GUI, expand **IBM Tivoli Storage Productivity Center** > **Monitoring** > **TPC Server Probes**.
2. Click the name of the TPC server probe you want to edit. The content pane displays information about that probe.
3. Edit the values on the What to PROBE, When to Run, and Alert pages to change the TPC server probe definition.
4. Click to save the updated TPC server probe definition.

Managing the runs and jobs for TPC Server Probes

To view information about the runs, jobs, and log files for TPC Server Probes, use the stand-alone GUI.

Tivoli Storage Productivity Center creates a run each time a TPC Server Probe schedule collects data from subordinate servers. A job is a unit of work within a run for a schedule. Use the Job Management panel in the stand-alone GUI to manage TPC Server Probe schedules and view information about the related runs and jobs. Specifically, use this panel to complete the following tasks for schedules:

- View or edit the definition of a schedule.
- Run a schedule immediately.
- Rename or delete a schedule.
- View basic information about a schedule, including the ID of the user who created it, description, the last time the schedule was run, and the next time the schedule is run.

Use the Job Management panel to complete the following tasks for the runs and jobs that are related to TPC Server Probe schedules:

- View the status of a run and the date and time when it was executed.
- View information about the jobs that are associated with each run, including the name of the computer on which a job was run, the type of the monitored device, and the status of the job.
- View the log files for one or more jobs at a time.
- Display jobs that have started or completed within a specified length of time. For example, display jobs that have completed in the last two weeks.
- Display jobs that have a certain status, such as Success, Running, or Failed.

To access the Job Management panel for a TPC Server Probe schedule, complete the following steps:

1. In the navigation tree, expand **IBM Tivoli Storage Productivity Center** > **Monitoring** > **TPC Server Probes**.
2. Select the probe schedule whose runs and jobs you want to manage. The content pane displays information about the schedule.
3. In the content pane, click **View Job History**. The Job Management panel is displayed and the selected schedule is highlighted in the **Schedules** section. View the entries in the **Jobs for Selected Schedules** section to manage the jobs and runs for the schedule.

You can also access the Job Management panel at any time by clicking the related icon ( ) in the toolbar. Use the **Filters** section on the Job Management panel to display the TPC Server Probe schedule that you want to manage. See “Managing schedules, jobs, and runs” on page 605 for more information.

Data collection groups

To create data collection groups, use the stand-alone GUI. With groups, you can define a set of related resources such as a computer, fabric, or storage subsystem one time. You can save these selections as a named group for future use in scans and pings to make data collection more efficient. For example, you can create a group that contains a number of related computers. You can then reference that group within a monitoring job when you want to collect data about all those computers at the same time.

In the navigation tree of the stand-alone GUI, each type of group is represented by a node under ***<manager name>*** > **Monitoring** > **Groups**. You can expand the tree for each type of group to see what groups exist for that type. Resources can belong to only one group at a time. For example, a computer already belongs to a computer group and you attempt to add it to a different computer group. The

system removes it from the first computer group before you add it to the new computer group. This process includes removing it from the default group if it is not previously assigned to a user-defined group. This action is to prevent a resource from being counted twice statistically.

Additionally, you can use Data Manager for Chargeback to create invoices for these groups. When you implement your chargeback system, it is advisable that you group the IT resources within your organization into functional groups. Accurately identifying the types and locations of your stored data and resources is the first step in implementing an effective and equitable chargeback policy. When you define groups with Data Manager, you can classify your data and resources into billable entities. This action provides information on the users and what storage they use. You can bill specific groups according to their usage of storage space.

The following table lists and describes the types of groups that you can create.

Group Type	Description
	Data Manager
Computer	A logical grouping of related computers. A computer can belong to one computer group only. A computer is automatically removed from a group if you add it to another group. You can use computer groups to monitor computers that are related to each other or belong to a common group as a department, location, territory, and so on. Examples of computer groups include Sales, Accounting, Development, IT, West, East, Midwest, NAS Filers, Remote Servers
File system	A grouping of file systems across multiple computers. A file system can belong to one file system group only. A file system is automatically removed from a group if you add it to another group. Here is an example of file system groups: `C:\`, `D:\`, `/boot`, `/home`, `/usr`
Directory	A grouping of directories across multiple computers. You can include a directory in multiple directory groups. Examples of directory groups: `WindowsTempDir` Included Computer Name:Directory - `a:\machine1\winnt\temp, b:\machine2\winnt\temp, c:\machine3\winnt\temp` **Note:** To monitor directories from the IBM Tivoli Storage SAN File Systems, you must create a directory group and add the wanted directories as follows: *`/<SAN File System cluster name>/<directory name>`* where *`<SAN File System cluster name>`* represents the name of the SAN File System cluster that contains the directory that you want to add and `<directory name>` represents the name of the directory you want to add (for example, `/IBM/myFilesetName/myDirectoryName`, `/IBM/myDirectoryName`.

Group Type	Description
User	A grouping of related users. A user can belong to one user group only. A user is automatically removed from a group if you add it to another group. Users that were discovered by Tivoli Storage Productivity Center can be included in a user group. **Note:** The user groups that you define in Data Manager for Databases will also available in Data Manager, and vice versa. Examples of user groups include Administrators, Developers, Marketing staff, HR staff
OS User Group	A grouping of OS user groups that were defined outside of Tivoli Storage Productivity Center at the operating system level. OS user groups are the operating system groups that Tivoli Storage Productivity Center finds that own files and directories in the environment. Examples of OS user groups include Windows, Linux, UNIX
	Data Manager for Databases
Computer	A logical grouping of related computers. A computer can belong to one computer group only. A computer is automatically removed from a group if you add it to another group. You can use computer groups to monitor computers that are related to each other or belong to a common group as a department, location, territory, and so on. Examples of computer groups include Sales, Accounting, Development, IT, West, East, Midwest, NAS Filers, Remote Servers
Databases-Table spaces	A grouping of RDBMS databases or table spaces. You can use database-tablespace groups to monitor entire databases or table spaces (depending on the RDBMS) across multiple instances. A database or table space is automatically removed from a group if you add it to another group. A database or table space can belong to one database-tablespace group only. Examples of database-tablespace groups include DB2 databases, system table spaces
Table	A grouping of related RDBMS tables.
User	A grouping of related users. A user can belong to one user group only. A user is automatically removed from a group if you add it to another group. Users that were discovered by Tivoli Storage Productivity Center can be included in a user group. **Note:** The user groups that you define in Data Manager for Databases will also available in Data Manager, and vice versa. Examples of user groups include Administrators, Developers, Marketing staff, HR staff
	Disk Manager
Storage Subsystem	A storage subsystem group can contain storage systems and storage virtualizers.
	Fabric Manager
Fabric	A grouping of fabrics.

Default groups

In the stand-alone GUI, default groups for each storage resource are provided within each group type node. As IBM Tivoli Storage Productivity Center discovers

these resources within your enterprise through discovery, probes, and scans, the resources are automatically added to the corresponding default groups.

The following table lists the names of the default groups for each group type.

Manager	Group Type	Default Group Name
Data Manager	Computer	TPCUser.Default Computer Group
	File System	TPCUser.Default FS Group
	Directory	n/a
	User	TPCUser.Default User Group
	OS User Group	TPCUser.Default OSGroup Group
Data Manager - Databases	Computer	TPCUser.Default Computer Group
	Databases-Tablespaces	TPCUser.Default Table Space Group
	Table	n/a
	User	TPCUser.Default User Group
Disk Manager	Storage Subsystems	TPCUser.Default Storage Subsystem Group
Fabric Manager	Fabrics	TPCUser.Default Fabric Group

Working with groups of resources

In the stand-alone GUI, create resource groups to logically arrange resources and to permit actions against the entire group.

Creating a group

Use the stand-alone GUI to create named groups that represent a set of resources, such as computers, databases, table spaces, tables, and users, to later use in pings, scans, and probes.

To create a new group, complete the following steps.

1. In the stand-alone GUI, expand **Monitoring** > **Groups**. For example, **Data Manager** > **Monitoring** > **Groups**.
2. Right-click the type of group you want to create and select **Create *group type***, where ***group type*** represents the group you want to create, such as a **Create Computer Group**, **Create Filesystem Group**, or **Create Table Group**.
3. In the Create *group type* Groups page, select the resources in the **Available** column that you want to include in the group. Hold down the Ctrl key as you click to select multiple resources.

   Note: For computer and file system groups in Data Manager, a **Group By Cluster** check box is displayed. When this option is not selected, the **Available** pane displays a cluster name next to any clustered computers and cluster resource groups. When this option is selected, any computer or cluster resource group that belongs to a IBM PowerHA® SystemMirror® for AIX or MSCS cluster is grouped under its cluster name. Cluster resource groups are identified with a special icon along with the name of the current host node. All other computers will be grouped under **Non-Clustered Computers**.

4. Click >> to move the selected resources to the **Current Selections** column. Click << to move resources from the **Current Selections** column and move them to the **Available** column.
5. Optional: Type a description of the group in the **Description** field.
6. Click the save icon in the menu bar to save the group.
7. Type a descriptive name for the group and click **OK**.

Grouping fabrics

In the stand-alone GUI, you can group fabrics to run reporting jobs against all the fabrics in the group. A fabric can belong to multiple groups at a time.

For more information on groups, see "Data collection groups" on page 334.

To create a fabric group, complete the following steps:

1. In the stand-alone GUI, expand **Fabric Manager** > **Monitoring** > **Groups**, right-click on **Fabric**, and select **Create Fabric Group**.
2. Enter a description for the group, and from the **Available** column, select the fabrics to be added to the group. Move them to the **Current Selections** column by clicking >>.
3. To save the group, click **File** > **Save**.
4. In the **Save As** window, enter the group name and click **OK**.

Editing groups

In the stand-alone GUI, when you update the resources that are included within a group, those updates are automatically picked up by the jobs that reference that group.

To edit an existing group, complete the following steps.

1. In the stand-alone GUI, expand **Monitoring** > **Groups**. For example, to edit groups within Disk Manager, expand **Disk Manager** > **Monitoring** > **Groups** > **Storage Subsystem**.
2. Expand a group type to view the list of groups that are defined for that type. For example, to view storage subsystem groups, expand the **Storage Subsystem** node under Disk Manager.
3. Click the name of the group you want to edit. Details for the selected group are displayed in the Create *group type* Groups page, where group type is the type of group (for example, Directory or Computer).
4. Edit the values that are defined for the group as desired. Click >> to move the selected resources to the **Current Selections** column. Click << to move resources from the **Current Selections** column and move them to the **Available** column.
5. Click the save icon in the menu bar to save the group.

Creating directory groups

To create directory groups, use the stand-alone GUI. Create directory groups within scans to monitor and gather information about specific directories within file systems across a number of systems.

Note: IBM Tivoli Storage SAN File Systems information: To monitor directories from the SAN File System, you must create a directory group and add the wanted directories as follows:

```
/<IBM Tivoli Storage SAN File Systems cluster name>/<directory name>
```

Where:

- `<IBM Tivoli Storage SAN File Systems cluster name>` represents the name of the IBM Tivoli Storage SAN File Systems cluster that contains the directory you want to add
- `<directory name>` represents the name of the directory you want to add

Examples:

- `/IBM/myFilesetName/myDirectoryName`
- `/IBM/myFilesetName`
- `/IBM/myDirectoryName`

To create a directory group, do the following steps.

1. In the stand-alone GUI, expand **Data Manager** > **Monitoring** > **Groups** > **Directory** or click **Create Directory Group** from the **Directory Groups** tab of the Scan window. The Create Directory Group window opens.
2. Use the New Directory window to edit or add directories to a directory group.

Field	Description
Creator	Displays the creator of the group.
Name	Displays the name of the group you are adding or editing. This name can be helpful in determining the content of the group. For example, a directory group named `UnixUserHomes` might consist of the `/home` directory on all the UNIX and Linux computers within an organization.
Description	A description of the group you are adding or editing.
New Computer, New Directory	Click these buttons to add a computer or directory to the group.

Viewing directory groups:
When adding directories to a group, you can determine how you want to view groups of directories in the stand-alone GUI. You can choose either of the following options:

- **Directories by computer.** Use this option to display the directories in the directory group according to the computer on which they exist.
  1. In the stand-alone GUI, expand Data Manager.
  2. Expand Monitoring.
  3. Expand Groups.
  4. Expand Directory.
  5. Double-click a directory object.
  6. In the content pane, under View, select Directories by Computer which is the default setting.

     Note: You might need to try more than one directory object to find one with data.

- **Computers by directory.** Use this option to display the computers in the group according to the directories included in the directory group.
  1. In the stand-alone GUI, expand Data Manager.
  2. Expand Monitoring.
  3. Expand Groups.
  4. Expand Directory.
  5. Double-click a directory object.
  6. In the content pane, under View, select Computers by Directory.

     Note: You might need to try more than one directory object to find one with data.

Tip: Choosing these different buttons enables you to add directories to a directory group in different ways. For example, if you want to monitor or receive alerts on different directories on multiple computers, choose the Directories by Computer option. If you want to monitor or receive alerts on the same directory across multiple computers, choose the Computers by Directory option.

Adding new directories by computer:
If you want to add new directories to a directory group by first selecting the computers on which they exist, complete the following steps:

1. In the stand-alone GUI, select **directories by computer** on the Create Directory Group window.
2. Click **New computer**. The Edit Directories for Computer window opens. This window enables you to select the computer for which you want to monitor specific directories, and enter the directories for the selected computer.

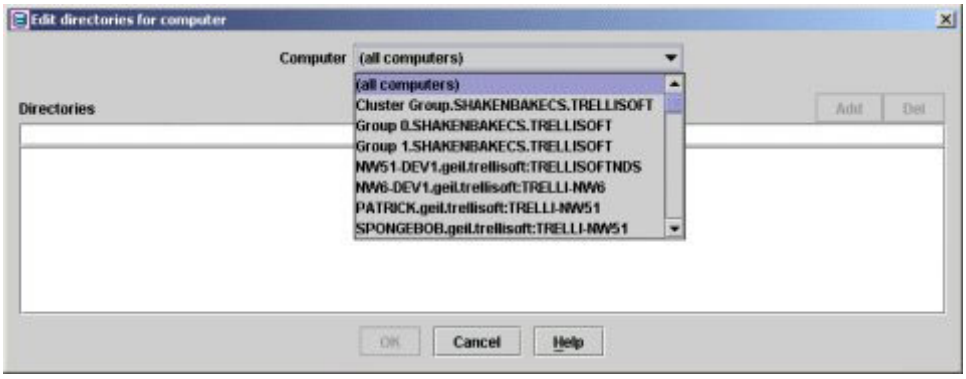

3. Click the **Computer** list box.
4. Select the computer for which you want to monitor specific directories.
5. Enter the directory you want to add to the group in the **Directories** field. For example: `c:\` or `/home/*`.
6. Click **Add**. The directory is added to the list box.
7. Click **OK** when you are done adding directories. The New Directory Group window opens. The computer you added will be displayed as a node in the display area, with the specified directories under the node.

Adding new computers by directory:
If you want to add new directories to a directory group by first specifying the directory and then selecting the computers on which they exist, complete the following steps:

1. In the stand-alone GUI, select **computers by directory** on the Create Directory Group window.
2. Click **New computer**. The Edit Computers for Directory window opens. This window enables you to enter the directories you want to add to the directory group and specify which computers contain those directories.
3. Enter the directory you want to add to the group in the **Directory** field. For example: `c:\`, `/home/*`, `\user\*`, `\windows\temp\`, `/usr/*`.
4. Select the computers you want to include in the group and click . The computer will be displayed in the **Current Selections** column. Select the **All computers** option to include all computers that contain the directory in the Directory Group.
5. Click **OK**. The New Directory Group window opens. The directory you added will be displayed as a node in the display area, with the selected computers under the node.

Alerting

Determining when and how you are alerted to conditions and violations of performance thresholds within your storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines the data about your resources for the conditions and performance thresholds that you specify.

Alerts are triggered by conditions and violations of performance thresholds that are detected during data collection and event processing. For some storage systems such as SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, or theXIV system, events are polled every minute from the resource. For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to Tivoli Storage Productivity Center when conditions change on the resource.

The conditions that trigger alert notifications depend on the type of resource that you are monitoring. Some triggering conditions, like performance thresholds, require you to enter values for triggering alerts. In general, the following types of conditions can trigger alerts:

- A performance threshold was violated
- A data collection job did not complete
- A change occurred in the storage infrastructure

For a complete list of the triggering conditions for alerts, see “Triggering conditions for alerts” on page 1665.

Setting up alerts

To set up alerts so that you are notified of conditions on monitored resources, you must complete the following tasks:

Tasks	GUI	Learn more
Add resources for monitoring by Tivoli Storage Productivity Center.	Web-based GUI[1]	“Adding resources” on page 257

Tasks	GUI	Learn more
Collect data about resources. Some alerts are triggered when conditions are detected during data collection. To detect violations of performance thresholds, you must run performance monitors to collect performance data about resources.	Web-based GUI: schedule probes and performance monitors[2] Stand-alone GUI: schedule pings and scans	"Collecting data" on page 295
Create alerts to specify the conditions on resources that trigger notifications.	Stand-alone GUI	"Creating alerts" on page 344
(optional) Configure how you want to be notified of alerts. By default, all alerts are recorded to a log and can be viewed in each GUI. Complete this task if you want to be notified of an alert in ways other than an entry in the log.	Stand-alone GUI	"Configuring alert destinations" on page 355
Tips: • [1] To add IBM SONAS systems and deploy Storage Resource agent, you must use the stand-alone GUI. • [2] In the web-based GUI, probes and performance monitors are automatically scheduled when resources are added for monitoring.		

Viewing and managing alerts

After alerts are set up, you can complete the following tasks to manage those alerts:

Tasks	GUI	Learn more
View, acknowledge, export, and remove alerts. For alerts that are triggered by performance threshold violations, you can view a chart of the historical performance of the resource that violated a threshold.	Web-based GUI†	• "Viewing all alerts" on page 356 • "Removing alerts" on page 362 • "Acknowledging alerts" on page 359
Modify, rename, and disable alerts.	Stand-alone GUI	• "Modifying alerts" on page 361 • "Renaming alerts" on page 362 • "Disabling or enabling alerts" on page 360
Tip: Some of these tasks are also available in the stand-alone GUI. For more information, access the online help in the stand-alone GUI.		

Alert notifications and actions

You can define the actions that are taken when alert conditions are detected on monitored resources. Use the stand-alone GUI to determine which of the following actions are taken when an alert is detected:

SNMP trap
Generate SNMP trap messages to any network management station (NMS), console, or terminal when an alert condition is detected. System

administrators must set up their SNMP trap ringer with the provided management information base (MIB) files to receive SNMP traps from the product.

Tivoli Enterprise Console® (TEC) or OMNIbus EIF Probe Server
Send alert notifications to a TEC server or OMNIbus EIF probe server within your environment that was configured to receive Tivoli Storage Productivity Center alerts.

Login notification
Send alert notifications to specific users when they view the alert log.

Windows event log or UNIX syslog
Write alert messages to the OS log. If you already have an administrator monitoring OS logs, this method is a way to centralize your priority messages for quick notification and viewing.

Run script
Run a script in response to an alert. You can run scripts that start other software tools for actions, such as archiving, backup and recovery, or provisioning.

Archive/Backup (Constraints only)
Define IBM Tivoli Storage Manager backup or archive jobs to run in response to an alert that is generated by a constraint violation. A constraint is a policy management feature that you can use to define the acceptable and unacceptable file types, file sizes, and file owners on monitored resources. For example, you can use constraints to alert you when users store certain files (such as MP3 or .avi files) on a monitored resource.

Email Send notifications to specific email addresses when an alert is detected on a monitored resource.

Related tasks:

"Adding resources" on page 257
You must add resources to Tivoli Storage Productivity Center before you can collect data, generate reports, and manage storage that is related to those resources.

Related reference:

"Triggering conditions for alerts" on page 1665
You can define alerts so that IBM Tivoli Storage Productivity Center notifies you when a specified event occurs. Such events are the *triggering conditions* for the alert. The specific triggering conditions that you can use to define an alert depend on the sort of storage resource that you are monitoring.

"Performance thresholds" on page 1690
Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

Alert severities

Tivoli Storage Productivity Center determines the severity of alert conditions that it detects on monitored resources. Use the severity level to help determine the priority in which you resolve alerts.

Each severity is represented by an icon in the web-based GUI.

Table 53. Possible severities of alerts

Icon	Status	Description
	Critical	An alert with a Critical severity represents a serious problem on a resource or on its internal resources. Resolve these problems as soon as possible.
	Critical - Acknowledged	An alert with a Critical severity was acknowledged. A Critical - Acknowledged severity indicates that an alert was reviewed and is either resolved or can be ignored.
	Warning	An alert with a Warning severity represents potential problems on a resource or on its internal resources. The condition that generated the alert is not critical.
	Warning - Acknowledged	An alert with a Warning severity was acknowledged. A Warning - Acknowledged severity indicates that an alert was reviewed and is either resolved or can be ignored.
	Informational	An alert with an Informational severity does not represent a problem, but is intended to provide information about actions related to a resource.
	Informational - Acknowledged	An alert with an Informational severity was acknowledged. An Informational - Acknowledged severity indicates that an alert was reviewed and can be ignored.

Creating alerts

In the stand-alone GUI, create alerts so that Tivoli Storage Productivity Center notifies you when a condition or violation of a performance threshold is detected on a resource. The specific triggering conditions and performance thresholds that you can specify are determined by the type of resource that you are monitoring.

Tip: When alerts are triggered, you can view information about those alerts in the web-based GUI.

Tivoli Storage Productivity Center includes a predefined set of alerts that correspond to events detected by Tivoli Storage Productivity Center for Replication. Only these alerts are available for replication. You can modify these alerts as required for your environment.

Creating storage system alerts

Create alerts that monitor storage systems for changes in configuration, status, and performance.

If necessary, complete the following prerequisites before you begin the task.

- If you want a script to be run after an event triggers an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.

To create a storage subsystem alert, complete the following steps:

1. Expand the **Disk Manager** > **Alerting** node in the navigation tree.
2. Right-click **Storage Subsystem Alerts**, and click **Create Storage Subsystem Alert**. The Create Storage Subsystem Alert window is opened.
3. In the **Description** field, type a brief description of the job.

4. Select the **Enabled** check box.
5. In the **Triggering Condition** pane, specify the event that triggers the alert:
   a. In the **Condition** field, select the triggering condition. See "Triggering conditions for alerts" on page 1665 for a list of triggering conditions that are available for each alert type.
   b. If you selected a performance threshold, specify the thresholds:
      1) In the **Critical Stress** field, type the upper-threshold boundary. When this boundary is exceeded, a critical condition is recognized.
      2) In the **Warning Stress** field, type the upper-threshold boundary. When this boundary is exceeded, a warning condition is recognized.
      3) In the **Warning Idle** field, type the lower-threshold boundary. When the value falls under this boundary, a warning condition is recognized.
      4) In the **Critical Idle** field, type the lower-threshold boundary. When the value falls under this boundary, a critical condition is recognized.

      For more information about thresholds, see "Performance thresholds" on page 1690.
6. In the **Alert Suppression** pane, choose the condition under which alerts are triggered, and then choose how to manage repeating alerts. For an explanation of the options in this pane, see "Alert suppression actions" on page 1681.
7. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
8. Click the **Storage Subsystems** tab. The Storage Subsystems page is displayed.
9. Select the storage systems and storage system groups to which you want the alert to apply:
   a. In the **Available** field, click the storage entities.
   b. Click . The storage entities are moved to the **Current Selections** field.
10. Save the alert:
    a. Click **File** > **Save**.
    b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

Creating switch alerts

Create alerts that monitor switches for changes in configuration, status, and performance.

If necessary, complete the following prerequisites before you begin the task.

- If you want a script to be run after an event triggers an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.

To create a switch alert, complete the following steps:

1. In the Navigation Tree pane, expand **Fabric Manager** > **Alerting**.
2. Right-click **Switch Alerts** and click **Create Switch Alerts**. The Create Switch Alerts page opens and the Alert page is displayed.
3. Select the **Enabled** check box.

4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** pane, specify the event that triggers the alert:
   a. In the **Condition** field, select the triggering condition. See "Triggering conditions for alerts" on page 1665 for a list of triggering conditions that are available for each alert type.
   b. If you selected a performance threshold, specify the thresholds:
      1) In the **Critical Stress** field, type the upper-threshold boundary. When this boundary is exceeded, a critical condition is recognized.
      2) In the **Warning Stress** field, type the upper-threshold boundary. When this boundary is exceeded, a warning condition is recognized.
      3) In the **Warning Idle** field, type the lower-threshold boundary. When the value falls under this boundary, a warning condition is recognized.
      4) In the **Critical Idle** field, type the lower-threshold boundary. When the value falls under this boundary, a critical condition is recognized.

      For more information about thresholds, see "Performance thresholds" on page 1690.
6. In the **Alert Suppression** pane, choose the condition under which alerts are triggered, and then choose how to manage repeating alerts. For an explanation of the options in this pane, see "Alert suppression actions" on page 1681.
7. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
8. Click the **Switches** tab. The Switches page is displayed.

   Note: If you selected a performance threshold as the triggering condition, only those switches for which performance data can be collected are listed.
9. Select the storage entities to which you want the alert to apply:
   a. In the **Available** field, click the storage entities.
   b. Click . The storage entities are moved to the **Current Selections** field.

   Note: You can select one or more switches from the **Switches** category. For device triggering conditions only, you can select one or more fabrics from the **Switches In Fabric** category. Selecting fabrics will include all switches in those fabrics. You can also select from both categories.
10. Save the alert:
    a. Click **File** > **Save**.
    b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

Creating fabric alerts

Create alerts that monitor fabrics for conditions that you specify. These alerts can be applied to fabrics, zones, and zone sets.

If necessary, complete the following prerequisites before you begin the task.

- If you want a script to be run after an event triggers an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must

specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.

1. In the navigation tree of the stand-alone GUI, expand **Fabric Manager** > **Alerting**.
2. Right-click **Fabric Alert** and click **Create Fabric Alerts**. The Create Fabric Alerts window opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that triggers the alert.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
7. Click the **Fabrics** tab; the Fabrics page opens.
8. Select the resources to which you want the alert to apply:
   a. In the **Available** field, click the resources.
   b. Click . The resources are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** > **Save**.
   b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

Creating computer alerts

Create alerts that monitor computers for changes in configuration or status.

If necessary, complete the following prerequisites before you begin the task.

- If you want a script to be run after an event triggers an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.

To create a computer alert, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Data Manager** > **Alerting**.
2. Right-click **Computer Alert** and click **Create Alert**. The Create Alert window opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering-Condition** area, specify the event that triggers the alert:
   a. In the **Condition** field, select the triggering condition.
   b. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.
   c. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
7. Click the **Computer** tab; the Computers page opens.

8. Select the resources to which you want the alert to apply:
   a. In the **Available** field, click the resources.
   b. Click . The resources are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** > **Save**.
   b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

Creating file system alerts

Create alerts that monitor file systems for changes in available free space and configuration, or for file systems that are removed or unmounted.

If necessary, complete the following prerequisites before you begin the task.

- If you want a script to be run after an event triggers an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.

File system alerts can be applied to file systems, file system groups, computers, computer groups, storage systems, storage system groups, and clusters. To create a file system alert, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Data Manager** > **Alerting**.
2. Right-click **File System Alert**, and then click **Create Alert**. The **Create Alert** panel opens and the Alertspage is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering-Condition** area, specify the event that triggers the alert:
   a. In the **Condition** field, select the triggering condition.
   b. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.
   c. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
7. Click the **File Systems** tab; the File Systems window opens.
8. Select the resources to which you want the alert to apply:
   a. In the **Available** field, click the resources.
   b. Click . The resources are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** > **Save**.
   b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

Creating directory alerts

Create alerts that monitor when a directory is not found or consumes more than a specified amount of storage space.

Before you can create a directory alert, you must have completed the following prerequisites:

- If you want a script to be run after an event triggers an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.
- Directories must be grouped into a directory group before they can be included in an alert.

To create a directory alert, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Data Manager** > **Alerting**
2. Right-click **Directory Alert** and click **Create Alert**. The Create Alert window opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering-Condition** area, specify the event that triggers the alert:
   a. In the **Condition** field, select the triggering condition.
   b. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.
   c. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
7. Click the **Directories** tab; the Directories windows opens.
8. Select the resources to which you want the alert to apply:
   a. In the **Available** field, click the resources.
   b. Click

. The resources are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** > **Save**.
   b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

Creating hypervisor alerts

Create alerts that monitor hypervisors for conditions that you specify.

If necessary, complete the following prerequisites before you begin the task.

- If you want a script to be run after an event triggers an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.

To create a hypervisor alert, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Data Manager** > **Alerting**.
2. Right-click **Hypervisor Alert** and click **Create Alert**.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering-Condition** area, specify the event that triggers the alert:
   a. In the **Condition** field, select the triggering condition.
   b. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.
   c. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
7. Click the **Hypervisors** tab to view a list of monitored hypervisors.
8. Select the resources to which you want the alert to apply:
   a. In the **Available** field, click the resources.
   b. Click . The resources are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** > **Save**.
   b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

Creating other NAS alerts

Use other network-attached storage (NAS) alerts to monitor systems that include NAS and are not included in other alert types. When a change in configuration or status is detected during a probe or from an SNMP event, an alert is generated.

If necessary, complete the following prerequisites before you begin the task.

- If you want a script to be run after an event triggers an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.

Other NAS alerts can be applied to NAS clusters, nodes, storage systems, and storage system groups. To create a NAS alert, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Data Manager** > **Alerting**.
2. Right-click **Other NAS Alerts**, and then click **Create Alert**. The Create Alert panel opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering-Condition** area, specify the event that triggers the alert:
   a. In the **Condition** field, select the triggering condition.
   b. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.

c. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.

6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
7. Click the **NAS** tab; the NAS page opens.
8. Select the resources to which you want the alert to apply:
   a. In the **Available** field, click the resources.
   b. Click . The resources are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** > **Save**.
   b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

Creating instance alerts

Create alerts that monitor instances for conditions that you specify. These alerts can be applied to instances, computers, and computer groups.

If necessary, complete the following prerequisites before you begin the task.

- If you want a script to be run after an event triggers an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.

To create an instance alert, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Data Manager for Databases** > **Alerting**.
2. Right-click **Instance Alerts** and click **Create Alert.**. The Create Alert window opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that triggers the alert:
   a. In the **Rdbms Type** field, select the type of database application. You can choose **All Rdbms**, **Oracle**, **SQL\Server**, **Sybase**, or **UDB**.
   b. In the **Condition** field, select the triggering condition.
   c. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.
   d. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
7. Click the **Instances** tab; the **Instances** window opens.
8. Select the resources to which you want the alert to apply:
   a. In the **Available** field, click the resources.
   b. Click 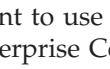 . The resources are moved to the **Current Selections** field.

9. Save the alert:
   a. Click **File** > **Save**.
   b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

Creating database and tablespace alerts

Create alerts that monitor database and tablespaces for conditions that you specify. These alerts can be applied to databases, database groups, table spaces, table space groups, instances, computers, and computer groups.

If necessary, complete the following prerequisites before you begin the task.

- If you want a script to be run after an event triggers an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.

Note: (DB2 only) Alerts triggered by table and user statistics will not function until the **runstats** utility is run. This utility ensures that table and user statistics are up-to-date. When you define a scan job for DB2 databases, select the **Perform runstats on all tables being scanned** check box on the **Options** page.

Note: (DB2 only) When the following conditions apply, an alert is not generated:

- The triggering condition selected is Database-Tablespace Free Space Less Than.
- The storage entities selected are DB2 databases.
- The table spaces affected by the triggering condition are System Managed Space (SMS) tablespaces.

The alert is not triggered because IBM Tivoli Storage Productivity Center considers SMS containers to be directories on a file system

To create a database-table space alert, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Data Manager for Databases** > **Alerting**.
2. Right-click **Database-Tablespace Alerts** and click **Create Alert.**. The Create Alert window opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that triggers the alert:
   a. In the **Rdbms Type** field, select the type of database application. You can choose **All Rdbms**, **Oracle**, **SQL\Server**, **Sybase**, or **UDB**.
   b. In the **Condition** field, select the triggering condition.
   c. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.
   d. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
7. Click the **Databases-Tablespaces** tab; the **Databases-Tablespaces** window opens.

8. Select the resources to which you want the alert to apply:
   a. In the **Available** field, click the resources.
   b. Click . The resources are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** > **Save**.
   b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

Creating table alerts

Create alerts that monitor tables for conditions that you specify. These alerts can be applied to table groups and tables only.

- If you want a script to be run after an event triggers an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.
- Tables must be grouped into a table group before they can be included in an alert.

To create a table alert, complete the following steps:

1. From the Navigation Tree pane, expand **Data Manager for Databases** > **Alerting**.
2. Right-click **Table Alerts** and click **Create Alert.**. The Create Alert window opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that triggers the alert:
   a. In the **Rdbms Type** field, select the type of database application. You can choose **All Rdbms**, **Oracle**, **SQL\Server**, **Sybase**, or **UDB**.
   b. In the **Condition** field, select the triggering condition.
   c. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.
   d. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
7. Click the **Tables** tab; the **Tables** window opens.
8. Select the resources to which you want the alert to apply:
   a. In the **Available** field, click the resources.
   b. Click . The resources are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** > **Save**.
   b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

Creating endpoint device alerts

Create alerts that monitor endpoint devices for conditions that you specify.

If necessary, complete the following prerequisites before you begin the task.

- If you want a script to be run after an event triggers an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.

To create an endpoint device alert, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Fabric Manager** > **Alerting**.
2. Right-click **Endpoint Device Alerts** and click **Create Endpoint Device Alerts**. The Create Endpoint Device Alerts window opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that triggers the alert.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
7. Click the **Endpoint Devices** tab; the Endpoint Devices page opens.
8. Select the resources to which you want the alert to apply:
   a. In the **Available** field, click the resources.
   b. Click . The resources are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** > **Save**.
   b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

Database connection alerts

Database connection alerts notify the Tivoli Storage Productivity Center administrator when the database connection used by the Tivoli Storage Productivity Center server is lost unexpectedly and when the database connection is restored.

If necessary, complete the following prerequisites before you begin the task.

- If you want a script to be run after an event triggers an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.

The database connection alerts are always enabled and configured to add the alert to the alert log and send an e-mail to the data administrator.

The database connection alerts are:

- DB connection failed

- DB connection successful

You cannot create, configure, delete, disable, edit, or rename these alerts. You can, however, change the destination for the database connection alerts. Expand **Administrative Services** > **Configuration**. Left-click **Alert Disposition**. In the content pane, under Email, you can change all of the fields.

Configuring alert destinations

In the stand-alone GUI, configure the control information that is used to generate alerts when a corresponding event is discovered. You can specify default configuration settings for SNMP, TEC/OMNIbus, and email-directed alerts.

If you want to send SNMP traps, IBM Tivoli Enterprise Console events, or email when an alert is triggered, you must configure IBM Tivoli Storage Productivity Center with the alert destinations and email settings.

To configure alert dispositions, complete the following steps:

1. In the navigation tree of the stand-alone GUI, expand **Administrative Services** > **Configuration.**
2. Click **Alert Disposition**. The Edit Alert Disposition window opens.
3. In the **SMNP** area, specify the information about where SNMP traps are sent:
   a. In the **Community** field, type the name of the SNMP community. By default, this is set to public.
   b. In the **Host field**, type the fully qualified name of the system that you have configured to receive SNMP traps.
   c. In the **Port** field, type the port number to which Tivoli Storage Productivity Center will direct SNMP traps. By default, this is set to 162.
   d. If you want to configure a second SNMP destination, type the required information in the fields in the **SNMP Destination 2** area.
4. In the **TEC** area, specify information about where Tivoli Enterprise Console events are sent:
   a. In the **Server** field, type the fully qualified name of the Tivoli Enterprise Console server.
   b. In the **Port** field, type the port number to which Tivoli Storage Productivity Center will direct Tivoli Enterprise Console events. By default, this is set to 5529.
5. In the **Email** area, configure email settings:
   a. In the **Mail Server** field, type the name of the mail server.
   b. In the **Mail Port** field, type the port number for the SMTP server.
   c. In the **Default Domain field**, type the domain, for example, MyCompany.com.
   d. In the **Return to field**, type the email address to which undeliverable mail is sent.
   e. In the **Reply to field**, type the email address to which email is sent when someone replies to an email notification about an alert. (It is a best practice to specify the email address of an Tivoli Storage Productivity Center administrator.)
   f. In the **Data Administrator** field, type the email address of the Data Administrator.
6. In the **Alert Log Disposition** field, select the length of time to keep entries in the alert logs. The default is 90 days.

Viewing all alerts

Use the Alerts page in the web-based GUI as a central location to view and manage all the alerts that were detected by Tivoli Storage Productivity Center.

The Alerts page is organized into two sections:

Alerts summary chart

The summary chart at the top of the page is a visual representation of the alerts that were detected on monitored resources. The chart consists of stacked bars with the number and severity of current alerts on the y-axis and the resource type on the x-axis. The y-axis labeling scales dynamically according to the highest number of alerts in any of the categories.

Alerts list

The alerts list shows information about all the alerts that were detected on monitored resources. Information in the list is organized into columns. These columns include information about the condition that was detected, the severity of the alert, when the alert was detected, and the names of the resources where the alert condition was detected.

The version of Tivoli Storage Productivity Center that detected an alert determines where information about the alert is displayed. Alerts that were generated in version 5.1 or later are displayed on the Alerts page and in the **Alerts** tab on the pages for the related resources. Alerts that were generated in versions of the product before 5.1, also called migrated alerts, are only displayed on the Alerts page.

Tip: Depending on how data is configured for a NetApp filer, alerts are shown on different pages in the interface:

- Alerts for the block data in a NetApp filer are shown in the **Alerts** tab on the Storage Systems page and in the list of alerts on the Alerts page.
- Alerts for the file data in a NetApp filer are only shown in the list of alerts on the Alerts page.

1. In the navigation pane, go to **Home** > **Alerts**. Information about alerts is shown in the list of alerts.
2. Optional: To view the number of alerts for a specific type of resource, hover the mouse pointer over the related icon in the summary chart. For example, if you want to view the number of Critical, Warning, and Informational alerts that were triggered by hypervisors, hover the mouse pointer over the **Hypervisors** icon .
3. Optional: To filter the list of alerts for a specific type of resource, click the icon for that resource in the summary chart. For example, if you only want to view a list of alerts for hypervisors, click the **Hypervisors** icon .

   Tip: To view descriptions of the icons and columns on the Jobs page, click the **Help** icon in the upper-right corner of the page.

Viewing alert logs

In the stand-alone GUI, you can also use the Alert Log to view alerts from a central location.

There are two alert logs that are available in the stand-alone GUI. The type of alert determines where information about the alert is shown:

- For alerts generated by Data Manager for Databases, expand **Data Manager for Databases** > **Alerting** > **Alert Log**. You can view:
  - All alerts or alerts organized by RDBMS such as by Oracle, SQL Server, Sybase, and UDB. If you do not license for a specific RDBMS, no nodes for that RDBMS are displayed under the **Alert Log** node.
  - Alerts by type (such as by Instance, Database-Tablespace, Table, User)
  - Alerts directed to you (the user currently logged on to the UI component)
- For alerts generated by all other components, expand **IBM Tivoli Storage Productivity Center** > **Alerting** > **Alert Log**. The alert types that are displayed depend on the Tivoli Storage Productivity Center license type. The following is a list of the possible alert types that are displayed under **Alert Log**:
  - All
  - Alerts Directed to <user-id> - Alerts directed to the user logged on to the interface.
  - Authentication Configuration
  - Replication
  - Storage Subsystem
  - Computer
  - Disk
  - Filesystem
  - Directory
  - User
  - OS User Group
  - Fabric
  - Switch
  - Endpoint Device
  - External
  - Configuration Analysis
  - Hypervisor

  For alert types that are not displayed for certain license types, see "Product licenses" on page 21

Each alert node has an icon next to it that indicates its status:

- Red-outlined circle: a negative alert occurred, such as job failed, new disk defect.
- Green-outlined box: a neutral alert was detected, such as a new disk detected.

When you click an alert node, the alert log window opens a table that contains information for all the alerts for that node. Click on a row to drill down for more information.

The following table describes the buttons on the Alert Log window:

Button	Description
Clear	Clear the selected alerts. When you click this button, the icon for the alert changes to indicate that the alert was cleared. The alert continues to be displayed in the alert log.
Delete	Delete the selected alerts from the alert log.
Clear ALL	Clear all of the alerts in the report. When you click this button, the icon for all the alerts changes to indicate that the alerts were cleared. The alerts continue to be displayed in the alert log.

Button	Description
Delete ALL	Delete all of the alerts in the current report from the alert log.
Refresh	Regenerate the report you are viewing to pick up any newly generated alerts that were triggered since you last generated the report.

Tip: Some alerts are shown automatically when a client application is registered to receive alerts from Tivoli Storage Productivity Center. You can identify these alerts by the keyword *Pre-defined* in their names. You can hide these alerts from view in the GUI. For more information, see the IBM Tivoli Storage Productivity Center Problem Determination Guide.

Restriction: For database tables, the values in the **Resource Type** and **Resource Name** columns are set to N/A for the following alerts:

- Total Table Size Greater Than
- Chained Row Count Greater Than

Viewing alerts for a type of resource

Use resource list pages in the web-based GUI to view only the alerts for a specific type of top-level resource. Each resource list page has an **Alerts** tab that lists only the alerts for monitored storage systems, servers, hypervisors, switches, or fabrics. For example, you can use the list page for storage systems to view only the alerts that were detected on monitored storage systems.

1. In the navigation pane of the web-based GUI, select a type of top-level resource. For example, if you want to view the alerts that are related to storage systems, select **Storage Resources** > **Storage Systems**.
2. Click the **Alerts** tab to view a list of all the alerts for the monitored resources of that type. The information about alerts is organized into columns. These columns include information about the condition that was detected, the severity of the alert, when the alert was detected, and the names of the resources where the alert condition was detected.

   Tip: The icon on the **Alerts** tab represents the most severe alert condition that was detected on the monitored resources. For example, if one of the jobs for storage systems failed, the following icon is shown on the tab:

   Alerts.
3. Optional: View the icons at the top of the **Alerts** tab for a summary of alert severities. This summary includes the number of alerts that have a Critical, Warning, Informational, or Acknowledged status.

   Tip: To view descriptions of the icons and columns on the **Alerts** tab, click the **Help** icon in the upper-right corner of the page.

Viewing alerts for a specific resource

Use resource detail pages in the web-based GUI to view only the alerts for a specific top-level resource. Each resource details page has an **Alerts** link that lists only the alerts for a monitored storage system, server, hypervisor, switch, or fabric. For example, you can use the details page for a storage system to view only the alerts that were detected for that storage system.

1. In the navigation pane of the web-based GUI, select a type of top-level resource. For example, if you want to view the alerts that are related to storage systems, select **Storage Resources** > **Storage Systems**.
2. Right-click a resource in the list and select **View Details**. In the **General** section, the number next to **Alerts** represents the number of alert conditions that were detected on the resource. For example, **Alerts (25)** indicates that 25 alerts were detected on the resource.

   The icon next to the Alerts link represents the most severe alert condition that was detected on the resource. For example, if alerts with an Error severity were detected the following icon is shown:

   Alerts (25)
3. Click the **Alerts** link in the **General** section to view the alerts that were detected on the resource. The information about alerts is organized into columns. These columns include information about the condition that was detected, the severity of the alert, when the alert was detected, and the names of the resources where the alert condition was detected.
4. Optional: View the icons at the top of the page for a summary of alert severities that were detected. This summary includes the number of alerts that have a Critical, Warning, Informational, or Acknowledged status.

   Tip: To view descriptions of the icons and columns that are shown for alerts, click the **Help** icon in the upper-right corner of the page.

Acknowledging alerts

Some alerts in Tivoli Storage Productivity Center are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify the other alerts in lists that must still be reviewed and resolved.

Tip: Acknowledging alerts does not affect the status of resources. The status of a resource is determined by its consolidated status and the propagated status of its internal resources, not by alerts that are generated. For information about how to acknowledge the status of a resource, see "Acknowledging the status of resources" on page 380.

To acknowledge an alert, complete the following steps:

1. In the web-based GUI, go to a list of alerts that includes the alerts that you want to acknowledge. For example, if you want to acknowledge the alerts for a specific server, complete the following steps.
   a. In the navigation pane, go to **Server Resources** > **Servers**.
   b. Right-click the server that has alerts you want to acknowledge and select **View Details**.
   c. Click the **Alerts** link.
2. Right-click the alert that you want to acknowledge and select **Acknowledge alerts**. The severity of the alert is changed to Acknowledged on the Alerts page, resource list page, and resource details page.

   For example, if you acknowledge a server alert that has a Warning severity, it is shown on the Alerts page, Servers page, and Server details page with the following severity:

   Warning - Acknowledged

3. Optional: To acknowledge multiple alerts at the same time, press Ctrl and click each alert. To select a series of alerts, select the first alert, and then press Shift and click the last alert. When the alerts are selected, right-click a row in the list and select **Acknowledge alerts**.
4. Optional: To acknowledge all the alerts that are shown in a list, select **Acknowledge all alerts** from the **Actions** menu.
5. Optional: To acknowledge all the alerts with an Informational severity in a list, select **Acknowledged informational alerts** from the **Actions** menu. You might want to acknowledge all Informational alerts when you want to view only alerts that represent potential or serious problems that require action to resolve.

Disabling or enabling alerts

Disable an alert so that Tivoli Storage Productivity Center does not generate notifications for its specified condition. You can enable alert that is disabled at any time.

To disable or enable an alert, complete the following steps:

1. In the navigation tree of the stand-alone GUI, complete one of the following actions:

Alert type	Action
For computer alerts	Expand **Data Manager** > **Alerting** > **Computer Alerts**
For file system alerts	Expand **Data Manager** > **Alerting** > **Filesystem Alerts**
For directory alerts	Expand **Data Manager** > **Alerting** > **Directory Alerts**
For hypervisor alerts	Expand **Data Manager** > **Alerting** > **Hypervisor Alerts**
For other NAS alerts	Expand **Data Manager** > **Alerting** > **Other NAS Alerts**
For instance alerts	Expand **Data Manager for Databases** > **Alerting** > **Instance Alerts**
For database-tablespace alerts	Expand **Data Manager for Databases** > **Alerting** > **Database-Tablespace Alerts**
For table alerts	Expand **Data Manager for Databases** > **Alerting** > **Table Alerts**
For storage subsystem alerts	Expand **Disk Manager** > **Alerting** > **Storage Subsystem Alerts**
For fabric alerts	Expand **Fabric Manager** > **Alerting** > **Fabric Alerts**
For switch alerts	Expand **Fabric Manager** > **Alerting** > **Switch Alerts**
For endpoint device alerts	Expand **Fabric Manager** > **Alerting** > **Endpoint Device Alerts**
For replication alerts	Expand **Replication Manager** > **Alerting** > **Replication Alerts**

   The list of existing alerts is displayed.
2. Click the alert that you want to disable or enable. The Edit Alert page opens.

3. To disable an alert, clear the **Enabled** check box. To enable an alert, select the **Enabled** check box.
4. Click **File** > **Save**.

Modifying alerts

Modify an alert if you want to change the triggering condition, triggered action, or the resources for which it monitors.

To modify an alert, complete the following steps:

1. In the navigation tree of the stand-alone GUI, complete one of the following actions:

Alert type	Action
For computer alerts	Expand **Data Manager** > **Alerting** > **Computer Alerts**
For file system alerts	Expand **Data Manager** > **Alerting** > **Filesystem Alerts**
For directory alerts	Expand **Data Manager** > **Alerting** > **Directory Alerts**
For hypervisor alerts	Expand **Data Manager** > **Alerting** > **Hypervisor Alerts**
For other NAS alerts	Expand **Data Manager** > **Alerting** > **Other NAS Alerts**
For instance alerts	Expand **Data Manager for Databases** > **Alerting** > **Instance Alerts**
For database-tablespace alerts	Expand **Data Manager for Databases** > **Alerting** > **Database-Tablespace Alerts**
For table alerts	Expand **Data Manager for Databases** > **Alerting** > **Table Alerts**
For storage subsystem alerts	Expand **Disk Manager** > **Alerting** > **Storage Subsystem Alerts**
For fabric alerts	Expand **Fabric Manager** > **Alerting** > **Fabric Alerts**
For switch alerts	Expand **Fabric Manager** > **Alerting** > **Switch Alerts**
For endpoint device alerts	Expand **Fabric Manager** > **Alerting** > **Endpoint Device Alerts**
For replication alerts	Expand **Replication Manager** > **Alerting** > **Replication Alerts**

   The list of existing alerts is displayed.
2. Click the alert that you want to modify. The Edit Alert page opens.
3. Make any necessary changes.
4. Click **File** > **Save**.

Tip: Some alerts are shown automatically when a client application is registered to receive alerts from Tivoli Storage Productivity Center. You can identify these alerts by the keyword *Pre-defined* in their names. You can hide these alerts from view in the Tivoli Storage Productivity Center. For more information, see the IBM Tivoli Storage Productivity Center Problem Determination Guide.

Renaming alerts

Rename alerts at any time. This action might be helpful if the original name is ambiguous or lacks clarity.

To rename an alert, complete the following steps:

1. In the navigation tree of the stand-alone GUI, complete one of the following actions:

Alert type	Action
For computer alerts	Expand **Data Manager** > **Alerting** > **Computer Alerts**
For file system alerts	Expand **Data Manager** > **Alerting** > **Filesystem Alerts**
For directory alerts	Expand **Data Manager** > **Alerting** > **Directory Alerts**
For hypervisor alerts	Expand **Data Manager** > **Alerting** > **Hypervisor Alerts**
For other NAS alerts	Expand **Data Manager** > **Alerting** > **Other NAS Alerts**
For instance alerts	Expand **Data Manager for Databases** > **Alerting** > **Instance Alerts**
For database-tablespace alerts	Expand **Data Manager for Databases** > **Alerting** > **Database-Tablespace Alerts**
For table alerts	Expand **Data Manager for Databases** > **Alerting** > **Table Alerts**
For storage subsystem alerts	Expand **Disk Manager** > **Alerting** > **Storage Subsystem Alerts**
For fabric alerts	Expand **Fabric Manager** > **Alerting** > **Fabric Alerts**
For switch alerts	Expand **Fabric Manager** > **Alerting** > **Switch Alerts**
For endpoint device alerts	Expand **Fabric Manager** > **Alerting** > **Endpoint Device Alerts**
For replication alerts	Expand **Replication Manager** > **Alerting** > **Replication Alerts**

   The list of existing alerts is displayed.
2. Right-click the alert, and click **Rename**. The Rename Alert window opens.
3. Type the new name, and click **OK**.

Removing alerts

Remove alerts when you no longer require to view them. By default, alerts are automatically removed based on retention settings that are defined in the stand-alone GUI. However, you can use web-based GUI to manually remove alerts at any time.

1. In the web-based GUI, go to a list of alerts that includes the alerts that you want to remove. For example, if you want to remove alerts for all different types of top-level resources, go to **Home** > **Alerts**. If you want to remove alerts for a specific type of resource such as storage systems, go to **Storage Resouces** > **Storage Systems** and click the **Alerts** tab.
2. Right-click the alert that you want to remove and select **Remove alerts**.

3. Click **Remove** to remove the alert. When you remove an alert, it is no longer shown on the Alerts page or in the **Alerts** tabs on the related resource list and resource detail pages.

   For example, if you remove an alert for a fabric, that alert is removed from the Alerts page, Fabrics page, and Fabric details page.
4. Optional: To remove all the alerts that are shown in a list, select **Remove all alerts** from the **Actions** menu.
5. Optional: To remove all the acknowledged alerts that are shown in a list, select **Remove acknowledged alerts** from the **Actions** menu.

Viewing information about resources

You can view detailed information about the resources that are monitored by Tivoli Storage Productivity Center. Resources include storage systems, volumes, pools, servers, hypervisors, fabrics, and switches. You can also view information about internal and related resources.

Remember: Data collection jobs retrieve information about monitored resources. To ensure that the latest information about resources is available in the web-based GUI, run data collection jobs regularly. To learn about how to collect data, see "Collecting data" on page 295.

How information is organized

Information about monitored resources is organized into different pages in the web-based GUI. You can use these pages to view summary information about all resources, detailed information about top-level resources, and detailed information about internal and related resources.

The dashboard

The dashboard is displayed when you start a new session of Tivoli Storage Productivity Center or when you click **Home** > **Dashboard** in the navigation pane. Use the dashboard to view the following summary information about a storage environment:

- A topology or resource diagram of the top-level resources that are monitored by Tivoli Storage Productivity Center. This diagram shows the number of top-level resources that are monitored and their overall status. You can also click the icons in the resource diagram to view more detailed information about a top-level resource on a resource list page.
- Capacity information that was aggregated for storage systems, servers, and hypervisors. Use this information to view the overall storage capacity for each type of resource. This information is presented as horizontal bar charts.
- The status of Tivoli Storage Productivity Center jobs that were run during a specific time. You can also click links to access a complete list of jobs.
- The number of unacknowledged alerts that were detected on resources during a specific time. You can also click a link to access a complete list of alerts that were detected on all monitored resources.

For more information about the dashboard, see "Viewing overall storage information on the dashboard" on page 365.

Jobs page

Use the Jobs page to view the following information about Tivoli Storage Productivity Center jobs:

- The overall status of all jobs. Use this information to identify the number of jobs that are encountering problems when running.
- Information about specific jobs, including status, related resource, type, and schedule.
- Information about performance monitors, including status, last successful collection attempt, start time, end time, and remaining duration.
- Information in the log files for jobs.

To access the Jobs page, click **Home** > **Jobs** in the navigation pane.

Alerts page

Use the Alerts page to view and manage the alert conditions that were detected on monitored resources. You can complete the following tasks on this page:

- View the overall number of alerts that were detected on monitored resources. This information is presented in a bar chart for the following types of alerts:
  - Database
  - Fabric
  - Hypervisor
  - Other
  - Performance
  - Replication
  - Server
  - Storage system
  - Tivoli Storage Productivity Center
- View information about specific alerts, including severity, last occurrence, and the resources where they were detected.
- Acknowledge alerts that were reviewed but are not yet resolved.
- Unacknowledge alerts.
- Remove alerts from the page.

To access the Alerts page, click **Home** > **Alerts** in the navigation pane.

Resource list pages

Use a resource list page as central location for viewing information about the top-level resources of a specific type. For example, use the Storage Systems page to view a list of the storage systems that are monitored by Tivoli Storage Productivity Center. Each resource list page corresponds to a specific type of resource, which includes storage systems, volumes, pools, servers, hypervisors, switches, and fabrics. You can view the following information on resource list pages:

- The overall status of the resources
- Key information about the resources
- Alerts that were generated for the resources
- Jobs that were run for the resources

To learn about how to access the list pages for resources, see "Viewing information about top-level resources" on page 371.

Note: Volumes and pools are internal resources of storage systems, but have their own resource list pages so that you can view them all in a central location.

Resource detail pages

Use resource details pages to view information about resources that are internal or related to a top-level resource. For example, use the Storage System details page to view detailed information about a specific storage system. You can view the following information on resource detail list pages:

- The overall capacity and storage usage for a resource
- Key information about a resource
- Alerts that were generated for a resource
- Jobs that were run for a resource
- Information about components that exist in top-level resources
- Information about the resources that are related to another resource

To learn about how to access detail pages for resources, see "Viewing information about internal and related resources" on page 371.

Reporting

Use the IBM Tivoli Common Reporting tool in the web-based GUI to view detailed reports about resources. The following reports are available:

Predefined reports
These reports are included with Tivoli Storage Productivity Center. The following types of predefined reports are available:

- Resources
- Performance
- Historical space
- Storage tiering

Custom reports
Unlike predefined reports, you can select the information that you want to include in custom reports. Custom reports can contain detailed information about the relationships between monitored resources, the properties of monitored resources, and detailed information about the performance of monitored resources.

To learn about how to use the IBM Tivoli Common Reporting tool, see Chapter 7, "Reporting," on page 789.

Viewing overall storage information on the dashboard

The dashboard is shown when you first log in to the GUI or when you select **Dashboard** from the **Home** icon in the navigation pane. Use the dashboard to monitor the overall status of monitored resources and identify potential problem areas in a storage environment. You can learn about the status and capacity of resources, the jobs that collect performance data, and the unacknowledged alert conditions that are detected on resources.

Specifically, the dashboard provides a concise, yet detailed overview of your storage environment that includes the following information:

- The status of monitored resources on the resource diagram
- The overall capacity of monitored resources

- The status of performance monitors
- The status of alerts that were detected on monitored resources

Important: The information that is shown on the dashboard is gathered by data collection jobs. To ensure that the latest information is available, run these jobs regularly.

Viewing the status of monitored resources on the resource diagram

The diagram in the top section of the dashboard provides a visual representation of the resource types that are monitored by Tivoli Storage Productivity Center. Use this diagram to view the following information:

A list of resources that are being monitored for each resource type

Each type of resource in the diagram is represented by an icon. To learn more detailed about the individual resources that are monitored for a type, click the related icon in the diagram. For example, if you click the icon for storage systems, a new page shows detailed information about the individual storage systems that are being monitored.

The number of resources that are monitored for a resource type

A number in parentheses is shown next to each resource icon. This number represents the number of resources for each type that is monitored. For example, if 20 storage systems are monitored, the value **Storage Systems (20)** is shown. This number excludes resources that are not monitored by Tivoli Storage Productivity Center.

Tip: If no resources are being monitored for a resource type, click its icon on the dashboard to add a resource of that type. For example, if the dashboard shows **Storage Systems (0)**, you can click the icon for storage systems to add a storage system. The GUI guides you though the process for adding the resource for monitoring. For more information about adding resources, see "Adding resources" on page 257.

The overall status for each resource type

A status symbol is shown next to each resource icon. This symbol represents the most critical status that was detected on the internal resources for a resource type. For example, if 20 storage systems are monitored, and an error was detected on a volume for one of those storage systems, a red symbol is shown next to the storage systems icon in the diagram. If no errors, warnings, or unreachable statuses were detected on the internal resources of monitored storage systems, then a green symbol is shown.

Use the status to quickly determine the condition of your storage systems and if any actions must be taken. On the dashboard, the following statuses might be shown for a resource type:

Error (red)

An error status was detected on one or more of the monitored resources for a resource type. Error statuses represent serious problems on a resource. Resolve these problems as soon as possible. View the details of a resource to learn more about the status of its internal resources.

Unreachable (orange)
One or more of the monitored resources for a resource type are not responding. This status might be caused by a problem in the network or by a Storage Resource agent that is no longer running and did not communicate that it was shutting down.

Warning (yellow)
A warning status was detected on one or more of the monitored resources for a resource type. Warning conditions are not critical, but represent potential problems. View the details of a resource to learn more about the status of its internal resources.

Normal (green)
No warnings or errors were detected on the monitored resources for a resource type.

Unknown status (gray)
One or more of the monitored resources for a resource type have an unknown status. This status might occur if no data was collected about a resource. To change an unknown status, use a data collection job to collect status information about the resource.

For information about how status is determined, see "How status is determined" on page 375.

The number of statuses that were detected and acknowledged for each resource type
To view the total number of statuses for a resource type, move the mouse pointer over the related status symbol. Hover help shows the total number of occurrences for each status, and the number of statuses that were acknowledged by a Tivoli Storage Productivity Center user. An acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.

Viewing the overall capacity of monitored resources

The Capacity section of the dashboard contains a visual summary of the storage capacity for each type of resource that Tivoli Storage Productivity Center monitors.

To cycle through the capacity information for different types of resources, click the arrows next to the name of the resource type. Use the chart information to quickly determine the overall storage capacity for monitored resources and their available space. Hover over areas in the charts that are not labeled for more information.

Storage Systems

Storage systems can be configured with block storage, file storage, or a combination of block and file storage. The type of storage systems that are monitored by Tivoli Storage Productivity Center determines the capacity charts that are shown on the dashboard:

- If only storage systems with block data are monitored, then a single chart is shown with the following information: pool space, volume capacity, available pool space, and volume space that is assigned and unassigned.

- If only storage systems with file data are monitored, then a single chart is shown with the following information: pool space and details about the space on file systems that are exported and not exported to servers.
- If storage systems with block data and file data are monitored, then charts for both block and file data are shown.

Use the information in the following horizontal bar charts to view details about the capacity of all the monitored storage systems in your environment:

Storage systems that are configured with block data
The following capacity information is shown for monitored storage systems that are configured with block data:

Pool Space
Total amount of storage space in the storage pools or pools that are associated with monitored storage systems. For SAN Volume Controller, this value represents the amount of managed disk space that is assigned to the managed disk group. For XIV systems, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Volumes
Total amount of unique storage system volume space or virtual storage volume space on monitored storage systems. For thin-provisioned volumes, this value is the physical capacity of the volume. In an XIV system, this value is referred to as the real space, or physical ("hard") size, and not the potentially much larger virtual ("soft") size, that is allocated. For NetApp devices, this value represents the total amount of storage space on the Fibre Channel (FC) LUNs that are associated with the devices.

Available (Pool Space)
Amount of unused storage space in a pool.

Assigned (Volumes)
Total volume space within monitored storage systems that is mapped or assigned to host systems.

Unassigned (Volumes)
Total volume space within monitored storage systems that is not mapped or assigned to host systems.

Storage systems that are configured with file data
The following capacity information is shown for monitored storage systems that are configured with file data:

File System Capacity
Total amount of storage space on all the file systems on a storage system or filer.

Exported FS (File Systems)
The amount of storage space from file systems on storage systems that is exported to servers. An export or share is a file system that is made available to remote servers over a network.

Exported Used
The amount of exported file system space that is used by objects in the file systems, such as files and directories.

Exported Available
The amount of exported file system space that is unused and available for use by objects in the file systems.

Unexported FS (File Systems)
The amount of storage space in file systems on storage systems that is not exported to remote servers.

Restriction: For a NetApp filer, Tivoli Storage Productivity Center uses a Storage Resource agent to determine the file systems that are exported from the filer. Depending on the operating system of the computer where the Storage Resource agent is deployed, some exported file systems might not be detected during a probe. Specifically, a UNIX Storage Resource agent might not detect the file systems that are exported with CIFS; a Windows Storage Resource agent might not detect the file systems that are exported with NFS.

Because of this limitation, the value for **Unexported FS** might include both the space of unexported file systems and the space of exported file systems. To work around this limitation, set up multiple Storage Resource agents on different operating systems to probe the NetApp filer. For example, deploy a Storage Resource agent to a Windows system and another Storage Resource agent to a UNIX system. Then, use both Storage Resource agents to probe the NetApp filer.

Unexported Used
The amount of unexported file system space that is used by objects in the file systems, such as files and directories.

Unexported Available
The amount of unexported file system space that is unused and available for use by objects in the file systems.

Servers

The following capacity information is shown for the monitored servers in your environment: .

Total Disk
The total disk space for all the local and SAN-attached storage (from storage systems) on monitored servers.

Storage Systems
The amount of storage space that is assigned to servers from storage systems. The value is only available when SAN-attached storage is assigned to servers.

Other (storage systems)
The amount of storage space on the local disks and other non-SAN-attached disks in servers.

Used Disk, Available Disk
The total amount of used and unused disk space on the local and SAN-attached storage for servers.

FS/LV (File Systems and Logical Volumes), Available FS (File System Space)
The amount of used and unused space on the file systems and logical volumes for servers.

Hypervisors

The following capacity information is shown for the monitored hypervisors in your environment:

Total Disk
The total disk space for all the local and SAN-attached storage (from storage systems) on hypervisors.

Storage Systems
The amount of storage space that is assigned to hypervisors from storage systems. The value is only available when SAN-attached storage is assigned to hypervisors.

Other (storage systems)
The amount of storage space on the local disks in the hypervisor.

Used Disk, Available Disk
The amount of used and unused disk space on the local and SAN-attached storage for the hypervisor.

FS (file system), Direct (direct-attached storage)
The proportion of file system space to direct-attached storage space that is used on the local and SAN-attached storage for the hypervisor.
Direct-attached storage is storage that not managed by the hypervisor, but is assigned directly to a virtual machine on the hypervisor.

FS/LV (File Systems and Logical Volumes), Available FS (File System Space)
The amount of storage space on file systems and logical volumes that is used and available.

Viewing the status of performance monitors

A performance monitor is a data collection job that gathers performance data about monitored storage systems and switches. The Performance section of the dashboard shows a status summary of the performance monitors that are active or were run in your environment. This summary includes the number of performance monitors that have the following statuses: Running, Running with problems, Failed, and Not running.

To view detailed information about all performance monitors, click **View all monitors**.

Viewing the status of alerts that were detected on monitored resources

Alerts are generated when Tivoli Storage Productivity Center detects certain conditions or events on monitored resources. The Unacknowledged Alerts section of the dashboard shows a status summary of the unacknowledged alerts that were detected in your environment during a specified time period. This summary includes the number of alerts that have the following statuses: Critical, Warning, and Informational. An unacknowledged alert is an alert that was not reviewed or resolved by a Tivoli Storage Productivity Center user.

Use the following actions to manage the display of statuses for unacknowledged alerts:

- To select the time period in which you want to view alert statuses, click the arrows to cycle through the choices. You can view the status of all unacknowledged alerts that were detected in the last hour, last day, or last week.

- To view more detailed information about alerts, click **View all alerts** to access the Alerts page.

Viewing information about top-level resources

Use a resource list page as a central location for viewing information about all the monitored resources of a specific type. Resources that have a list page include storage systems, volumes, pools, servers, hypervisors, switches, and fabrics.

1. In the navigation pane, hover the mouse pointer over the resource icon for the type of resource that you want to view. You can use the icons to access the following types of resources:
   - **Storage Resources**: storage systems, volumes, pools
   - **Server Resources**: servers, hypervisors
   - **Network Resources**: switches, fabrics
2. Select the type of resource that you want to view. For example, if you want to view information about servers, select **Server Resources** > **Servers**. A list of monitored servers is displayed on Servers page.
3. Locate the resource that you want to view in the list. Each column in the list shows information about the resource.
4. Optional: To view key attributes of a specific top-level resource without leaving the page, right-click the resource and select **View Properties**. A properties notebook is displayed. Information in a properties notebook can include storage statistics, hardware attributes, volume assignments, disk assignment, and performance metrics.
5. Optional: To view more detailed information about a specific top-level resource, including its internal and related resources, right-click the resource and select **View Details**. A resource details page is displayed. This page includes information about storage statistics, alerts, jobs, internal resources, and related resources.

A resource list page shows information about the monitored resources for the selected type. You can view the following information on this page:

- The overall status of the resources for a specific type
- Key information about the resources
- Alerts that were generated for the resources
- Jobs that were run for the resources

Tip: To view descriptions of the information that is available on a resource list page, click the **Help** icon in the upper-right corner of the page.

Related reference:

"Resources that you can monitor in the web-based GUI" on page 258
You can use the web-based GUI to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

Viewing information about internal and related resources

Use resource details pages to view information about a specific top-level resource and its internal and related resources.

1. In the navigation pane, hover the mouse pointer over the resource icon for the type of resource that you want to view. The icons provide access to the following types of resources:
   - **Storage Resources**: storage systems, volumes, pools
   - **Server Resources**: servers, hypervisors
   - **Network Resources**: switches, fabrics
2. Select the type of the resource that you want to view. A resource list page shows the monitored resources for the selected type. For example, if you want to view a storage system, select **Storage Resources** > **Storage Systems**. A list of monitored storage systems is displayed on the Storage Systems page.
3. In the list of resources, right-click the resource that you want to view and select **View Details**. A resource details page is displayed. This page includes information about storage statistics, alerts, jobs, internal resources, and related resources.
4. To view information about the internal resources of the top-level resource, click the name of the internal resource in the **Internal Resources** section. A list of the internal resources is displayed. Each column in the list shows information about the resources. For example, on the details page for a server, click **Controllers** to view a list of the controllers that are associated with the server.
5. Optional: To view key attributes of a specific internal resource without leaving the page, right-click the resource and select **View Properties**.
6. To view information about resources that are related to the top-level resource, click the name of the related resource in the **Related Resources** section. A list of the related resources is displayed. For example, on the server details page, click **Storage Systems** to view a list of the storage systems that have at least one disk (volume) assigned to the server.
7. Optional: To view key attributes of a specific related resource without leaving the page, right-click the resource and select **View Properties**.
8. Optional: To view the resource details page for a related resource, right-click the resource and select **View Details**.

A resource details page shows information about a top-level resource. The information is available in the following links and sections on the page:

- **Overview**: The overall capacity and storage usage for a resource.
- **Properties**: Key information about a resource.
- **Alerts**: Alerts that were generated for a resource.
- **Jobs**: Jobs that were run for a resource.
- **Internal Resources**: Information about the components that exist in a higher-level resource.
- **Related Resources**: Information about the resources that are related to another resource. A related resource is external to a higher-level resource, but is related to it through assigned storage, a network connection, or virtual hosting.

Tip: To view descriptions of the information that is available on a resource details page, click the **Help** icon in the upper-right corner of the page.

Related reference:

"Resources that you can monitor in the web-based GUI" on page 258
You can use the web-based GUI to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

Display differences in the interfaces

Tivoli Storage Productivity Center stores information about monitored resources in a database repository. You can view this information in the stand-alone GUI or in the web-based GUI. Depending on the interface that you use, there might be some differences in how the information is presented.

The web-based GUI does not show a fabric for a physical, Cisco SAN

The stand-alone GUI shows a "fabric" for a physical, Cisco SAN, whereas the web-based GUI does not show a "fabric" for a physical, Cisco SAN.

- In the stand-alone GUI, Tivoli Storage Productivity Center automatically generates a worldwide name (WWN) to represent a physical, CISCO SAN as a fabric. This WWN typically starts with the numbers "3000". The WWN for the fabric is shown on the **Topology** > **Fabrics** panel and in reports for fabrics.
- In the web-based GUI, a fabric is not shown for physical, Cisco SANs on the Fabrics page. Only the virtual fabrics, physical switches, and logical switches for Cisco Fibre Channel switches are shown in the web-based GUI.

The number of unidentified endpoints on fabrics is different from the stand-alone GUI

The web-based GUI and stand-alone GUI might display a different number of unidentified endpoints on a fabric:

- In the web-based GUI, unidentified endpoints are labeled as **Unknown Resources** and each endpoint *port* is considered one **Unknown Resource**.
- In the stand-alone GUI, unidentified endpoints are labeled as **Other** and each endpoint *node* is considered one **Other**.

For example, an endpoint with two ports is shown as one **Other** resource in the stand-alone GUI, but in the web-based GUI it is shown as two **Unknown Resources**.

Pool space for NetApp filers

For NetApp filers that are configured with block data, the values for pool space that are shown in the web-based GUI are different from the values that are shown in a storage management tool for NetApp filers, such as FilerView. The values that are different include total pool capacity and available pool space.

This discrepancy occurs because the pool space information that is collected by SMI-S agents for Tivoli Storage Productivity Center includes reserved space for the Write Anywhere File Layout (WAFL) file system and the Snapshot session type. Tools such as FilerView do not include this additional space.

Incorrect values for storage systems in the web-based GUI after upgrade

When you upgrade Tivoli Storage Productivity Center from Version 4.2.x to Version 5.2, some columns on the Storage Systems page might display incorrect values. For example, the Raw Disk Capacity column, the Disks column, and the Volumes column might display a value of 0. To resolve this issue, you must run a probe for the resource when the upgrade completes. When the probe successfully completes, the columns display the correct values.

Monitoring the status of resources

You can use the web-based GUI to monitor the status of resources and identify potential problem areas in a storage environment. Monitored resources include top-level resources and their internal and related resources.

Use status information to quickly determine the condition of your storage and if any actions must be taken. Tivoli Storage Productivity Center provides a number of different statuses to help you determine the condition of resources. Each status is represented by an icon in the web-based GUI.

Table 54. Possible statuses of monitored resources

Icon	Status	Description
	Error	A serious problem was detected on a resource or on its internal resources. Resolve these problems as soon as possible.
	Error - Acknowledged	An Error status was detected and acknowledged. An Error - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored. An acknowledged status is not used when determining the status of related, higher-level resources. For example, if the status of a volume is Error, the status of the related storage system is also Error. If the Error status of the volume is acknowledged, its status is not used to determine the overall status of the storage system.
	Unreachable	A resource is not responding to requests from the Tivoli Storage Productivity Center server. This status might be caused by a problem in the network or by a Storage Resource agent that is no longer running and did not communicate that it was shutting down.
	Unreachable - Acknowledged	An Unreachable status was detected and acknowledged. An Unreachable - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored. An acknowledged status is not used when determining the status of related, higher-level resources. For example, if the status of a controller is Unreachable, the status of the related server is also unreachable. If the Unreachable status of the controller is acknowledged, its status is not used to determine the overall status of the server.
	Warning	A Warning status represents potential problems on a resource or on its internal resources. This status is not critical.

Table 54. Possible statuses of monitored resources (continued)

Icon	Status	Description
	Warning - Acknowledged	A Warning status was detected and acknowledged. An Warning - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored. An Warning - Acknowledged status is not used when determining the status of related, higher-level resources. For example, if the status of a controller is Warning, the status of the related server is also Warning. If the Warning status of the controller is acknowledged, its status is not used to determine the overall status of the server.
	Normal	No warnings or errors were detected on a monitored resource.
	Unknown	A resource is known to Tivoli Storage Productivity Center but is not monitored. To change an Unknown status, use a data collection job to collect status information about the resource.

For top-level resources, use the resource details page to view the status of their internal resources.

How status is determined

Top-level resources that are monitored in the web-based GUI include storage systems, servers, hypervisors, fabrics, and switches. The status of a top-level resource is determined by its consolidated status and the propagated status of its internal resources.

A consolidated status is the status that is reported by the top-level resource for its own hardware. For example, if the cooling fans in a storage system are stopped and the internal temperature is too high, that status is reported to Tivoli Storage Productivity Center.

A propagated status is the combined statuses of the monitored, internal resources for a top-level resource. For example, if a Warning status is detected on a volume in a storage system, that Warning status is used to determine the overall status of the storage system.

Tivoli Storage Productivity Center uses the values for consolidated status and propagated status to calculate the overall status of a top-level resource. The statuses of the following internal resources are used to calculate the status of top-level resources.

Table 55. Internal resources that are used to determine the status of top-level resources

Top-level resource	Internal resources that are used to determine the status of a top-level resource
Fabric	, Blades, and ports on blades Ports that are not blades Switches

Table 55. Internal resources that are used to determine the status of top-level resources (continued)

Top-level resource	Internal resources that are used to determine the status of a top-level resource
Hypervisor	Controllers Ports on controllers Virtual machines
Server	, Controllers, and ports on controllers Volume groups Monitored directories
Storage system	Clusters Controllers Disks External disks Modules Network shared disks Nodes Pools, and the internal resources that are associated with pools: • Volumes • RAID arrays • Managed disks Ports
Switch	Blades Logical switches Ports

Tivoli Storage Productivity Center uses the following statuses of internal resources to help calculate the status top-level resources:

- Unknown
- Normal

- Error
- Warning
- Unreachable

Tip: Statuses that are acknowledged are not used to calculate the overall status of a top-level resource.

Internal resources for a top-level resource might have different statuses. Tivoli Storage Productivity Center uses the most severe status of an internal resource when determining the overall status for a top-level resource. For example, in a storage system, a port might have an Error status, a pool might have a Warning status, and a number of controllers might have an Unknown status. In this case, the overall status for the storage system is Error because it is the most severe status that was detected on the internal resources.

Statuses for internal resources are propagated in the following way to determine the status of a top-level resource:

- If an internal resource has an Error status, then the related top-level resource has an Error status.
- If an internal resource has an Unreachable status, and no other internal resources have an Error status, the related top-level resource has an Unreachable status.
- If an internal resource has a Warning status, and no other internal resources have an Error or Unreachable status, the related top-level resource has a Warning status.
- If an internal resource has a Normal status, and no other internal resources have an Error, Unreachable, or Warning status, the related top-level resources has a Normal status.
- If an internal resource has an Unknown status, and no other internal resources have an Error, Unreachable, Warning, or Normal status, the related top-level resource has an Unknown status.

The following table shows some of the possible combinations of statuses and their resulting, propagated status.

Table 56. Propagation of the statuses for internal resources in the web-based GUI

Error	Unreachable	Warning	Normal	Unknown	Resulting status for top-level resource
				X	Unknown
			X		Normal
			X	X	Normal
		X			Warning
		X		X	Warning
		X	X	X	Warning
	X				Unreachable
	X			X	Unreachable

Table 56. Propagation of the statuses for internal resources in the web-based GUI (continued)

Error	Unreachable	Warning	Normal	Unknown	Resulting status for top-level resource
	X		X	X	Unreachable
	X	X	X	X	Unreachable
X					Error
X				X	Error
X			X	X	Error
X		X	X	X	Error
X	X	X	X	X	Error

Viewing the overall status of resources

Use the dashboard to monitor the overall status of monitored resources and identify potential problem areas in a storage environment. The dashboard shows the overall status of storage systems, servers, hypervisors, fabrics, and switches, and the total number of different statuses that were detected on the individual resources for each type.

1. In the navigation pane, click **Home** > **Dashboard**. Information about the overall status of top-level resources is shown on the resource diagram.
2. View the status icon that is shown with each resource icon to determine the overall status of a resource type. The status icon represents the most critical status that was detected on the resources for a specific type.

   For example, if Tivoli Storage Productivity Center monitors 20 storage systems, and an error condition was detected for one of those storage systems, the following status icon is shown with the storage systems icon:

   Storage Systems (20)

   If no Error, Warning, or Unreachable statuses were detected on the monitored storage systems, the following status icon is shown with the storage systems icon:

   Storage Systems (20)
3. Move the mouse pointer over the status icon for a resource type to view the number of statuses that were detected and acknowledged for that resource type. Hover help shows the total number of occurrences for each status, and the number of statuses that were acknowledged by a Tivoli Storage Productivity Center user.

   For example, move the mouse pointer over the status symbol for fabrics to view the number of monitored fabrics that have an Error, Unreachable, Warning, Normal, or Unknown status. If five fabrics have an Unreachable status, three fabrics have a Normal status, and no fabrics have an Error or Warning status, the following information is shown when you hover over the status icon for fabrics:

0 Error

5 Unreachable

0 Warning

3 Normal

Related reference:

"How status is determined" on page 375
Top-level resources that are monitored in the web-based GUI include storage systems, servers, hypervisors, fabrics, and switches. The status of a top-level resource is determined by its consolidated status and the propagated status of its internal resources.

Viewing the status of top-level resources

Use a resource list page as a central location for viewing the propagated statuses of top-level resources. Resources that have a list page with status information include storage systems, volumes, pools, servers, hypervisors, switches, and fabrics.

1. In the navigation pane, select the type of top-level resource that you want to view. For example, if you want to view the status of monitored switches, select **Network Resources** > **Switches**.
2. To view the total number of different statuses that were detected for the type of resource, check the status icons at the top of the resource list page. The number next to the icon represents how many of each status was detected.

   For example, if 20 switches are monitored, and 10 have a Normal status while the other 10 have a Warning status, the following status information is displayed at the top of the page:

   10 Normal

   10 Warning

   0 Error
3. To view the status for specific resources, check the Status column in the list of resources. The value in the Status column represents the propagated status for a top-level resource.
4. Optional: Click the **Search** icon to filter the list so that only resources with a specified status are shown. You can use filtering to identify the resources with the most serious Error conditions in your environment.
5. Optional: Select **Status**.
6. Optional: In the **Choose Value** list, select the status on which you want to filter. The list of resources is automatically refreshed and only resources with the filtered status are shown. For example, select **Warning** to view only the resources on which a Warning status was detected.

Viewing the status of internal resources

Use a resource details page to view the propagated status of a top-level resource and the status of its internal resources.

1. In the navigation pane, select a type of top-level resource. For example, if you want to view the status of monitored storage systems and their internal resources, select **Storage Resources** > **Storage Systems**.

2. Locate the resource in the list and check the Status column. The value in the Status column represents the propagated status for the top-level resource.
3. To view the status of internal resources, right-click the top-level resource and select **View Details**. A resource details page shows the internal resources for the selected resource.
4. Check for any status icons that are associated with an internal resource. If Tivoli Storage Productivity Center detects a status other than normal, an icon is displayed for the most critical status.

   For example, if a server has three controllers with an Error status and two controllers with a Warning status, the Error status icon is shown for controllers in the **Internal Resources** section:

   Controllers (5) .
5. To view the status of specific internal resources, click the name of the internal resource in the **Internal Resources** section. A list of the monitored internal resources is displayed. For example, on the details page for a server, click **Controllers (5)** to view a list of the five controllers that are associated with the server.
6. Check the Status column to view the status for each internal resource.

   Tip: When viewing information about some internal resources, the Status column might show values that are more specific than Error, Unreachable, Warning, Unreachable, Unknown, and Normal. Use this additional status to determine the cause of the problem for an internal resource.

   For example, in the **Internal Resources** section on the details page for a storage system, an Error status is shown for managed disks:

   Managed Disks (5)

   This status means that one or more managed disks have an Error condition. When you view the list of managed disks, the Status column might show an Offline status for a specific managed disk. When an Offline status is detected, the web-based GUI shows it as an Error status on the details page for the storage system.

Acknowledging the status of resources

The web-based GUI shows Error, Unreachable, and Warning statuses for top-level resources and internal resources. Sometimes these statuses might represent problems that commonly occur but can be ignored. In such cases, you can acknowledge the resources so that they are not used when propagating statuses for top-level resources or when calculating the total number of Error, Warning, and Unreachable conditions that were detected.

Acknowledging the status of top-level resources

The web-based GUI shows a summary of the statuses for top-level resources on the dashboard and resource list pages. You can acknowledge the status of a top-level resource so that its status is not shown as Error, Unreachable, or Warning on those pages.

1. In the navigation pane, select the type of resource that has a status you want to acknowledge. For example, if you want to acknowledge the status of a switch, select **Network Resources** > **Switches**.

2. In the list of resources, locate the resource with the Error, Unreachable, or Warning status that you want to acknowledge.
3. Right-click the resource and select **Acknowledge status**. A window shows that the status for the specified resource was acknowledged.
4. Optional: To acknowledge multiple resources at the same time, press Ctrl and click each resource. To select a series of resources, select the first resource, and then press Shift and click the last resource. When the resources are selected, right-click a row in the list and select **Acknowledge status**.
5. Click **Close**.

When you acknowledge the status of a top-level resource, the following actions are taken on pages in the web-based GUI:

- The dashboard:
  - The status of the acknowledged resource is not used to determine the status icon that is displayed for that resource on the resource diagram. For example, if you acknowledge the Unreachable status for a switch, that Unreachable status is not used to determine the status icon that is displayed for the switch icon on the dashboard.
  - The acknowledged status is shown as **Acknowledged** when you point to the status icon for a resource. For example, if 10 switches have an Unreachable status and you acknowledge the status for one of them, the following information is shown when you hover the mouse over the switch icon:

    9 Unreachable (1 Acknowledged)

    Restriction: The dashboard shows status icons and hover information for servers, hypervisors, switches, fabrics, and storage systems. The statuses of pools and volumes are used to calculate the status of these top-level resources, so their individual statuses are not shown on the dashboard.
- The resource details page:
  - The status icons at the top of the page are updated. For example, if 10 switches have an Unreachable status, and you acknowledge one of them, the total number Unreachable statuses decreases by one and the number of Unreachable - Acknowledged statuses increases by one. The following information is updated at the top of the Switches page:

    9 Unreachable **1 Unreachable - Acknowledged**
  - The value in the **Status** column for the resource is updated to show an acknowledged status. For example, if you acknowledge the Unreachable status for a switch, the following value is shown in the **Status** column:

    Unreachable

Related reference:

"Resources that you can monitor in the web-based GUI" on page 258
You can use the web-based GUI to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

Acknowledging the status of internal resources

The web-based GUI uses the status of internal resources to determine the status of top-level resources. Sometimes these statuses might represent problems that commonly occur or are known, but can be ignored. In such cases, you can acknowledge the status of an internal resource so that it is not used when determining the status of a top-level resource.

For example, if the status of a volume is Error, the status of the related storage system is also Error. If the Error status of the volume is acknowledged, its status is not used to determine the overall status of the storage system. In this case, if the other internal resources of the storage system are Normal, then the status of the storage system is also Normal.

1. In the navigation pane, select the type of resource in which the internal resource is located. For example, if you want to acknowledge the status of internal resources for a switch, select **Network Resources** > **Switches**.
2. Right-click a resource and select **View Details**. For example, if you want to acknowledge the status of a port on a switch, right-click that switch and select **View Details**.
3. In the **Internal Resources** section of the resource details page, click the name of the internal resource. For example, if you want to acknowledge the status of a port on a switch, click **Ports**.
4. Right-click the resource and select **Acknowledge status**. A window shows that the status for the specified resource was acknowledged. For example, if you want to acknowledge the status of a port, right-click the port and select **Acknowledge status**.
5. Optional: To acknowledge multiple resources at the same time, press Ctrl and click each resource. To select a series of resources, select the first resource, and then press Shift and click the last resource. When the resources are selected, right-click a row in the list and select **Acknowledge status**
6. Click **Close**. The status of the related, top-level resource might take several minutes to update.

When you acknowledge the status of an internal resource, the following actions occur:

- The status of the internal resource is no longer used to determine the status of the related, top-level resource.

  For example, if you acknowledge the Error status for a controller on a server, the status of that controller is not used when determining the overall status of the server.
- The status icons at the top of the page for the internal resource are updated.

  For example, if you acknowledge the Error status for a controller on a server, the total number of Error statuses decreases by one and the number of Error - Acknowledged statuses increases by one on the Controllers page. If originally there were five controllers with Error statuses, the following updated status information is shown at the top of the page:

  4 Error **1 Error - Acknowledged**

Related reference:

"Resources that you can monitor in the web-based GUI" on page 258
You can use the web-based GUI to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

Monitoring the performance of resources

Tivoli Storage Productivity Center can collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations. Metrics can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage. Threshold violations alert you when the performance of a monitored resource falls outside of a specified range.

In the web-based GUI, you can use the performance monitoring to complete the following tasks:

- Measure, compare, and troubleshoot the performance of switches, storage systems, and their internal resources.
- Review the threshold violations and alerts that were triggered when the performance if a resource fell outside of a specific range.
- Customize views of performance so that you can analyze specific resources and metrics during time ranges that you specify.
- View performance information in a chart or table format to help you quickly identify where and when performance issues are occurring. The chart is a visual representation of how the performance of resources trend over time.
- Drill down into resources to view detailed information about the performance of internal and related resources. For example, if a SAN Volume Controller is shown in the chart, you can quickly view and compare the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.
- Implement server-centric monitoring of SAN resources without requiring a Storage Resource agent. When you add an agentless server, Tivoli Storage Productivity Center correlates the server with the known host connections on monitored storage systems. If matches are found, you can view the amount of storage that is assigned to the server, and trace that storage back to the storage system. You can then view details about the internal resources of the related storage system, including performance information.
- Export performance information to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application.

Before you begin

Before you can view performance information for resources, you must complete the following tasks:

Task	GUI	Learn more
Add storage systems and switches for monitoring by Tivoli Storage Productivity Center.	Web-based GUI†	For information about how to add these resources, see the following topics: • "Adding storage systems" on page 232 • "Adding fabrics and switches" on page 241
Ensure that performance data is being collected for a resource. Typically, data collection is scheduled when resources are added for monitoring.	Web-based GUI	For information about verifying data collection and scheduling performance monitors, see the following topics: • "Verifying that a performance monitor is running for a resource" on page 301 • "Creating performance monitors" on page 300 .
Define performance thresholds to monitor if the performance of a resource falls outside a specified range and might represent a potential problem. When you define a performance threshold for a resource, select a specific metric that you want to measure. For example, you can define a threshold that notifies you when the total I/O port rate for a storage system falls outside a specified range.	Stand-alone GUI	For information about how to define performance thresholds, see "Defining performance thresholds" on page 397.
Tip: † To add IBM SONAS systems, you must use the stand-alone GUI.		

Related reference:

"Performance thresholds" on page 1690
Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

"Performance metrics" on page 1701
Tivoli Storage Productivity Center can collect information about the performance of monitored storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks.

Viewing performance information

View the performance of storage systems, switches, and their internal resources. You can also view the performance of resources that violated a performance threshold and generated an alert.

Before you view the performance of resources, ensure that performance data was collected for those resources during the time ranges that you want to analyze. Tivoli Storage Productivity Center uses performance monitors to collect metrics for measuring the performance of storage systems and switches. For information about

how to verify that performance monitors were run, see “Verifying that a performance monitor is running for a resource” on page 301.

You can access the performance view for a resource from different locations in the web-based GUI:

Task	Accessing the performance view in the web-based GUI
View the performance of specific storage systems, switches, and their internal resources.	1. In the navigation pane, select the type of resource that you want to view. For storage systems, select **Storage Resources** > **Storage Systems**. For switches, select **Network Resources** > **Switches**. 2. Right-click a resource and select **View Performance**. By default, the following information is displayed in a chart for each type of resource: • For storage systems, the top 5 storage systems with the highest total I/O rate are shown. • For switches, the top 5 switches with the highest total port data rate are shown. 3. To view the performance of an internal resource, right-click the top-level resource in the chart legend and select the option for the internal resource that you want to view. For example, for a SAN Volume Controller, right-click its name in the chart legend and select **Mdisk Performance** to view the performance of its managed disks. For managed disks, the top 5 contributors to the storage system workload are shown on the chart. The default metrics are total back-end I/O rate and total back-end response time. **Tips:** • To add or change the metrics that are displayed for a resource, click the select metrics icon and select the metrics to display. • You can also view the performance of internal resources on the details page for the related, top-level resource. For more information, see Table 57 on page 388. • To view the performance of multiple resources at the same time, press Shift or Ctrl and click the resources that you want to view. Press Shift and click to select consecutive rows in the list of resources; press Ctrl and click to select non-consecutive rows. Then, right-click any of the selected resources and select **View Performance**. Each resource is represented by a separate line in the performance chart and separate rows in the table view.

Task	Accessing the performance view in the web-based GUI
View the overall performance of resources of the same type.	1. In the navigation pane, select the type of resource that you want to view. For storage systems, select **Storage Resources** > **Storage Systems**. For switches, select **Network Resources** > **Switches**. 2. Click the **Performance** tab. By default, the following information is displayed in a chart: • For storage systems, the top 5 storage systems with the highest total I/O rate are shown. • For switches, the top 5 switches with the highest total port data rate are shown. **Tip:** To add or change the metrics that are displayed for a resource, click the select metrics icon and select the metrics to display.
View the threshold violations that were detected during data collection and the performance of the resources where the violation occurred. You can use this method to view all threshold violations, not just the violations that generate alerts.	For information about how to view threshold violations, see "Viewing threshold violations" on page 400.
View the alerts that were generated by threshold violations and the performance of the resources where the violation occurred. Use this method to only view the alerts that were generated by threshold violations.	For information about how to view the alerts that are generated by threshold violations, see "Viewing alerts for threshold violations" on page 401.
View the performance of a resource that is included in a specific performance monitor.	1. In the navigation pane, click **Home** > **Performance Monitors**. 2. In the Name column, locate the name of the resource that you want to view. 3. Right-click the performance monitor and select **View Performance**.
View predefined reports and create custom reports that show the performance of resources. In Tivoli Common Reporting, you can view reports about storage systems, components of storage systems, switches, and switch ports.	1. In the navigation pane, hover the mouse pointer over the **Reporting** icon. 2. Select **View predefined reports** or **Create performance reports**. 3. To view predefined reports, navigate to the performance report that you want to view and click the name of the report. To create a custom report about the performance of resources, select **Reporting** > **Create performance reports**.

In the web-based GUI, you can also view the performance of internal resources on the details pages for the related, top-level resources. For example, you can view the performance of the controllers for a storage system on the details page for that

storage system. To access the performance view for an internal resource on a details page, select the type of internal resource and click the **Performance** tab.

The following information summarizes the internal resources and locations where you can access a performance view.

Table 57. Internal resources and locations where you access the performance view

Internal resource	Locations in the GUI where you can access the performance view
Pools	• **Storage Resources** > **Pools** • **Storage Resources** > **Storage Systems** > **Pools** Available only for pools that are associated with the following storage systems and virtualizers: System Storage DS series, ESS, SAN Volume Controller, the IBM Storwize family, and XIV systems.
Volumes	**Storage Resources** > **Volumes**
Controllers	**Storage Resources** > **Storage Systems** > **Controllers**
Disks	**Storage Resources** > **Storage Systems** > **Disks** Available only for disks that are associated with storage virtualizers such as SAN Volume Controller and the IBM Storwize family.
Host Connections	**Storage Resources** > **Storage Systems** > **Host Connections** Available only for host connections that are associated with the following storage systems and virtualizers: System Storage DS series, ESS, SAN Volume Controller, the IBM Storwize family, and XIV systems.
I/O Group	**Storage Resources** > **Storage Systems** > **I/O Group**
Modules	**Storage Resources** > **Storage Systems** > **Modules**
Nodes	**Storage Resources** > **Storage Systems** > **Nodes**
Managed Disks	**Storage Resources** > **Storage Systems** > **Managed Disks** Available only for nodes that are associated with storage virtualizers such as SAN Volume Controller and the IBM Storwize family.
Ports on storage systems	**Storage Resources** > **Storage Systems** > **Ports**
Ports on switches	**Network Resources** > **Switches** > **Ports**
RAID Arrays	**Storage Resources** > **Storage Systems** > **RAID Arrays** Available only for RAID arrays that are associated with System Storage DS series and ESS storage systems.

How performance information is displayed

The performance view is displayed when you view the performance of resources, which includes resources in threshold violations, alerts, and performance monitors. Information in the performance view is organized into two main sections: a chart or table and a legend.

Performance chart

The top section of the performance view shows information about the selected resources. You can view this information in the following formats:

Chart

The chart shows a visual representation of how the performance of a resource trends over time. Each line on the chart represents a metric and a resource. For example, if you select two metrics and three resources, six lines are shown on the chart. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.

Table

The table shows performance and asset information that is formatted into rows and columns. Each row represents a resource and a time stamp from the chart; each column represents a metric from the chart. For example, if two metrics and three resources are displayed on a chart, and each of the lines on the chart have 10 data points, 30 rows and 2 columns are shown in the table.

To view other metrics and asset information for a resource, right-click anywhere in the header row of the table and select extra columns. The type of resource determines the metrics and information that is available.

Performance chart legend

The bottom section of the performance view shows more information about the resources in the chart. The information is formatted into rows and columns. Each row represents a resource that was selected for the view. Each column provides asset and performance information about a resource.

Each metric on the chart is represented by two columns in the legend. The first column includes the weighted average of the metric for each resource in the selected time range. The second column represents the peak data point of the metric for each resource in the selected time range. For example, if a metric in the chart is Total I/O Rate, the following columns are displayed in the legend: Avg Total I/O Rate, Worst Total I/O Rate. The worst value can be the minimum or the maximum number, depending on the metric. For example, the worst value for an I/O rate metric is a maximum number, while the worst value for the Cache Hit Percentage metric is a minimum number.

You can also view information about the resources that are related to the resources in the legend. For example, if a SAN Volume Controller is listed in the legend, right-click it to view the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage. For more information about the resources that you can view in the chart legend, see “Resources in the chart legend” on page 394.

Tip:

- If you hide a resource in the performance chart, the row for that resource remains visible in the chart legend.
- In the chart legend, information about a resource represents the state of that resource during the selected time range. Specifically, this information does not represent the current state of a resource, but instead shows the final state of the resource during the selected time range.

  For example, if the current date is January 1, and you view a volume with the time range set to December 1 to December 7, the legend shows

the capacity of that volume on December 7. If the capacity of the volume was changed between December 7 and January 1, this historical capacity is different from the current capacity of the volume on January 1.

Related concepts:

"Customizing lists of resources, jobs, and alerts" on page 443
The web-based GUI uses lists to show information about resources, jobs, and alerts. The information in each list is organized into columns and rows. You can customize these lists by filtering information, sorting rows, and showing, hiding, and reordering columns.

Controls:

Each performance view includes controls for customizing how information is displayed. The type of view and the resources that you are viewing determine which controls are available.

Table 58. Controls for the performance view

Control	Description
 View chart	View performance information in a chart format. The chart shows a visual representation of how resource performance trends over time. Each line on the chart represents a metric and a resource. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window. Hover the mouse pointer over points on a line to view a snapshot of performance information at a specific time.
 View table	View performance information in a table format. Each row represents a resource and a time stamp. Each column represents a metric. You can view other metrics and information for a resource by right-clicking anywhere in the header row for the table and selecting more columns. The type of resource determines the metrics and information that is available. For information how to filter, sort, and customize the columns in a table, see "Customizing lists of resources, jobs, and alerts" on page 443.
 Metrics	Add metrics to the performance chart. On the Select Chart Metrics window, metrics are organized into three tabs: Volume Metrics, Disk Metrics, and Port Metrics. The metrics and tabs that are available depend on the type of resource that is being shown in the chart. The number next to the name of the tab represents the number of metrics that are currently selected from that tab. You can select multiple metrics at the same time and from different tabs, but you cannot include more than two unit types in the same view. For example, if you select metrics that use % and ops/s as units of measurement, you cannot select more metrics that use different units of measurement such as KiB/op or MiB/s. For a list of metrics that you can view for resources, see "Performance metrics" on page 1701.

Table 58. Controls for the performance view (continued)

Control	Description
 Hide and show resources	When the chart includes multiple resources, you can click the icon next to a resource to show only the line for that resource. Each icon is shown in a different color to match the color of line for the resource. You can also show and hide resources in the chart by selecting resources in the chart legend. To select multiple resources at the same time, press Shift or Ctrl and click those resources. Press Shift and click to select consecutive rows in the legend; press Ctrl and click to select non-consecutive rows. You can show up to 10 resources in a chart at the same time.

Table 58. Controls for the performance view (continued)

Control	Description
 Specify granularity	Determine the granularity of the data that is shown in a performance view. Granularity determines the points that are shown on the y-axis of a chart, the points on the lines in a chart, and the rows in a table. When you first view performance information, the default time range is the last 12 hours and the granularity matches the frequency of data collection by the performance monitor (sample). To change the granularity of chart or table, select one of the following options from the granularity menu: **(daily)** Set the granularity to show 1-day increments for performance information that is shown on the chart and table. This granularity is only available if the time range of the chart is more than 1 day. **(hourly)** Set the granularity to show one-hour increments for performance information that is shown on the chart and table. This granularity is only available if the time range of the chart is more than 1 hour. **(Sample)** Set the granularity to match the frequency of data collection by a performance monitor. For example, if a performance monitor is scheduled to collect sample data every 15 minutes, each point on the chart is shown in 15-minute increments. This granularity is only available if the time range of the chart is less than 2 weeks. **Restriction:** Each granularity has a maximum time range for displaying data. This time range is dynamically calculated based on the resources and metrics that are shown on the chart, and the data points per line. When the maximum range is exceeded, the granularity icon is disabled and a message indicates that there are too many lines on the chart. To avoid this problem, reduce the time range or the number of resources or metrics that are selected.

Table 58. Controls for the performance view (continued)

Control	Description
Specify a time range	The time range of the performance information is shown below the chart. You can change this range to display information for different times when data was collected. When you first access the performance view, the default time range is the last 12 hours. Ensure that performance data was collected during the time range that you select. If data was not collected during the time range, the chart and table are blank. If data collection was interrupted during the time range, the chart and table show gaps for the time increments when data was not collected. For example, if you select a time range for the last 7 days, but data was not collected on days 4 and 5, the lines in the chart do not show data for days 4 and 5.
 Hide and show chart controls	The performance view lists the names of resources that are being shown in the chart and the controls for customizing that chart. Click the vertical line to hide the controls and display the chart in the entire window. Click the line again to show the controls.
 Export information about the chart to a file	Export information on a performance view to a CSV file.
 Open the performance view in a separate web browser window	Open a duplicate of the current performance view in a separate web browser window. You can change the information that is displayed in this separate window while retaining the original performance view for comparison.
 Synchronize the time range across all the open performance views	Synchronize the time range across all the performance views that are displayed in separate browser windows. Use this action when you change the time range in a performance view and want to apply the same time range to the other performance views. **Restriction:** This action affects only the performance views that are displayed in separate browser windows. The time range of the performance view that is shown as part of the main window for the GUI is not affected. For example, you can view the storage system volumes that are assigned to a server. You can then open separate browser windows for the host connections, pools, and managed disks that are related to one of the volumes. If you change the time range in the performance view for host connections, click the synchronize button to apply the same time range to the views for pools and managed disks. The performance view of the original volume is not changed.

Table 58. Controls for the performance view (continued)

Control	Description
Actions for managing the resources in the chart legend	The chart legend in bottom section of the view shows more information about the selected resources. This information is organized into rows and columns, where each row represents a resource. When you select one or more resources in the legend, the following actions are available in the **Actions** menu: **View Properties** View key details about a resource, including asset, status, configuration, capacity, and performance information. **View** *resource* **Performance** View the performance of resources that are internal or related to a resource in the chart legend. For example, when you view the performance of a SAN Volume Controller, you can right-click it and view the performance of its internal resources, such as disks, volumes, ports, managed disks, and back-end storage. Information about an internal or related resource is shown in a separate web browser window. This window uses the same time range as the performance view in the main window of the GUI. **Restriction:** If performance data was not collected for an internal or related resource, that resource is not shown in the view. For example, if you right-click a storage system and select **Volume performance**, only the volumes for which performance data was collected are shown. For more information about the resources that you can view in the chart legend, see "Resources in the chart legend." **Key Metrics View** View a set of predefined metrics for resources in the chart legend. **Custom View** Customize the set of metrics that are shown for resources in the chart legend.

Resources in the chart legend:

In the performance view, performance metrics and related information for resources are shown in a chart and in the chart legend. For resources in the chart legend, you can open separate performance views for their internal or related resources. The separate performance views use the same time range as the performance view in the main window of the GUI.

For example, when you view the performance of a SAN Volume Controller, you can right-click it in the chart legend to open a performance view for its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.

The type of resource determines the internal and related resources that you can view. The following tables show the internal and related resources that are available when you right-click a resource in the chart legend.

Table 59. Resources that you can view for DS6000, DS8000, and ESS storage systems in the chart legend

Resource	Internal Resources	Related Resources
Storage system	Volume, RAID array, pool, controller, host connection, port	Storage virtualizer, switch
Volume	Storage system, RAID array, pool, controller, host connection	Managed disk
RAID array	Storage system, volume, pool, controller, host connection	None
Pool	Storage system, volume, RAID array, controller, host connection, port	None
Controller	Storage system, volume, RAID array, pool, host connection	None
Host connection	Storage system, volume, RAID array, pool, controller	None
Port	Storage system	Switch port

Table 60. Resources that you can view for SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified storage systems in the chart legend

Resource	Internal Resources	Related Resources
Storage system	Volume, pool, node, I/O group, disk, managed disk, host connection, port	Back-end storage system, storage virtualizer, switch
Volume	Storage system, pool, node, I/O group, disk, managed disk, host connection	None
Pool	Storage system, volume, node, I/O group, disk, managed disk, host connection	None
Node	Storage system, volume, pool, I/O group, disk, managed disk, host connection, port	None
I/O group	Storage system, volume, pool, node, disk, managed disk, host connection, port	None
Disk	Storage system, volume, pool, I/O group, node, managed disk, host connection	None

Table 60. Resources that you can view for SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified storage systems in the chart legend (continued)

Resource	Internal Resources	Related Resources
Managed disk	Storage system, volume, pool, node, I/O group, disk, host connection	Back-end volume
Host connection	Storage system, volume, pool, node, I/O group, disk, host connection	None
Port	Storage system, I/O group, node	Switch port

Table 61. Resources that you can view for XIV systems in the chart legend

Resource	Internal Resources	Related Resources
Storage system	Module, pool, volume, host connection, port	Storage virtualizer, switch
Module	Storage system, pool, volume, host connection, port	None
Pool	Storage system, module, volume, host connection	None
Volume	Storage system, module, pool, host connection	Managed Disk
Host connection	Storage system, module, pool, volume	None
Port	Storage system, module	Switch port

Table 62. Resources that you can view for SMI-S BSP systems in the chart legend

Resource	Internal Resources	Related Resources
Storage system	Volume, port	Storage virtualizer, switch
Volume	Storage system	Managed disk
Port	Storage system	Switch port

Table 63. Resources that you can view for switches in the chart legend

Resource	Internal Resources	Related Resources
Switch	Port	Storage system, switch
Port	Switch	Storage system port, switch port

Table 64. Resources that you can view for servers in the chart legend

Resource	Internal Resources	Related Resources
Server	None	Storage system, volume, pool, controller, module, node, I/O group, disk, managed disk, host connection, switch, storage system port, switch port

Defining performance thresholds

Define performance thresholds to detect if the performance of the internal resources within a storage system or switch falls outside a specified range and might represent a potential problem. You can specify to be automatically alerted when thresholds are violated.

Thresholds are checked when performance data is collected. Before you can check a resource for threshold violations, ensure that performance monitors are collecting data about that resource. Tivoli Storage Productivity Center uses performance monitors to collect metrics for measuring the performance of resources. For information about how to verify that performance monitors are being run, see "Verifying that a performance monitor is running for a resource" on page 301.

When you define a performance threshold, select a specific metric that you want to measure and its boundary values. When the performance of a resource falls outside the boundary values, a threshold violation occurs. For example, you can define a performance threshold that alerts you when the overall back-end response times for managed disks on a SAN Volume Controller exceed a certain value. The overall back-end response time is a metric that measures the average number of milliseconds that it takes to service each I/O operation on a managed disk.

Tip: Thresholds are checked against all the internal resources of a specific type; you cannot apply them only to specific internal resources. For example, you cannot set a threshold for port A and a different threshold for port B, if both are in the same storage system. You can set a threshold only for all ports in that storage system.

To define a performance threshold, complete the following steps:

1. Start the stand-alone GUI.
2. Choose one of the following options to determine which type of resource to check for a threshold violation:

Option	Description
Storage systems	1. Expand **Disk Manager** > **Alerting**. 2. Right-click **Storage Subsystem Alerts** and select **Create Storage Subsystem Alerts**.
Switches	1. Expand **Fabric Manager** > **Alerting**. 2. Right-click **Switch Alerts** and select **Create Switch Alerts**.

3. In the **Condition** menu, select the performance threshold that you want to monitor for a resource. Each threshold is related to a metric that Tivoli Storage Productivity Center collects about a resource. The value for a metric is checked at each collection interval of a performance monitor to determine if a threshold is violated. Threshold conditions have names that end with the word "Threshold".
4. Specify the boundary values that would cause a threshold to be violated. You can specify boundary values in the following fields:

   Critical Stress, Warning Stress
   In general, stress boundaries are upper-threshold boundaries that are violated when performance equals or exceeds a value that you specify. A critical stress boundary represents a value that is the most severe

violation of a performance threshold. A warning stress boundary is considered less severe than a violation of a critical boundary. If either boundary is violated, a threshold violation occurs and an alert might be generated.

Warning Idle, Critical Idle
In general, idle boundaries are lower-threshold boundaries that are violated when performance falls below a value that you specify. A critical idle boundary represents a value that is the most severe violation of a performance threshold. A warning idle boundary is considered less severe than a violation of a critical boundary. If either boundary is violated, a threshold violation occurs and an alert might be generated.

Tips:

- For some metrics and performance thresholds, lower values indicate more stress and higher values indicate idle behavior. For example, a lower value in the critical stress boundary for the metric Cache Holding Time might indicate a performance problem.
- If you do not want to monitor violations for certain boundaries, leave those boundaries blank.

How to determine threshold values: Tivoli Storage Productivity Center provides recommended values for thresholds that do not vary much between environments. For example, the default critical and stress boundaries for the disk utilization threshold are 80% and 50%. However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine boundary values for a resource, collect performance data over a period of time to establish a base-line of the normal and expected performance behavior for that resource. After you determine a set of base-line values, define thresholds to trigger if the measured performance behavior falls outside the normally expected range.

5. In the **Alert Suppression** section, determine the conditions under which an alert is triggered when a threshold is violated and how repeated alerts are managed. For information about the options for alert suppression, see "Alert suppression actions" on page 1681.
6. In the **Triggered actions** section, determine the actions that occur as a result of a threshold being violated. For information about triggered actions, see "Triggered actions for alerts" on page 1682.
7. Choose one of the following to select the type of resource to check for a threshold violation:

Option	Description
Storage systems	1. Click the **Storage Systems** tab. 2. Select the storage systems that you want to check for violations of the threshold. 3. Click the right-arrow icon >> to add the storage systems to the threshold.

Option	Description
Switches	1. Click the **Switches** tab. 2. Select the switches that you want to check for violations of the threshold. 3. Click the right-arrow icon to add the switches to the threshold.

If the threshold condition is not available for a resource, that resource is not shown on the **Storage Systems** or **Switches** tab.

8. Save the threshold alert:
   a. In the menu bar, click **File** > **Save**.
   b. In the Save As window, type a name for the alert and click **OK**. An entry for the threshold alert is added to the navigation tree. The name of the alert includes a prefix that represents the user name of the user who created the alert. For example, if "user1" creates a threshold alert that is named "total_iorate" for a storage system, the entry is displayed as "user1.total_iorate" under the **Disk Manager** > **Alerting** > **Storage Subsystem Alerts** node. You can edit or rename the alert at any time.

Tivoli Storage Productivity Center checks the value for a metric at each collection interval of a performance monitor to determine if a threshold is violated. If a violation occurs, you can view information about that violation in the following locations in the web-based GUI:

- The **Threshold Violations** tab on the **Home** > **Performance Monitors** page. This tab automatically shows all the threshold violations that were detected.
- The **Home** > **Alerts** page. This pages shows the threshold violations that were defined to generate an alert.

In the stand-alone GUI, you can view information about threshold violations in Constraint Violation reports. To view alerts that are generated by threshold violations, expand **IBM Tivoli Storage Productivity Center** > **Alerting** > **Alert logs** in the navigation tree and click **All**.

Tip: If a performance monitor is already collecting data about a resource when you add, modify, or remove a threshold for that resource, changes are applied dynamically. You do not have to stop and restart the performance monitor to apply the changes. A confirmation message is recorded in the log of the performance monitor when the threshold is updated.

Related concepts:

"Starting the Tivoli Storage Productivity Center stand-alone GUI" on page 41
The stand-alone GUI contains functions for monitoring the condition of storage resources, and all the tools for managing data, disk, fabric, and replication.

Related reference:

"Performance metrics" on page 1701
Tivoli Storage Productivity Center can collect information about the performance of monitored storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks.

"Performance thresholds" on page 1690
Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

Viewing threshold violations and alerts

View the performance thresholds that were violated and the alerts that were generated as a result of violations.

To view threshold violations and alerts for a resource, you must collect performance data and define performance thresholds for that resource. For information about how to collect performance data, see "Creating performance monitors" on page 300. For information about how to define performance thresholds, see "Defining performance thresholds" on page 397.

A threshold violation is detected when a performance monitor discovers that the measured value of a metric falls outside of the specified threshold range for a resource. In the web-based GUI, you can view a list of all the threshold violations that are detected during data collection. You can also view all of the alerts that are generated for threshold violations.

Viewing threshold violations

View all the threshold violations that were detected during data collection. You can also open a threshold violation to view information about the resource where the violation occurred.

1. In the navigation pane of the web-based GUI, click **Home** > **Performance Monitors**.
2. Click the **Threshold Violations** tab.
3. Right-click the threshold violation that you want to view and select **View Threshold Violation**. In this view, a chart shows the performance of the resource that violated the threshold. The time range of the chart is 2 hours before and 2 hours after the violation occurred. The lines on the chart represent the following values:
   - The blue line represents the performance of the resource for the selected metric.
   - The yellow line represents the warning boundary value for the threshold.
   - The red line represents the critical boundary value for the threshold.
   - The dotted line represents the time when the violation occurred.

   For certain storage systems, a ranked list of the 25 volumes that were the top contributors to the workload on the resource where the threshold violation occurred might be displayed below the chart. This list is shown under the following conditions:
   - A stress boundary was violated. In general, stress boundaries are upper-threshold boundaries that are violated when performance equals or exceeds a specified value.
   - The threshold was violated on an internal resource other than a port.
   - The threshold was violated on any of the following storage systems:
     - DS6000
     - DS8000
     - XIV system
     - ESS
     - SAN Volume Controller
     - Storwize V3500
     - Storwize V3700
     - Storwize V7000
     - Storwize V7000 Unified

You can also view the servers or hosts to which those volumes are assigned.

4. Optional: To access the full performance view of a resource that violated the threshold, click the open performance view icon on the chart.

Viewing alerts for threshold violations

View the alerts that were generated by threshold violations. You can also open an alert to view information about the resource where the threshold violation occurred.

Tip: Not all threshold violations generate alerts. A performance threshold must be defined to generate an alert before an alert can be displayed when the threshold is violated. For information about how to define a performance threshold with an alert, see "Defining performance thresholds" on page 397.

1. In the navigation pane of the web-based GUI, click **Home** > **Alerts**.
2. Right-click the threshold violation that you want to view and select **View Alert**. In this view, a chart shows the performance of the resource that violated the threshold and generated an alert. The time range of the chart is 2 hours before and 2 hours after the violation occurred. The lines on the chart represent the following values:
   - The blue line represents the performance of the resource for the selected metric.
   - The yellow line represents the warning boundary value for the threshold.
   - The red line represents the critical boundary value for the threshold.
   - The dotted line represents the time when the violation occurred.

   For certain storage systems, a ranked list of the 25 volumes that were the top contributors to the workload on the resource where the threshold violation occurred might be displayed below the chart. This list is shown under the following conditions:
   - A stress boundary was violated. In general, stress boundaries are upper-threshold boundaries that are violated when performance equals or exceeds a specified value.
   - The threshold was violated on an internal resource other than a port.
   - The threshold was violated on any of the following storage systems:
     - DS6000
     - DS8000
     - XIV system
     - ESS
     - SAN Volume Controller
     - Storwize V3500
     - Storwize V3700
     - Storwize V7000
     - Storwize V7000 Unified

   You can also view the servers or hosts to which those volumes are assigned.
3. Optional: To access the full performance view of a resource that generated an alert, click the open performance view icon on the chart.

Viewing the data paths of resources

A data path shows the connectivity between host connections, servers, hypervisors, virtual machines, storage systems, and the fabrics through which they share storage. You can use the data path view to monitor status and pinpoint problem areas in selected data paths. This view includes graphical and tabular representations of the top-level resources in a data path.

If a resource does not map storage to other resources, or the resource is not assigned storage from other resources, the data path view is not available.

To view the data path for a resource, complete these steps:

1. In the navigation pane, select the type of resource that you want to view. You can view the data path for servers, hypervisors, and storage systems. For example, if you want to view the data path for a server, select **Server Resources** > **Servers**.
2. Right-click the resource and select **View Data Path**.

Depending on how resources are configured, a data path can show between three to six types of top-level resources. For example, a data path with three types of top-level resources might show the following relationships:

```
server <> fabric <> storage system
```

A data path with six types of top-level resources might show the following relationships:

```
virtual machine <> hypervisor <> fabric <>
storage virtualizer <> fabric <> storage system
```

Before you can view a data path, you must complete the following tasks:

- Add the resource that has the data path you want view. For information about how to add a resource, see "Adding resources" on page 257.

  Tip: Resources that are not monitored by Tivoli Storage Productivity Center are also shown in a data path view. However, detailed information about their status and internal resources is not available.
- Schedule a probe to collect data about your monitored resources to ensure that the most recent information is shown for a data path. For information about probes, see "Collecting asset and status data" on page 296.

You can use the data path view to complete the following actions:

- View the path of data that is shared between resources. The data path view shows only the resources that have storage mapped between them. For example, if you view the data path for a storage system, only the host connections that are assigned volume space from that storage system are shown. Host connections that are connected to the storage system, but do not have storage assigned storage to them, are not shown.
- View the fabrics through which resources in a data path are communicating.
- View the propagated status of the top-level resources that are in a data path.
- View the status of the internal resources for top-level resources that are in a data path.
- Customize the appearance of the data path view to suit the needs of your environment.
- Export the data paths view as an image or CSV file.

Information about a data path is organized into two views: a graphical view and a table view. The graphical view shows icons for connected resources and uses lines to represent the path of data through them. The table view includes a row for each resource in a data path. The row is divided into columns that contain information about the resource.

Resources that you can view in a data path

The **Topology View** tab in the data path view includes a visual representation of the data path for a top-level resource. In a data path, icons represent the connected resources and lines represent the path of data through them. For example, resources that are storage systems are shown with an icon that resembles a storage system.

The following table shows the icons for each type of resource that can be shown in a data path.

Table 65. Resource icons in a data path

Resource	Icon
Server (with agent)	
Server (without agent)[1]	
Hypervisor	
Virtual machine	
Fabric	
Virtual fabric	

Table 65. Resource icons in a data path (continued)

Resource	Icon
DS8000	
DS6000	
SAN Volume Controller	
Storwize V7000[2]	
XIV system	
ESS	
Other storage systems	
Unmanaged host connections	

Table 65. Resource icons in a data path (continued)

Resource	Icon
Tips and restrictions: • [1] When agentless servers and servers with agents are grouped together, the icon for servers with agents is used to represent that group. • [2] If a Storwize V7000 is not monitored by Tivoli Storage Productivity Center, but is connected to a monitored hypervisor, it is shown in the data path with the icon for a SAN Volume Controller. For example, when a virtual machine is monitored, a probe collects information about its disks. If the disks are attached to a SAN-attached storage system that is not monitored, Tivoli Storage Productivity Center uses the probe information to determine the type of storage system. If the storage system is a Storwize V7000, the icon for SAN Volume Controller is shown instead of the icon for a Storwize V7000. • Only the resources that are sharing storage in a data path are shown. • Servers and hypervisors that are not monitored by Tivoli Storage Productivity Center, but are assigned storage from a monitored storage system, are shown as unmanaged host connections. For example, if a storage system has a volume that is mapped to an unknown server, the server is represented by an unmanaged host connection in the view. • For storage systems with multiple host connections, the host connections are included in the same resource box if they communicate with the storage system through the same fabric. For example, if a storage system has volumes that are mapped to 20 host connections that use the same fabric, a single resource box represents the 20 host connections. • Storage systems can have volumes that are provisioned to more than one server in a clustered environment. Because all the clustered servers can be part of the data path, each server is shown in the view.	

The status of resources

Each resource in a data path is represented by a resource box. A resource box contains an icon that represents the overall status of a resource. This status is determined by the propagated status of the internal resources for the resource. For example, in a storage system, a port might have an Error status and a pool might have a Warning status. In this case, the Error icon is shown on the resource box for the storage system because the Error condition is the most severe status that was detected on its internal resources.

Some resources in a data path might be shown with an unknown status. Resources can have an unknown status if they not monitored by Tivoli Storage Productivity Center, but are connected to other, monitored resources. Tivoli Storage Productivity Center uses the information that it collects about monitored resources to determine basic information about the connected resources that are not monitored.

For example, if a host connection is assigned volumes, but Tivoli Storage Productivity Center is not monitoring the fabric through with the host connection communicates with the storage system, the fabric is shown with an unknown status.

Navigating in a data path

A data path might contain more resources that can be fit in the view at the same time. In these cases, you can scroll the view to show different parts of a data path.

1. In the **Topology View** tab, click the **Pointer** button .
2. Click an empty area in the view and drag the mouse pointer to scroll to different areas of a data path.
3. Optional: To view where you are in the overall topology of a data path, click the **Topology overview** button in the upper-right corner of the topology view. The red box shows the visible area within the overall topology view, and status icons are displayed for the resources in the data path.

Viewing information about top-level resources

You can view detailed information about the top-level resources in a data path. Top-level resources that can be shown in a data path include storage systems, servers, hypervisors, and fabrics.

1. To view key attributes of a specific top-level resource without leaving the **Topology View** tab, right-click the resource and select **View Properties**. A properties notebook is displayed. Information in a properties notebook can include storage statistics, hardware attributes, volume assignments, disk assignment, and performance metrics.
2. To view more detailed information about a specific top-level resource, including its internal and related resources, right-click the resource and select **View Details**. A resource details page is displayed and includes information about storage statistics, alerts, jobs, internal resources, and related resources.

Viewing information about internal resources

You can view information about the internal resources that are associated with the top-level resources in a data path. Internal resources are components that exist in a top-level resource.

1. In the **Topology View** tab, hover the mouse pointer over a top-level resource in the data path.
2. View the following information about internal resources:
   - The number of internal resources that are associated with a top-level resource. For example, if a server has 20 disks, the value **Disks (20)** is shown.
   - The overall status of the internal resources. For information about how the overall status of internal resources is determined, see "How status is determined" on page 375.

Zooming in

You can use the zoom buttons in the toolbar or the mouse to magnify areas in a data path.

1. To zoom in on a specific area in a data path, click the **Zoom to selected** button and click and drag the mouse pointer over the area.
2. Optional: To use the mouse to zoom in on a specific area in a data path, move the mouse pointer over the area where you want to zoom in and scroll the middle mouse button forward.
3. Optional: To zoom in on the area that is in the center of the view, click the **Zoom in** button .

4. Optional: To adjust the zoom level of the view so that all the resources in a data path are shown at the same time, click the **Fit to view** button .

Zooming out

You can use the zoom out button in the toolbar or the mouse to zoom out in the view.

1. To zoom out from the area that is in the center of the view, click the zoom out button .
2. Optional: To use the mouse to zoom out from a specific area in a data path, move the mouse pointer over the area where you want to zoom out and scroll the middle mouse button backward.
3. Optional: To adjust the zoom level of the view so that all the resources in a data path are shown at the same time, click the **Fit to view** button .

Changing the layout

You can change the layout of the resources in a data path. The layout determines how resources are arranged in the view.

1. Click the **Layout** button in the toolbar. The icon that is displayed for the **Layout** button represents the current layout of the data path.
2. Select one of the following layouts:

   Hierarchical
   Organize resources in a vertical format. Resources that use storage are shown at the top of the view. Resources that provide storage are shown at the bottom of the view.

   Tree
   Organize resources in a vertical and horizontal layout.

   Customize
   Customize the layout of a data path. When you customize a layout, use the hierarchical or tree layouts as the base and specify values for the associated options.
3. Optional: If you selected **Customize**, click **OK** to apply your settings.

Saving the display settings

You can save the display settings that are defined for a data path view. The next time that you view the data path for the resource, the same display settings are applied.

1. In the **Topology View** tab, click the **Save** button .
2. Optional: To use the keyboard to save the display settings, press Ctrl+S.
3. Click **Save Customizations**.

Printing a view

You can print the current view of a data path in the **Topology View** tab. The printing capability is not available in the **Table View** tab.

1. Click the **Print** button .
2. Specify the following options when printing:

- In the **Page margin** field, enter the amount of space between the edges of the page and the view. The default setting is 1 inch.
- In the **Maximum pages in width or height** field, specify the number of pages across which the content of a view is printed. This number determines the maximum number of pages that are printed. For example, if you enter **1**, the content of a view is printed on one page. If you enter **2**, the view is printed across four pages: two pages horizontally and two pages vertically.

3. Click **OK**.
4. Select the printer settings and click **OK** to print the view.

Exporting a data path to a file

You can save the contents of a data path to an image file or CSV file for offline viewing.

1. In the **Topology View** tab, right-click on an empty area in the data path.
2. Complete one of the following actions to select an output format for the file:
   - Select **Export as Image** to save the data path as a JPEG or PNG file.
   - Select **Export** to save the data path to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application. The information for a data path is organized into the following column in the CSV files: Name, Type, From, RelType, To.
3. Specify the location and name of the file.
4. Save the file.

Configuring for storage cloud provisioning

To take advantage of the simplified provisioning and optimization capabilities that are available by using the web-based GUI, some configuration is required.

To provision storage, you are required to specify only the storage capacity and storage quality that is required. After volumes are created, Tivoli Storage Productivity Center can analyze and optimize volume performance.

To configure Tivoli Storage Productivity Center for provisioning and optimization, you complete the following tasks:

1. Assign storage pools to tiers by setting the tier level of each storage pool. The tier levels are number tags that can reflect any tier structure in the environment.

   You must assign storage pools to tiers to later provision volumes that require a certain tier level and to optimize the placement of volumes. Storage pools do not need to be assigned to tiers to provision volumes that do not require a tier level, or to provision file shares.
2. Create service classes. A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.
3. Optional: Create capacity pools. You can restrict provisioning requests to a capacity pool. Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. For example, a capacity pool might contain the

storage resources that are allocated to a particular department or division of your business, or to a particular application.

Service class configuration

Configure service classes to represent the different levels or types of storage quality that you want to provision. When you are configuring service classes for your installation, consider whether certain storage requests must always be satisfied from a particular set of resources. Also, consider whether you want to allow users without administrator privileges to provision storage, and whether you want their provisioning requests to require administrator approval.

Setting access path host replacements

If you are going to use the VMware vSphere Web Client extension to provision file shares, verify whether you must provide Tivoli Storage Productivity Center with an access path host replacement.

You need to provide an access path host replacement if the cluster name of an IBM SONAS or Storwize V7000 Unified is not a valid host name that is registered with a Domain Name System (DNS) server.

Complete the following steps to verify cluster host names and to provide access host name replacements where necessary:

1. Open the **File Storage** tab of the Storage Systems page. From the navigator pane, select **Storage Resources** > **Storage Systems**, and click the **File Storage** tab.
2. For each IBM SONAS and Storwize V7000 Unified system, complete the following steps to determine whether you must specify a replacement host name.
   a. From the **File System** tab of the Storage System page, select the storage system, and then select **Actions** > **View Details**.
   b. From the list of File Internal Resources, click **Clusters**.
   c. Determine whether the cluster name is a valid host name that is registered with a DNS server. If the cluster name is registered with a DNS server, the access path that is returned by the Tivoli Storage Productivity Center will be valid. If the cluster name is not valid, the VMware vSphere Web Client extension will not be able to access the share unless you specify a host name replacement for the access path.
3. For each IBM SONAS and Storwize V7000 Unified system that requires a host name replacement, complete the following steps:
   a. From the **File Storage** tab of the Storage Systems page, select the storage system, and then select **Actions** > **View Properties**.
   b. In the properties notebook for the storage system, click **Edit**.
   c. In either the Custom Tag 1, Custom Tag 2, or Custom Tag 3 field, enter a fully qualified host name that is mapped to one or more interface nodes of the storage system.

      Requirement: For all storage systems that require a replacement host name, ensure that you set the replacement host name by using the same custom tag. For all IBM SONAS and Storwize V7000 Unified systems that do not require a replacement host name, ensure that the custom tag is empty.
4. When you are creating or modifying file-storage service classes that might be used by the VMware vSphere Web Client extension, set the access host

replacement to the custom tag that is set on the file storage systems. To set the access host replacement, complete these steps:

a. From the list of general properties in the Create Service Class wizard or the service class properties window, click **Advanced**.

b. On the **General** tab of the Advanced Properties window for the file-storage service class, set the access path host replacement to the custom tag that specifies the replacement host name. When the share is provisioned, the access path that is returned by Tivoli Storage Productivity Center includes the replacement host name if one is specified on the custom tag for the storage system. If the storage system does not specify a value for the custom tag, the cluster name is used in the access path.

Tagging storage resources to satisfy custom requirements of a service class

A service class can specify tags to create custom requirements for provisioning. Add the same tags to storage resources that satisfy the custom requirements. When the service class is specified during provisioning, only the tagged resources are candidates for provisioning.

You must have Administrator privileges to tag storage resources.

When you provision storage, you indicate your storage requirements by specifying a service class. Certain properties of the service class describe capabilities that storage resources must have to be a candidate for provisioning. In addition to the default properties, a service class can define custom requirements by specifying up to three custom tags. To provide the service class, a storage resource must have all the same tags as the service class. The ability to define custom requirements by tagging service classes and storage resources gives you more control over the selection of storage resources during provisioning. By using tags, you can filter the candidates for provisioning by using any criteria that you choose.

When you provision volumes, you specify the requirements by using a block-storage service class. Based on the requirements of the service class, IBM Tivoli Storage Productivity Center identifies a storage pool for the volume. If the service class specifies tags, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

When you provision shares, you specify the requirements by using a file-storage service class. Based on the requirements of the service class, Tivoli Storage Productivity Center identifies a file system or Network Shared Disk (NSD) for the share. If the service class specifies tags, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

Tag only the storage resources that satisfy the intended custom requirements of the service class. During provisioning, only the tagged resources can provide the service class.

1. Open the resource list page for the type of storage resource that you want to tag. To satisfy the custom tag requirements of a block-storage service class, you can tag a storage pool or its containing block storage system. To satisfy the custom tag requirements of a file-storage service class, you can tag a file system or NSD, or the containing file storage system.

Important: The tags on a containing storage system apply to an internal pool, file system, or NSD resource only if the internal resource is not tagged.

- If you are tagging a storage pool, from the navigation pane, select **Storage Resources** > **Pools**.
- If you are tagging the block storage system that contains the pools, complete the following steps:
  a. From the navigation pane, select **Storage Resources** > **Storage Systems**.
  b. Click the **Block Storage** tab.
- If you are tagging a file system or NSD, complete the following steps:
  a. From the navigation pane, select **Storage Resources** > **Storage Systems**.
  b. Click the **File Storage** tab.
  c. From the list of storage systems, select the storage system that contains the file system or NSD, and then select **View Details**. Alternatively, you can select **Actions** > **View Details**.
  d. In the Internal Resources area, click **File Systems** or **Network Shared Disks**.
- If you are tagging the file storage system that contains the file systems or NSDs, complete the following steps:
  a. From the navigation pane, select **Storage Resources** > **Storage Systems**.
  b. Click the **File Storage** tab.

2. In the table of storage systems, pools, file systems, or NSDs, select the storage resource that you want to tag, and then select **View Properties**. Alternatively, you can select **Actions** > **View Properties**.
3. In the properties notebook for the storage system, pool, file system, or NSD, click **Edit**.
4. In the custom tag fields of the properties notebook, specify tags that match the tags that are specified in the service class. To provide the service class, the resource must have all of the same tags.
5. In the properties notebook, click **Save**.

Creating a service class

Create service classes to specify storage requirements for future provisioning requests. A service class describes attributes and capabilities of block storage or file storage resources, and typically describes particular quality of service.

You must have Administrator privileges to create a service class.

The Create Service Class wizard guides you through the steps to create a service class. When you provision a volume or share, you define storage requirements by specifying a service class. A block-storage service class specifies the requirements for the block-storage resource from which a new volume is allocated. A file-storage service class defines requirements for the file storage resources from which a new share is allocated.

You can associate a service class with one or more capacity pools. If you do, future provisioning requests can be satisfied only by storage resources in the specified capacity pools.

You can selectively grant users who do not have administrator privileges permission to provision storage by using the service class.

Restrictions:

- To create a block-storage service class, you must have the IBM Tivoli Storage Productivity Center Advanced license.
- To associate the service class with a capacity pool, one or more capacity pools must be defined in the Tivoli Storage Productivity Center database. If no capacity pool is defined, the Specify Capacity Pool page of the wizard is not displayed.

Start this task at the Service Classes page. To open the Service Classes page, select **Advanced Analytics** > **Cloud Configuration** > **Work With Service Classes**. To create a service class, complete the following steps:

1. From the Service Classes page, click **Create Service Class**. The Create Service Class wizard opens.
2. Select the type of service class that you want to create:
   - To create a block-storage service class, click the **Block** icon.
   - To create a file-storage service class, click the **File** icon.
3. In the Define Properties page, complete the following steps:
   a. Specify a name for the service class.
   b. Optional: Specify a service class description. Descriptions are displayed in the table of service classes on the Service Classes page.
   c. Specify the general attributes of the service class. Most of the general properties specify requirements for provisioning. Some of the general properties specify configuration instructions for the storage resources or for the volumes or shares that are provisioned from them.

      Tip: To display help information about a service class property, move the mouse pointer over the field or control. Then, move the mouse pointer over the question mark icon displayed next to the field.
4. Optional: Specify or modify the advanced properties of the service class.
   a. From the Define Properties page, click **Advanced**.
   b. If you are creating or modifying a block-storage service class, and thin provisioning is enabled for the service class, specify configuration instructions for thin provisioning.
   c. If you are creating or modifying a block-storage service class, specify how the multipath driver on the client host is configured.
   d. Create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.
5. Optional: In the Specify Capacity Pool page, associate the service class with one or more capacity pools. If capacity pools are associated with the service class, future provisioning requests that specify the service class can be satisfied only by storage resources in the capacity pools.
6. Optional: Grant permission to individual users, who do not have administrator privileges, to provision storage by using the service class. By default, users must be assigned to the Administrator role to provision storage. However, you can selectively grant users who are limited to the Monitor or External Application role permission to provision by using this service class. You can also specify whether administrator approval is required to run provisioning tasks that are created by these users.
7. Click **Finish** The new service class is listed in the table of service classes on the Service Classes page, and can be specified in future provisioning requests.

Viewing details of a service class

You can view the provisioning requirements and storage constraints that a service class defines. You can also view the volumes or shares that were provisioned by using the service class.

From the Service Classes page, you can open a properties window to view the following information for a service class:

- The properties of a service class, such as the required RAID level or storage tier for a block-storage service class.
- The storage constraints that are imposed by the service class. The service class can allow provisioning from all available storage, or can limit provisioning to one or more capacity pools.
- The users who can provision storage by using the service class.
- The volumes or shares that were provisioned by using the service class.

To open the service class properties window from the Service Classes page, complete the following steps:

1. From the navigation pane in the web-based GUI, select **Advanced Analytics** > **Cloud Configuration** > **Work With Service Classes**. The Service Classes page is displayed.
2. From the list of service classes, select the service class, and then select **Actions** > **View / Modify**. The service class properties window is displayed.

Modifying a service class

You can modify the attributes of a service class to change the provisioning requirements, capacity pool storage constraints, user permissions, and service class name.

You must have Administrator privileges to modify a service class.

Important: When volumes or shares are provisioned by specifying a service class, they are associated with the service class. The Volumes page and the Shares page in the web-based GUI can display the name of the service class that was used to provision the volume or share. If you modify a service class, the volumes or shares are still associated with the service class, but might be on storage resources that no longer satisfy the requirements of the service class. Depending on your changes to the service class, users might incorrectly assume that the new volumes or shares have attributes that they do not possess.

Start this task at the Service Classes page. To open the Service Classes page, select **Advanced Analytics** > **Cloud Configuration** > **Work With Service Classes**. To modify a service class, complete the following steps:

1. From the list of service classes, select a service class and then select **Actions** > **View / Modify**. The service class properties window is displayed.
2. Optional: Click the **General** tab to modify the general attributes of the service class. Most of the general properties specify requirements for provisioning. Some of the general properties specify configuration instructions for the storage resources or for the volumes or shares that are provisioned from them.

   Tip: To display help information about a service class property, move the mouse pointer over the field or control. Then, move the mouse pointer over the question mark icon displayed next to the field.
3. Optional: Specify or modify the advanced properties of the service class.

a. From the **General** tab, click **Advanced**.
b. If you are creating or modifying a block-storage service class, and thin provisioning is enabled for the service class, specify configuration instructions for thin provisioning.
c. If you are creating or modifying a block-storage service class, specify how the multipath driver on the client host is configured.
d. Create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.

4. Optional: Modify storage constraints. If one or more capacity pools are defined in the IBM Tivoli Storage Productivity Center database, you can associate the service class with one or more capacity pools. If capacity pools are associated with the service class, future provisioning requests that specify the service class can be satisfied only by storage resources in the capacity pools. If a capacity pool was previously associated with the service class, you can disassociate it from the service class.
5. Optional: Modify user permissions. By default, users must be assigned to the Administrator role to provision storage. However, you can selectively grant users who are limited to the Monitor or External Application role permission to provision by using this service class. You can also specify whether administrator approval is required to run provisioning tasks that are created by these users.

   If users in the Monitor or External Application roles have permission to provision by using the service class, you can revoke that permission. You can specify that only administrators can provision by using the service class, or you can selectively remove users. To remove a user, right-click the user name in the table and then select **Remove** from the context-sensitive menu.

Deleting a service class

You can delete a service class that is no longer needed.

You must have Administrator privileges to delete a service class.

Deleting a service class does not affect the storage volumes or NAS shares that are already provisioned by using the service class. However, the storage volumes or NAS shares are no longer associated with any service class. When a volume that is not associated with a service class is optimized by the Optimization wizard, the volume can be migrated to any available storage pool. The storage pool to which the volume is migrated might not satisfy the requirements of the service class that was specified when the volume was provisioned.

You delete a service class from the Service Classes page.

1. From the navigation pane in the web-based GUI, select **Advanced Analytics** > **Cloud Configuration** > **Work With Service Classes**. The Service Classes page is displayed.
2. From the list of service classes, select the service class, and then select **Actions** > **Delete**.

Capacity pool configuration

Configure capacity pools to track the used and available space for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

Adding resources to a capacity pool

Divide storage resources into capacity pools to later restrict provisioning or optimization to a specific set of resources. You can also use capacity pools to track the used and available space for block and file storage on any set of storage resources. You can add resources to an existing capacity pool, or you can create a new capacity pool.

You must have Administrator privileges to perform this task.

You can add storage resources to a capacity pool from resource list pages such as the Storage Systems page and the Pools page.

You can add the following resource types to a capacity pool:

- Storage systems
- Storage pools
- File systems of file storage systems
- IBM SONAS Network Shared Disks (NSDs)

When you add a storage system to a capacity pool, any storage pool, file system, or IBM SONAS NSD that is an internal resource of that storage system is also indirectly assigned to the capacity pool.

Restriction: A resource can be assigned to only one capacity pool. When you attempt to add one or more resources to a capacity pool that are already assigned to different capacity pool, IBM Tivoli Storage Productivity Center displays a message that lists the resources and their capacity pool assignments. When a message is displayed, ensure that you want to change the capacity pool assignments for the listed resources.

To add one or more resources to a capacity pool, complete these steps:

1. Open the appropriate list page for the resource you want to add. You can add an entire storage system, or you can add individual pools, file systems, or IBM SONAS NSDs.
   - To add a storage system, from the navigation pane select **Storage Resources** > **Storage Systems**.
   - To add a storage pool, from the navigation pane select **Storage Resources** > **Pools**.
   - To add a file system or IBM SONAS NSD, complete the following steps:
     a. From the navigation pane, select **Storage Resources** > **Storage Systems**.
     b. Click the **File Storage** tab.
     c. Select the storage system that contains the file system or IBM SONAS NSD, and then select **View Details** from the context-sensitive menu.
     d. From the Internal Resources list, click **File Systems** or **Network Shared Disks**.
2. To add one or more resource to a capacity pool, select the resources in the list, and then select **Actions** > **Add to Capacity Pool**. If one or more capacity pools are defined, a list of capacity pools is displayed. If no capacity pools are defined, the Create Capacity Pool window is displayed so that you can add resources to a new capacity pool.

   Tip: The list of capacity pools and the Create Capacity Pool window shows capacity information. For the list of capacity pools, the Capacity column shows a bar that summarizes the used and free space in each capacity pool. In the

Create Capacity Pool window, a capacity bar summarizes the used and free space of the resources that are being added to the new capacity pool. Move the mouse pointer over the bar to display a text summary and details of the capacity usage.

3. Take one of the following actions:
   - If the Create Capacity Pool window is displayed, go to step 4.
   - If a list of capacity pools is displayed, add the selected resources to an existing capacity pool or to a new capacity pool.
     - To add the selected resource or resources to an existing capacity pool, select the appropriate row in the list of capacity pools, and click **Save**.
     - To add the selected resource to a new capacity pool, click **Create Capacity Pool**. The Create Capacity Pool window is displayed.
4. In the Create Capacity Pool window, specify a name for the capacity pool.
5. Optional: Specify a capacity pool description. Capacity pool descriptions appear in the list of capacity pools on the Capacity Pools page.
6. Optional: Specify custom tags to associate any significant information with the capacity pool. You can specify up to three custom tags. The custom tags can be displayed in the list of capacity pools on the Capacity Pools page.
7. Save the new capacity pool by clicking **OK**.

Viewing details of a capacity pool

You can view details of a capacity pool, such as its total capacity, used and available space, and member resources.

The Capacity Pools page lists all of the capacity pools that are defined in IBM Tivoli Storage Productivity Center. For any capacity pool in the list, you can open a window to view more details about the capacity pool and its storage resources.

To view details of a capacity pool, complete these steps:

1. From the navigation pane in the web-based GUI, select **Advanced Analytics** > **Cloud Configuration** > **Work With Capacity Pools**. The Capacity Pools page is displayed.
2. From the list of capacity pools, select the capacity pool, and then select **Actions** > **View / Modify**. A window is displayed, showing the following details:
   - A bar that summarizes the used and free space in the capacity pool in gibibytes. The bar represents the total capacity of the capacity pool and is divided into sections that represent the amount of space that is used for file storage, the amount of free space for file storage, the amount of space that is used for block storage, and the amount of free space for block storage. Move the mouse pointer over the bar to display a text summary and details of the capacity usage.
   - A table that lists the storage systems, storage pools, file systems, and IBM SONAS Network Shared Disks (NSDs) that are in the capacity pool. To view details of any of these resources, select the resource, and then select **Actions** > **View Properties**.

Modifying a capacity pool

You can modify the attributes of a capacity pool such as its name and description, and you can remove storage resources from a capacity pool.

You must have Administrator privileges to perform this task.

To modify a capacity pool, complete these steps:

1. From the navigation pane in the web-based GUI, select **Advanced Analytics** > **Cloud Configuration** > **Work With Capacity Pools**. The Capacity Pools page is displayed.
2. From the list of capacity pools, select the capacity pool, and then select **Actions** > **View / Modify**. A window that shows details of the capacity pool is displayed. A table lists the storage systems, storage pools, file systems, and IBM SONAS Network Shared Disks (NSDs) that are in the capacity pool.
3. Optional: Modify attributes of the capacity pool.
4. Optional: Remove one or more resources from the capacity pool. Removing a storage resource from a capacity pool does not delete the storage resource, but does remove its capacity from the capacity pool. To remove resources, select the appropriate resources in the list, and then select **Actions** > **Remove**. The resources are removed from the list, but are not removed from the capacity pool until you save your changes.
5. Click **Save**.

Deleting a capacity pool

You can delete a capacity pool that is no longer needed.

You must have Administrator privileges to perform this task.

When you delete a capacity pool, the storage resources that were added to the capacity pool are no longer members of any capacity pool. Deleting a capacity pool does not affect any volumes or NAS shares that were provisioned from the capacity pool. However, the volumes or NAS shares are no longer associated with a capacity pool.

To delete a capacity pool, complete the following steps:

1. From the navigation pane in the web-based GUI, select **Advanced Analytics** > **Cloud Configuration** > **Work With Capacity Pools**. The Capacity Pools page is displayed.
2. From the list of capacity pools, select the capacity pool, and then select **Actions** > **Delete**
3. In the confirmation window, click **Delete**.

Provisioning storage

To provision volumes or shares to servers or hypervisors, use the Provision Storage wizard.

When you use the Provision Storage wizard, you concern yourself only with which service class is needed for the new volumes or shares, and how much capacity you require. IBM Tivoli Storage Productivity Center uses the service class to identify storage requirements and determines the best location for the storage.

Technology Preview: The provisioning feature is part of a Technology Preview program. Use this preview to evaluate the functions and usefulness of the provisioning feature as it is being developed and enhanced for future releases. To report issues, send feedback, and submit enhancement requests, join the conversation on Service Management Connect.

Setting the zoning policy

You can set a zoning policy to enable automatic zoning. When automatic zoning is enabled, Tivoli Storage Productivity Center can create zones during storage provisioning to connect a server to a storage system.

When Tivoli Storage Productivity Center creates a provisioning task, it identifies the best location for the new storage that satisfies the requirements of the service class. If automatic zoning is enabled, then, during provisioning, existing zones are used if the server already has connectivity to the storage system. Otherwise, one or more zones are created between a host initiator port and a controller, node, or module port.

To set or modify the zoning policy, complete the following steps:

1. Click **Advanced Analytics** > **Provisioning** > **Set Zoning Policy**.
2. Enable or disable automatic zoning:
   - Enable automatic zoning to allow Tivoli Storage Productivity Center to create zones during storage provisioning. The naming convention for new zones is *host_storage-system_suffix-number*. When you enable automatic zoning, you can specify the following options:

     Zone name prefix
     If you specify a zone name prefix, all zones that are automatically created by Tivoli Storage Productivity Center are prefixed with a string you specify. This prefix can help you identify which zones were automatically created, and which were created manually.

     Make changes to the active zone set
     Whether changes are made to the active zone set, or to a new inactive active zone set. If this check box is selected, changes are made to the active zone set. If the check box is cleared, changes are made to the inactive zone set. The new inactive zone will contain only the new zones.
   - Disable automatic zoning if you want Tivoli Storage Productivity Center to use only existing zones.
3. Click **Save**.

Opening the Provision Storage wizard

The Provision Storage wizard guides you through the steps for provisioning volumes or shares.

Using the Provision Storage wizard, specify the capacity required and a service class. For shares, also specify at least one file access protocol. Tivoli Storage Productivity Center identifies the storage resources that can provide the requested capacity and service class. From the set of storage resources that can provide the capacity and service class, Tivoli Storage Productivity Center identifies the best location for the storage and creates a provisioning task. Depending on your level of permission, you can save, execute, or schedule the provisioning task.

Provisioning volumes to servers

To provision volumes to one or more servers, use the Provision Storage wizard. The wizard guides you through the steps of requesting capacity and specifying requirements by using a service class.

You must have Administrator privileges or have permission to provision storage by using at least one block-storage service class. In a service class, an administrator can grant users permission to provision storage by using the service class.

At least one block storage resource must be known to IBM Tivoli Storage Productivity Center and must be probed.

You can provision volumes to one or more servers. If a server is a virtual machine, the storage is provisioned to the hypervisor that is managing the server.

Restrictions:

- To provision volumes, you must have the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license.
- If the service class requires multipathing or if automatic zoning is enabled, the fabrics and switches must be managed by Tivoli Storage Productivity Center. If the fabric is not managed by Tivoli Storage Productivity Center, fabric-related configuration options are ignored during provisioning.
- You cannot provision storage volumes to servers and virtual machines at the same time.
- To provision volumes to multiple servers, the servers must all run the same operating system. If the servers are virtual machines, the hypervisors that manage the servers must all run the same operating system.
- When you request multiple volumes, they are provisioned from pools in the same storage system.

You open the Provision Storage wizard from the Servers page.

1. From the navigation pane in the web-based GUI, click **Server Resources** > **Servers**. Alternatively, you can click **Advanced Analytics** > **Provisioning** > **Provision to Servers**.
2. Select one or more servers, and then select **Actions** > **Provision Storage**.

The Provision Storage wizard opens.

Provisioning volumes to hypervisors

To provision volumes to one or more hypervisors, use the Provision Storage wizard. The wizard guides you through the steps of requesting capacity and specifying requirements by using a service class.

You must have Administrator privileges or have permission to provision storage by using at least one block-storage service class. In a service class, an administrator can grant users permission to provision storage by using the service class.

At least one block storage resource must be known to IBM Tivoli Storage Productivity Center and must be probed.

You can provision volumes to one or more hypervisors.

Restrictions:

- To provision volumes, you must have the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license.
- If the service class requires multipathing or if automatic zoning is enabled, the fabrics and switches must be managed by Tivoli Storage Productivity Center. If the fabric is not managed by Tivoli Storage Productivity Center, fabric-related configuration options are ignored during provisioning.

- To provision volumes to multiple hypervisors, the hypervisors must all run the same operating system.
- When you request multiple volumes, they are provisioned from pools in the same storage system.

You open the Provision Storage wizard from the Hypervisors page.

1. From the navigation pane in the web-based GUI, click **Server Resources** > **Hypervisors**. Alternatively, you can click **Advanced Analytics** > **Provisioning** > **Provision to Hypervisors**.
2. Select one or more hypervisors, and then click **Actions** > **Provision Storage**.

The Provision Storage wizard opens.

Provisioning shares to servers

To provision a network-attached storage (NAS) file share to one or more servers, use the Provision Storage wizard. The wizard guides you through the steps of requesting capacity, specifying requirements by using a service class, and providing information about how to export the share.

You must have Administrator privileges or have permission to provision storage by using at least one file-storage service class. In a service class, an administrator can grant users permission to provision storage by using the service class.

At least one file storage resource must be known to IBM Tivoli Storage Productivity Center and must be probed.

You can provision a NAS file share to one or more servers. The share is made available to the selected servers, but is not automatically mounted from the selected servers.

Restriction: You can provision NAS file shares only from IBM Scale Out Network Attached Storage and IBM Storwize V7000 Unified systems.

You open the Provision Storage wizard from the Servers page.

1. From the navigation pane in the web-based GUI, click **Server Resources** > **Servers**. Alternatively, you can click **Advanced Analytics** > **Provisioning** > **Provision to Servers**.
2. Select one or more servers, and then select **Actions** > **Provision Storage**.

The Provision Storage wizard opens.

Provisioning shares to hypervisors

To provision a network-attached storage (NAS) file share to one or more hypervisors, use the Provision Storage wizard. The wizard guides you through the steps of requesting capacity, specifying requirements by using a service class, and providing information about how to export the share.

You must have Administrator privileges or have permission to provision storage by using at least one file-storage service class. In a service class, an administrator can grant users permission to provision storage by using the service class.

At least one file storage resource must be known to IBM Tivoli Storage Productivity Center and must be probed.

You can provision a NAS file share to one or more hypervisors. The share is made available to the selected hypervisors, but is not automatically mounted from the selected hypervisors.

Restriction: You can provision NAS file shares only from IBM Scale Out Network Attached Storage and IBM Storwize V7000 Unified systems.

You open the Provision Storage wizard from the Hypervisors page.

1. From the navigation pane in the web-based GUI, click **Server Resources** > **Hypervisors**. Alternatively, you can click **Advanced Analytics** > **Provisioning** > **Provision to Hypervisors**
2. Select one or more hypervisors, and then select **Actions** > **Provision Storage**.

The Provision Storage wizard opens.

How available space for a block-storage service class is calculated and how the placement of volumes is determined

For each block-storage service class, IBM Tivoli Storage Productivity Center calculates the amount of available space for volumes of that service class. When you request volumes of a particular service class, Tivoli Storage Productivity Center recommends the best location for the volumes from the available space.

The Service Classes page in the web-based GUI displays the available space for each block-storage service class. The available space is the amount of space that is available across storage pools for new volumes of the service class. Typically, not all pools are used in the available-space calculation. Some pools might be eliminated because of the consolidated status of the pool or its storage system. Other pools might be eliminated because they do not satisfy the requirements of the service class. The set of storage pools that are used to calculate the available space are also the candidates for provisioning by using the service class.

When you request storage volumes by using the Provision Storage wizard, you specify the capacity that is required for one or more volumes, and you specify a service class. You can also specify a capacity pool from which the volumes must be provisioned. Based on the volume size that is requested and if a capacity pool is specified, the candidates for provisioning are further refined to include only the pools that have enough available space and are in the specified capacity pool. From the remaining candidates for provisioning, Tivoli Storage Productivity Center creates a provisioning task that identifies the best location for the new volumes.

How available space for a block-storage service class is calculated

For each block-storage service class, Tivoli Storage Productivity Center calculates the available space for volumes of that service class. Only storage pools that are candidates for provisioning by using the service class are used to calculate the available space.

Identifying candidates for provisioning is a filtering process. For example, if the service class specifies that RAID 10 is required, then any storage pool that is not configured at RAID 10 is eliminated as a candidate for provisioning. Tivoli Storage Productivity Center filters the set of available storage resources against certain properties of the service class until only the storage resources that can satisfy all of the requirements remain. These remaining storage resources are the candidates for provisioning. Only certain properties of a service class are used to identify

candidates for provisioning. More advanced properties contain configuration instructions, and are not used to identify candidates for provisioning.

The candidates for provisioning are identified by eliminating storage pools that cannot provide the service class, as follows:

- If a storage pool is a primordial pool or if it is formatted for count key data (CKD) volumes, it is not a candidate for provisioning.
- If a storage system has a consolidated status of Error or Unreachable, none of its internal storage pools are candidates for provisioning. If a storage pool has a consolidated status of Error, it is not a candidate for provisioning.

  Remember: A consolidated status is different from a propagated status, which is shown in the web-based GUI. A consolidated status is the status reported by a storage resource without considering the status of its internal resources. A propagated status is the combined status of the monitored, internal resources for the higher-level resource, and the consolidated status of the higher-level resource itself. The status for a storage system that is shown on the Storage Systems page and the status for a pool that is shown on the Pools page is the propagated status. The propagated status might not have the same value as the consolidated status, which is used to eliminate pools as candidates for provisioning.

  Although the propagated status of a resource might not have the same value as its consolidated status, the propagated status is always at least as severe as the consolidated status. The propagated status reflects the most severe consolidated status of any internal resource or the consolidated status of the resource itself. If the propagated status of a resource is Normal, then its consolidated status is also Normal.

  Tip: The Service Classes page in the web-based GUI displays both the available space and the unavailable space for a service class. The Unavailable Space column shows the amount of storage that satisfies the requirements of the service class, but is unavailable because of the consolidated status of the storage system or its internal resources.
- If a service class allows provisioning only from certain capacity pools, any storage pool that is not in a specified capacity pool is eliminated as a candidate for provisioning.

  Tip: You can see whether a service class restricts provisioning to one or more capacity pools from the Service Classes page. On the Service Classes page, show the Capacity Pools column, which is hidden by default.
- A storage pool must satisfy the requirements of every service class property in the following table to be a candidate for provisioning. Other properties of a service class are instructions for provisioning and do not affect whether the storage pool is a candidate for provisioning.

Table 66. Filtering candidates for provisioning based on service class properties

Service class property:	Candidates for provisioning:
Storage tier	If specified in the service class, storage pools that are assigned to any storage tier within the specified range are candidates for provisioning.

Table 66. Filtering candidates for provisioning based on service class properties (continued)

Service class property:	Candidates for provisioning:
RAID level	If a RAID level is specified, only storage pools that are configured at the specified RAID level are candidates for provisioning. **Exception:** The IBM XIV Storage System uses a nontraditional approach to data redundancy and does not support any of conventional RAID levels. Tivoli Storage Productivity Center always considers XIV system storage pools to be configured at RAID 10. **Restriction:** Storage pools that support multiple RAID levels are not candidates for provisioning when a service class specifies a RAID level. Storage pools that support multiple RAID levels are not candidates even when the specified RAID level is one of the RAID levels that are supported by the storage pool. Storage pools that support multiple RAID levels are only candidates for provisioning when the service class specifies that any RAID level is allowed.
Virtualization	If virtualization is required by the service class, only storage pools on IBM storage virtualizers are candidates for provisioning. The IBM storage virtualizers are SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified If virtualization is not allowed by the service class, only storage pools that are not on IBM storage virtualizers are candidates for provisioning.
Encryption	If encryption is required by the service class, only storage pools on DS8000 storage systems that are encrypted and configured to use the encryption group that is specified in the service class are candidates for provisioning. If encryption is not allowed by the service class, only pools that are not encrypted are candidates for provisioning.
Thin provisioning	If thin provisioning is required by the service class, only storage pools that are configured for thin provisioning are candidates for provisioning. Thin provisioning is supported by the XIV system, DS8000, SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified. During provisioning, a thin-provisioned volume is created. If thin provisioning is not allowed by the service class, any storage pool that is not thin provisioned is a candidate for provisioning. In addition, any DS8000, SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified storage pool is a candidate for provisioning, regardless of whether it is configured for thin provisioning. During provisioning, the volume that is created is not thin provisioned.
Overallocation limit	If thin provisioning is required by the service class, Tivoli Storage Productivity Center identifies the storage pools that can thin provision volumes as candidates for provisioning. If an overallocation limit is specified, Tivoli Storage Productivity Center refines that list to include only the storage pools that have an overallocation percentage that does not exceed the overallocation limit. The overallocation percentage for a pool is determined by using the following measurements, as shown on the Pools page: • For an XIV system storage pool, the Soft Space divided by the Capacity. • For other storage system pools, the Virtual Allocation.

Table 66. Filtering candidates for provisioning based on service class properties (continued)

Service class property:	Candidates for provisioning:
Resource Tags	Resource tags further refine the candidates for provisioning. A service class can specify up to 3 resource tags. If resource tags are specified for the service class, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

After the storage pools that can provide the service class are identified, Tivoli Storage Productivity Center calculates the amount of space that is available for provisioning new volumes of the service class. The calculated volume capacity is shown in the Available Space column on the Service Classes page.

The available-space calculation differs depending on whether the block-storage service class requires thin provisioning, as follows:

- If thin provisioning is not allowed, the available space is the sum of the following measurements, as shown on the Pools page:
  - For each XIV System pool, the Available Soft Space
  - For other storage system pools, the Available Pool Space
- If thin provisioning is required, the available space is the total of the following measurements, as shown on the Pools page, and calculations:
  - For each XIV System pool, the Available Soft Space
  - For other storage system pools, the Available Pool Space multiplied by the overallocation limit that is specified in the service class. For DS8000 storage pools, any space that is used by Track Space Efficient (TSE) volumes is subtracted from the calculated available space.

How the recommended placement of storage volumes is determined

When you request one or more new volumes by using the Provision Storage wizard, the candidates for provisioning are further refined and a provisioning task is created which shows the recommended placement of the new volumes. The same filtering that was used to determine the available space for the service class is used, and is refined further as follows:

- A provisioning request specifies the set of storage pools from which the storage can be provisioned. This set can be all the storage pools that are known to Tivoli Storage Productivity Center, or can be limited to the storage pools in a particular capacity pool. Only the storage pools in the specified set are candidates for provisioning. This set might be more restrictive than the set used to calculate the available space for the service class.
- Any storage pool that cannot provide the requested volume capacity is eliminated as a candidate for provisioning. If multiple volumes are requested, they will be provisioned from pools in the same storage system. If there is not enough available space across pools in a particular storage system, all of the pools are eliminated as candidates for provisioning.
- If an overallocation limit is specified in the service class, only the storage pools that will not exceed the limit after the new volume or volumes are provisioned are candidates for provisioning.

Remember: Storage pools that have an overallocation percentage that does not exceed the overallocation limit, are included in the available-space calculation. However, if the overallocation percentage would exceed the overallocation limit after the requested volume space is included, the pools are not candidates for provisioning.

After the storage pools that can satisfy the provisioning request are identified, Tivoli Storage Productivity Center finds the best location for the storage. The best location for the storage is based on the unallocated volume space in the pool and performance data. Preference is given to systems that have available performance data. When you request multiple volumes, they are provisioned from pools in the same storage system.

How available space for a file-storage service class is calculated and how placement of shares is determined

For each file-storage service class, IBM Tivoli Storage Productivity Center calculates the amount of available space for NAS file shares of that service class. When you request file shares of a particular service class, IBM Tivoli Storage Productivity Center recommends the best location for the shares from the available space.

For each file-storage service class, Tivoli Storage Productivity Center calculates the amount of available space for NAS file shares of that service class. When you request file shares of a particular service class, Tivoli Storage Productivity Center recommends the best location for the shares from the available space.

The Service Classes page in the web-based GUI displays the available space for each file-storage service class. The available space is the amount of space that is available across all IBM SONAS and IBM Storwize V7000 Unified file systems and unused IBM SONAS system NSDs that are known to Tivoli Storage Productivity Center, and that can satisfy the requirements of the service class. The set of file storage resources that are used to calculate the available space are also the candidates for provisioning by using the service class.

When you request a file share by using the Provision Storage wizard, you specify the capacity that is required for the share, and you specify a service class. You can also specify a capacity pool from which shares must be provisioned. Based on the requested size of the share and whether a capacity pool is specified, the candidates for provisioning are further refined to include only the file storage resources that have enough available space and are in the specified capacity pool. From the remaining candidates for provisioning, Tivoli Storage Productivity Center creates a provisioning task that identifies the best location for the new share.

How available space for a file-storage service class is calculated

For each file-storage service class, Tivoli Storage Productivity Center calculates the available space for file shares of the service class. Only the file storage resources that are candidates for provisioning by using the service class are used to calculate the available space.

Identifying candidates for provisioning is a filtering process. For example, if the service class requires dedicated storage, then the share must be created on an unused NSD within an IBM SONAS system. File systems that could be candidates for provisioning if shared storage is allowed are eliminated as candidates for provisioning when dedicated storage is required. Tivoli Storage Productivity Center filters the set of available storage resources against each attribute of the

service class until only the storage resources that can satisfy all of the requirements remain. These remaining storage resources are the candidates for provisioning.

The candidates for provisioning are identified by eliminating storage resources that cannot provide the service class, as follows:

- Only IBM SONAS and IBM Storwize V7000 Unified systems are supported for provisioning file shares. No other storage system type is a candidate for provisioning file shares.
- If a storage system has a consolidated status of Error or Unreachable, none of its internal NSDs or file systems are included in the available-space calculation.
- If an NSD or file system has a consolidated status of Error, it is not included in the available-space calculation.

  Tip: The status of file systems is not displayed n the web-based GUI. View the status on the NSD on which a file system resides to determine the file system status.
- If the storage volume or NSD on which a file system resides has a consolidate status of Error, the file system is not included in the available-space calculation.

  Remember: A consolidated status is different from a propagated status, which is shown in the web-based GUI. A consolidated status is the status reported by a storage resource without considering the status of its internal resources. A propagated status is the combined status of the monitored, internal resources for the top-level resource. The status for a storage system that is shown on the Storage Systems page, the status for volume that is shown on the Volumes page, and the status of an NSD that is shown on the storage system details page is the propagated status. The propagated status might not have the same value as the consolidated status, which is used to eliminate file storage resources as candidates for provisioning.

  Although the propagated status of a resource might not have the same value as its consolidated status, the propagated status is always at least as severe as the consolidated status. The propagated status reflects the most severe consolidated status of any internal resource or the consolidated status of the resource itself. If the propagated status of a resource is Normal, then its consolidated status is also Normal.

  Tip: The Service Classes page in the web-based GUI displays both the available space and the unavailable space for a service class. The Unavailable Space column shows the amount of storage that satisfies the requirements of the service class, but is unavailable because of the consolidated status of the storage system or its internal resources.
- If a service class allows provisioning only from certain capacity pools, any storage resource that is not in a specified capacity pool is eliminated as a candidate for provisioning.

  Tip: You can see whether a service class restricts provisioning to one or more capacity pools from the Service Classes page. On the Service Classes page, show the Capacity Pools column, which is hidden by default.
- A file system or NSD must satisfy the requirements of every service class property in the following table to be a candidate for provisioning. Other properties of a service class are instructions for provisioning and do not affect the storage placement determination.

Table 67. Filtering candidates for provisioning based on service class properties

Service class attribute:	**Candidates for provisioning:**
Shared storage / Dedicated storage	If the service class specifies dedicated storage, the share can be created only on an unused NSD within an IBM SONAS system. File systems and other file storage resources are eliminated as candidates for provisioning.
Resource tags	Resource tags further refine the candidates for provisioning. A service class can specify up to 3 resource tags. If resource tags are specified for the service class, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

After the file storage resources that can provide the service class are identified, Tivoli Storage Productivity Center calculates the amount of space that is available for providing new files shares of the service class.

The available-space calculation differs depending on whether the file-storage service class requires dedicated storage or allows shared storage, as follows:

- If the file-storage service class requires dedicated storage, the available space is the total available space on unused NSDs within each IBM SONAS system. You can verify that an NSD is unused, and view its available space, from the details page of an IBM SONAS system. To view information about the NSDs that are associated with an IBM SONAS system, click **Network Shared Disks**. If the NSD is unused, the File Systems column shows 0. The Available Space column shows the amount of space on the NSD, which is the measurement that is used in the available-space calculation for the service class.
- If the file-storage service class allows shared storage, the file system in which the NAS share is provisioned can contain other NAS shares. In this case the available space is the total of the following measurements:
  - The available space on unused NSDs within each IBM Scale Out Network Attached Storage (IBM SONAS) system.
  - The available space for each file system within each IBM SONAS and IBM Storwize V7000 Unified storage system. You can view the file system available space from the details page of a file storage system. To view information about the file systems that are in a storage system, click **File Systems**. The Available Space column shows the amount of available space on each file system, which is the measurement that is used in the available-space calculation for the service class.

How the recommended placement of file shares is determined

When you request a file share by using the Provision Storage wizard, the candidates for provisioning are further refined and a provisioning task is created which shows the recommended placement of the new share. The same filtering that was used to determine the available space for the service class is used, and is refined further as follows:

- A provisioning request specifies the set of file storage resources from which the storage can be provisioned. This set can be all the storage resources that are known to Tivoli Storage Productivity Center, or can be limited to the storage resources in a particular capacity pool. Only the storage resources in the

specified set are candidates for provisioning. This set might be more restrictive than the set used to calculate the available space for the service class.
- Any storage resource that cannot provide the requested capacity for the share is eliminated as a candidate for provisioning.

The shares are created on file systems or IBM SONAS NSDs. After the file systems and NSDs that can satisfy the provisioning request are identified, Tivoli Storage Productivity Center identifies the best location for the storage. The best location for the storage is based on the available space on the file system or NSD.

Optimizing storage tiering

To optimize the placement of volumes on storage tiers, use the Analyze Tiering wizard.

Before you open the Analyze Tiering wizard, complete the following tasks:
- Probe the storage virtualizers
- Set the tier level of storage pools
- Collect performance data

You must probe the storage virtualizers to collect information about the logical devices, such as volumes and virtual disks, that are associated with the storage virtualizer.

You must set the tier level of both the source storage pools and the target storage pools. Source storage pools contain the volumes that you want to analyze. Target storage pools are the storage pools that you select for the relocation of the volumes.

The Analyze Tiering wizard uses performance data, information about back-end storage technology, and conditions that are set in tiering policies to provide recommendations for tiering storage. If you implement the recommendations, you ensure that volumes are placed on the tiers that best match the workload requirements of the volumes.

You can analyze the tiering of volumes on the following storage virtualizers:
- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

Volumes can be moved to tiered storage pools on the same storage virtualizer, but volumes cannot be moved from one storage virtualizer pool to another storage virtualizer pool.

1. Select the resources that you want to analyze. The source storage pools that are related to the resources that you selected are analyzed to determine whether they meet the workload requirements of the volumes. If the workload requirements of the volume in its current tier are not met, the volume is a candidate for relocation.
2. Specify the target storage pools for the volumes.
3. Include or exclude volumes in mirrored volume relationships from the analysis.
4. Optional: Provide more information about storage pools on back-end storage systems. Tivoli Storage Productivity Center might require more information to estimate the workload capability of the source and target storage pools.

5. To ensure that the performance of the target pools is not degraded when volumes are added, you specify a maximum utilization percentage for the pools. The performance data that is collected on the previous day is used to estimate the average daily utilization of the physical resources, such as controllers, nodes, and disks, that are associated with a pool. The physical resources that are associated with a pool vary depending on the type of storage system.
6. Specify whether you want to use the automatic or create an advanced tiering policy.

   More information: See the topics about analyzing tiering in online help.

When you complete entering the information to run the Analyze Tiering wizard, recommendations for tiering the volumes are shown on the Tiering Analysis page. You can create a schedule to run the analysis based on the criteria that you selected in the Analyze Tiering wizard.

Tiering storage overview

To manage storage efficiently and reduce storage costs, optimize the placement of volumes in storage virtualizer pools.

The Analyze Tiering wizard uses performance data, information about back-end storage technology, and conditions that are set in tiering policies to provide recommendations for tiering storage. The purpose of analyzing tiering is to tier volumes based on the criteria that you set in tiering policies. For example, you can tier volumes based on the volume workload or on file usage, or both. Depending on the conditions that are set in the tiering policy, recommendations are generated. The recommendations show how storage costs can be reduced by moving volumes, for example, with low workloads to lower or less expensive tiers. The recommendations also show how performance can be improved and how storage can be used more efficiently by moving volumes with heavy workloads to the tiers that meet their workload requirements.

You can choose one of the following tiering policies to determine which volumes are moved to lower or higher tiers of storage:

- Automatic tiering
- Advanced tiering

The purpose of automatic tiering is to balance the workload of volumes that are selected for analysis across pools with the same tier level. The objective is to redistribute the workload of volumes across each tier.

The purpose of advanced tiering is to tier volumes based on threshold values such as minimum I/O density per second per GiB of storage, or file usage, or both. The objective is to ensure that volumes are placed on the tier that matches the workload and usage requirements of the volumes.

To maintain the performance level of destination pools, you set a maximum utilization percentage for the destination pools. Volumes are only moved to destination pools that can handle the additional workload without exceeding the maximum utilization percentage.

Automatic tiering

In automatic tiering, the analysis process begins with the source pools on the lowest tier and ends with the source pools on the highest tier. The workload of the volumes is balanced across pools tier by tier until the source pools on the highest tier are balanced.

Because automatic tiering involves balancing the workload of volumes across pools on the same tier level, it is assumed that you selected volumes in multiple source pools on multiple tier levels. For example, the volumes in the following source pools are selected for analysis:

Table 68. Source pools selection

Source pool	Tier
ds81_pool_a1	1
ds81_pool_a2	1
ds81_pool_b1	2
ds81_pool_b2	2
ds81_pool_c1	3

The volumes in tier 1 and tier 2 source pools can be analyzed. The volumes in the tier 3 source pool cannot be analyzed because only one tier 3 source pool is selected.

You select target pools to relocate the volumes in pools that cannot be balanced. The process to balance pools is as follows:

1. Starting with the lowest tier of storage, the source pools are balanced.
2. If the pools cannot be balanced without exceeding the maximum utilization percentage of the target pools, the most active volumes are moved to a higher tier of storage. This step does not apply to pools on the highest tier of storage.
3. The source pools are balanced again to ensure that the maximum utilization percentage of the target pools is not being exceeded.
4. If the maximum utilization percentage of the target pools is still being exceeded, the least active volumes are moved to a lower tier. This step does not apply to pools on the lowest tier of storage.
5. The source pools are balanced again to ensure that the maximum utilization percentage of the target pools is not being exceeded.
6. Steps 1 - 5 are repeated until the highest tier is balanced.

To ensure that volumes in source pools can be moved to higher or lower tiers, you select multiple target pools on multiple tier levels. For example, the following pools are selected as target pools for the source pools:

Table 69. Target pool selection

Target pool	Tier
ds82_pool_10	1
ds82_pool_11	1
ds82_pool_20	2
ds82_pool_21	2

Because the selection of the target pools is restricted to tier 1 and tier 2, the option

to move volumes in the source pools on tier 2 to target pools on tier 3 is not available.

Advanced tiering

In advanced tiering, you specify the conditions for tiering storage by specifying minimum threshold levels for each tier. You can select one or both of the following conditions to tier storage:

- The minimum I/O density per second per GiB of storage for each tier
- The minimum percentage of files that were last accessed within a specified period for each tier

The process for tiering volumes comprises two phases:

- In the first phase, the advanced tiering conditions are used to identify which volumes require relocation to lower or higher tiers
- In the second phase, the recommended destination pools are identified for the volumes that require relocation

The volumes are analyzed in their current pools to determine whether they meet the tiering conditions that are specified.

Volume tiering by I/O density

You want to tier volumes by the number of I/O operations per second, per GiB of storage.

For example, you selected a source pool with four volumes on tier 2 for analysis, and then selected target pools on tier 1, tier 2, tier 3, and tier 4. You set the following minimum thresholds for I/O density for the target pools on tier 1, tier 2, tier 3, and tier 4:

Table 70. Minimum thresholds for the I/O density of the storage tiers

Tier	Minimum I/O density thresholds
1	4 operations per second per GiB of storage
2	3 operations per second per GiB of storage
3	2 operations per second per GiB of storage
4	1 operation per second per GiB of storage

The I/O density of a volume is calculated by dividing the highest number of daily I/O operations of the volume by the capacity of the volume. The I/O densities of the volumes in the source pool on tier 2 are as follows:

Table 71. I/O densities of the volumes in the source pool

Volume	I/O density
volume_1	4 operations per second per GiB of storage
volume_2	2 operations per second per GiB of storage
volume_3	2 operations per second per GiB of storage
volume_4	1 operations per second per GiB of storage

The volumes that are identified as candidate volumes for relocation are volume_1, volume_2, and volume_4. Based on the minimum thresholds that you set for I/O density for pools on each tier, the recommendations are as follows:

- Move volume_1 from the pool on tier 2 to a pool on tier 1
- Move volume_2 from the pool on tier 2 to a pool on tier 3
- Move volume_4 from the pool on tier 2 to a pool on tier 4

The I/O density of volume_3 in the pool on tier 2 is 2 operations per second per GiB of storage. The minimum threshold for pools on tier 2 is 2.0 operations per second per GiB of storage. No recommendations are generated to relocate volume_3 to a lower or higher tier of storage.

Volume tiering by file access

You want to tier volumes based on how frequently a specified percentage of files were last accessed. You selected a pool on tier 1 for analysis, and then selected target pools on tier 1, tier 2, and tier 3.

For example, you set the following minimum thresholds for the tier 1, tier 2, and tier 3 target pools:

Table 72. Minimum thresholds for file access frequency for each storage tier

Tier	File usage threshold
1	Eighty percent of files were last accessed within 1 day.
2	Eighty percent of files were last accessed within 1 week.
3	Eighty percent of files were last accessed within 1 month.

The source pool on tier 1 has three volumes with the following file statistics:

Table 73. Last accessed statistics for the source pool on tier 1

Volume	Minimum file usage threshold
volume_a	Ninety percent of the files were last accessed within one day.
volume_b	Eighty percent of the files were last accessed within four days.
volume_c	Ninety percent of the files were last accessed within three weeks.

The volumes that are identified as candidate volumes for relocation are volume_b and volume_c. Based on the thresholds that you set for file usage on each tier, the recommendations are as follows:

- Move volume_b from tier 1 to tier 2
- Move volume_c from tier 1 to tier 3

The file usage statistics for volume_a in the pool on tier 1 is ninety percent of the files were last accessed within one day. The minimum threshold for pools on tier 1 is eighty percent of files last accessed within 1 day. No recommendations are generated to relocate volume_a to a lower or higher tier of storage.

Related tasks:

"Re-tiering volumes to balance pools" on page 96
In this scenario, you must re-tier some of the volumes to help resolve the performance hot spots that you identified. To re-tier the volumes, you run the Analyze Tiering wizard and select the default or automatic tiering policy.

"Tiering volumes by I/O density and file usage" on page 97
In this scenario, you use the Analyze Tiering wizard to regulate the placement of volumes in storage pools on tier 1. To ensure that only volumes that require the performance capability of tier 1 storage, you add minimum thresholds for I/O density and file usage for each storage tier.

Setting the tier level of storage pools

To analyze tiering and to balance pools, you must set the tier level of storage pools.

Before you open the Analyze Tiering wizard or the Balance Pools wizard, you must set the tier level of pools.

To analyze tiering, you must set the tier level of the source pools and the target pools. The source pools are the pools that you select for analysis. The target pools are the pools that are used to relocate volumes to lower or higher tiers of storage. To balance pools, you must set the tier level of the pools that you want to balance.

If you do not set the tier level of the pools, you are prompted to enter the tier level on the More Information Required page of the Analyze Tiering or the Balance Pools wizard.

Tip: You can also set or modify the tier level of pools on the **General** tab of the View Properties page for the storage pool.

1. On the Navigation pane, click **Storage Resources** > **Pools**.
2. On the Pools page, select one or more storage pools.
3. Click **Set** > **Tier**, and then select a tier level. To remove the tier level, click **None**.

The tier level that you selected for the storage pools is shown in the Tier column.

Analyzing tiering by servers

To analyze the tiering of volumes on storage virtualizers that are connected to servers, use the Analyze Tiering wizard.

The servers that you select must be connected to one or more of the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You must configure performance monitors to collect performance data about the storage subsystems that you want to analyze.

1. On the Navigation pane, click **Server Resources** > **Servers**.
2. Click one or more servers, and then select **Analyze Tiering** from the context-sensitive menu. Alternatively, you can click **Actions** > **Analyze Tiering**.

The Analyze Tiering wizard opens.

Analyzing tiering by hypervisors

To analyze the tiering of volumes on storage virtualizers that are connected to hypervisors, use the Analyze Tiering wizard.

The hypervisors that you select must be connected to one or more of the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You must configure performance monitors to collect performance data about the storage subsystems that you want to analyze.

1. On the Navigation pane, click **Server Resources** > **Hypervisors**.
2. Click one or more servers, and then select **Analyze Tiering** from the context-sensitive menu. Alternatively, you can click **Actions** > **Analyze Tiering**.

The Analyze Tiering wizard opens.

Analyzing tiering by storage virtualizers

To analyze the tiering of volumes on storage virtualizers, use the Analyze Tiering wizard.

You selected one or more of the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You must configure performance monitors to collect performance data about the storage subsystems that you want to analyze.

1. On the Navigation pane, click **Storage Resources** > **Storage Systems**.
2. Click one or more storage virtualizers, and then select **Analyze Tiering** from the context-sensitive menu. Alternatively, you can click **Actions** > **Analyze Tiering**.

The Analyze Tiering wizard opens.

Analyzing tiering by storage pools

To analyze the tiering of volumes on storage pools, use the Analyze Tiering wizard.

The storage pools that you select are on a back-end storage system that is connected to one of the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You must configure performance monitors to collect performance data about the storage subsystems that you want to analyze.

1. On the Navigation pane, click **Server Resources** > **Pools**.

2. Click one or more storage pools, and then select **Analyze Tiering** from the context-sensitive menu. Alternatively, you can click **Actions** > **Analyze Tiering**.

The Analyze Tiering wizard opens.

Analyzing tiering by volumes

To analyze the tiering of volumes, use the Analyze Tiering wizard.

The volumes that you select are in storage pools on back-end storage systems that are connected to one of the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You must configure performance monitors to collect performance data about the storage subsystems that you want to analyze.

1. On the Navigation pane, click **Server Resources** > **Volumes**.
2. Click one or more volumes, and then select **Analyze Tiering** from the context-sensitive menu. Alternatively, you can click **Actions** > **Analyze Tiering**.

The Analyze Tiering wizard opens.

Modifying the period for analyzing performance data

To change the default period for analyzing performance data in storage virtualizer pools, issue the **setdscfg** command. In the Analyze Tiering wizard, the period that is set for analyzing performance data is used to calculate the average workload activity of the volumes that are selected for analysis.

You must have Administrator authority to use the **setdscfg** command.

When you run the Analyze Tiering wizard, the performance data that is collected about the pools on storage virtualizers is analyzed. By default, the period for analyzing performance data is set to 14 days. To provide a more accurate analysis of the workload activity of the volumes in your storage environment, you can extend or shorten the period.

To change the analysis period, complete the following steps:

1. To start the Tivoli Storage Productivity Center CLI, issue the **tpctool** command.
2. Issue the **setdscfg** command, and then set the following values by using the **-property** and the **-context** parameters:
   a. Specify *tiering_pm_days* as the property value.
   b. Specify *tiering* as the context value, and then specify the number of days for the analysis period.

   For example, to change the analysis period from the default value of 14 days to 28 days, issue the following command:

```
tpctool setdscfg -url localhost:9550 -user ***** -pwd *****
-property tiering_pm_days -context tiering 28
```

The default value is changed to the number of days that you specified, for example 28.

Tip: To verify that the analysis period is changed, you enter the following command:

```
tpctool getdscfg -url localhost:9550 -user ***** -pwd ***** -context tiering
```

Balancing pools

To distribute the workload of volumes across pools on the same tier, use the Balance Pools wizard. To improve the performance of the pools, recommendations are generated to move volumes to other pools on the same tier.

Before you open the Balance Pools wizard, complete the following tasks:

- Probe the storage virtualizers. You must probe the storage virtualizers to collect information about the logical devices, such as volumes and virtual disks, that are associated with the storage virtualizer.
- Collect performance data. To gauge the average workload activity of the volumes and the utilization percentage of pools, you must collect sufficient performance data for the storage virtualizers.
- Assign a tier level to the pools that you want to balance. If you select pools that are not assigned a tier level, you are prompted to set the tier level on the More Information Required page of the Balance Pools wizard.

You can balance the workload of volumes in pools on the same tier on the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

To balance pools, you must select two or more pools on the same tier level. For example, you select three tier 1 pools, and four tier 3 pools. The pools on tier 1 are analyzed and balanced, and the pools on tier 3 are analyzed and balanced.

Before you select the pools, you must decide whether you want to balance pools in capacity pools, or pools that contain volumes with service classes. If you balance pools in a capacity pool, the recommended destination pool might not be in the same capacity pool as the volume. If you balance pools that contain volumes with service classes, the recommended destination pool might not meet the requirements of the service class.

If you select pools that contain volumes with mirrored volumes, you are prompted to include or exclude the volumes from the analysis.

1. On the Navigation pane, click **Server Resources** > **Pools**.
2. Click two or more storage pools with the same tier level, and then select **Balance Pools** from the context-sensitive menu. Alternatively, you can click **Actions** > **Balance Pools**.
3. To ensure that the performance of the pools is not degraded when volumes are added, you specify the maximum utilization percentage for the pools. The performance data that is collected is used to estimate the average daily utilization of the physical resources, such as controllers, nodes, and disks, that are associated with a pool. The physical resources that are associated with a pool vary depending on the type of storage system. The value for the utilization of a pool is estimated based on the performance data that was collected on the previous day or by the number of days that are specified by the user.

The Balance Pools wizard opens.

When you close the Balance Pools wizard, a task is created that shows the recommendations on the Balance Analysis page. You can implement the recommended changes, or you can create a schedule to implement the changes later.

Related tasks:

"Setting the tier level of storage pools" on page 433
To analyze tiering and to balance pools, you must set the tier level of storage pools.

Transforming and migrating volumes

You can transform fully allocated volumes to compressed or thin-provisioned volumes, or transform compressed or thin-provisioned volumes to fully allocated volumes. You can move volumes to other pools or to pools that are enabled for Easy Tier.

Before you open the Transform Storage wizard, you must probe the storage virtualizers to collect information about the logical devices, such as volumes and virtual disks that are associated with the storage virtualizer.

You can convert or move volumes in pools on the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

Restrictions: The restrictions for converting or moving volumes are as follows:

- To move volumes, the destination pool must be on the same storage virtualizer. You cannot move a volume from a pool on one storage virtualizer to a pool on another storage virtualizer.
- To convert or move volumes, image mode volumes must be converted to managed mode volumes.
- To convert fully allocated volumes to compressed volumes, you must have the IBM Real-time Compression™ license.

1. On the Navigation pane, click **Server Resources** > **Volumes**.
2. Click two or more storage volumes, and then select **Transform Storage** from the context-sensitive menu. Alternatively, you can click **Actions** > **Transform Storage**.

The Transform Storage wizard opens.

When you close the Transform Storage wizard, the Transform Plan page is opened. A task is created which you can run to implement the recommendations.

Modifying the properties of resources

In the web-based GUI, you can modify the properties of the resources that are monitored by Tivoli Storage Productivity Center.

The following information summarizes the resources and properties that you can modify, and the steps for modifying each type of resource.

Resource	Properties that can be modified	Steps
Storage systems	• Name (label) • Location • Custom tags (3)	1. In the navigation pane of the web-based GUI, go to **Storage Resources** > **Storage Systems**. 2. Right-click the storage system that you want to modify and select **View Properties**. 3. On the properties notebook, click **Edit**. 4. Modify the properties and click **Save**.
Pools	• Tier • Back-end storage system type • Back-end storage RAID level • Back-end storage disk type • Back-end storage disks **Tip:** You can edit the back-end values if the back-end storage system for the related storage virtualizer was not probed and the values for the pool are unknown. If you manually enter values for these properties, Tivoli Storage Productivity Center uses the values to help calculate the approximate read I/O capability of the pool.	1. In the navigation pane of the web-based GUI, go to **Storage Resources** > **Pools**. 2. Right-click the pool that you want to modify and select **View Properties**. 3. On the properties notebook, click **Edit**. 4. Modify the properties and click **Save**.

Resource	Properties that can be modified	Steps
Servers	**Servers with Storage Resource agents** You can modify the following properties for servers where agents are deployed: • Location • Custom tags (3) • Fabric functions • Agent trace • Trace level • Number of trace files • Trace file size • Run scripts on agent **Servers without Storage Resource agents (agentless)** You can modify the following properties for agentless servers: • Name (label) • OS type • IP address • Virtual machine • Location • Custom tags (3)	1. In the navigation pane of the web-based GUI, go to **Server Resources** > **Servers**. 2. Right-click the server that you want to modify and select **View Properties**. 3. On the properties notebook, click **Edit**. 4. Modify the properties and click **Save**.
Hypervisors	• Location • Custom tags (3)	1. In the navigation pane of the web-based GUI, go to **Server Resources** > **Hypervisors**. 2. Right-click the hypervisor that you want to modify and select **View Properties**. 3. On the properties notebook, click **Edit**. 4. Modify the properties and click **Save**.

Resource	Properties that can be modified	Steps
Switches	• Name (label) • Location • Custom tags (3)	1. In the navigation pane of the web-based GUI, go to **Network Resources** > **Switches**. 2. Right-click the switch that you want to modify and select **View Properties**. 3. On the properties notebook, click **Edit**. 4. Modify the properties and click **Save**.
Fabrics	• Name (label) • Location • Custom tags (3)	1. In the navigation pane of the web-based GUI, go to **Network Resources** > **Fabrics**. 2. Right-click the fabric that you want to modify and select **View Properties**. 3. On the properties notebook, click **Edit**. 4. Modify the properties and click **Save**.

Tip: To view the description of a property that you can modify for a resource, hover the mouse pointer over the hover icon for that property in the web-based GUI.

Opening the management GUI for a resource

Resources such as storage systems and switches can have their own management GUIs or element managers. In Tivoli Storage Productivity Center, you can open the start page in these GUIs from resource list pages, such as the Storage Systems page and Switches page.

The action for opening a management GUI is available for storage systems and switches under the following conditions:

- The resource is monitored by Tivoli Storage Productivity Center.
- The resource has a management GUI that is web-based or can be accessed with Java Web Start.

1. In the navigation pane of Tivoli Storage Productivity Center, select a type of top-level resource. For example, if you want open the GUI for a storage system, select **Storage Resources** > **Storage Systems**. If you want open the GUI for a switch, select **Network Resources** > **Switches**.
2. Right-click the resource and select the action to open the GUI. For storage systems, select **Open Storage System GUI**. For a switch, select **Open Switch GUI**. The start page for the management GUI is opened in a separate web browser window.

Opening the DS Storage Manager GUI for a DS8000

If the firmware level for a DS8000 is earlier than R6.2, you must complete extra configuration steps before you can open its management GUI. To determine the firmware version of a DS8000 and enable the ability to open its management GUI, complete the following steps:

1. Start the Tivoli Storage Productivity Center web-based GUI.
2. In the navigation pane, select **Storage Resources** > **Storage Systems**.
3. Right-click the DS8000 and select **View Properties**.
4. Check the version in the **Firmware** field. If the firmware version is earlier than R6.2, continue to the next step. If the version is R6.2 or later, you do not need to complete this procedure.
5. Download and start the PuTTY utility. PuTTY is a free implementation of Telnet and SSH for Windows and UNIX. You can download PuTTY from http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html.
6. On the Session page in the PuTTY tool, complete the following actions:
   a. In the **Host Name (or IP address field)**, enter the IP address of the Hardware Management Console (HMC) for the DS8000. You can determine this IP address by viewing the properties of the DS8000 in the Tivoli Storage Productivity Centerweb-based GUI.
   b. Select **SSH** for a **Connection type**.
   c. Click **Open**.
7. Enter the user name and password for logging in to the HMC.
8. Run the following command: `touch /opt/esshmc/data/config/ntvbrsen`
9. In the web-based GUI, right-click the DS8000 and select **Open Storage System GUI**.

Exporting information about resources, jobs, and alerts

You can save information about resources, jobs, or alerts to a PDF, CSV, or HTML file. Information that you can export from the web-based GUI includes all the values that are being shown in the columns for a list of resources, jobs, or alerts.

The information that you export to a file is organized according to the sorting, filtering, and column order that is defined for a list. Before you export information, complete the following tasks to configure a list:

- Change the order of columns in the list as you want them to appear in the generated file.
- Hide columns in the list that you do not want to include in the generated file. Information in hidden columns is not exported.
- Sort the rows in the list as you want them to appear in the generated file. Each list in the web-based GUI has a column or set of columns that determines the order of its rows.
- Filter the list to show only the resources, jobs, or alerts that you want to export. Use filtering to limit the rows that are shown in a list based on value in a specific column or set of default columns.

You can export information that is shown on the following pages:

Resource list and resource details pages
You can export information about the top-level resources, internal resources, or related resources that are shown on resource list pages and

resource details pages. For top-level resources, you can also export information about related jobs or alerts.

For example, you can export information about monitored switches on the Switches page. When you select the export action, the generated file includes the column values for each monitored switch in the list. In the **Jobs** or **Alerts** tabs on the Switches page, you can export information about all the jobs or alerts that are related to switches.

The Alerts page
On the **Home** > **Alerts** page, you can export information about all the alerts that were detected for monitored resources. Information that you can export includes all the values that are being shown in the columns for the list of alerts.

The Jobs page
On the **Home** > **Jobs** page, you can export information about all the jobs that are associated with monitored resources. Information that you can export includes all the values that are being shown in the columns for the list of jobs.

1. In the web-based GUI, go to a list of resources, jobs, or alerts for which you want to export information.
2. Click **Actions** and select **Export** > *file_format*, where *file_format* represents the format that you want to use for a file. The following formats are available:

   CSV (comma-separated values)
   A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application. The information in a CSV file has the following format:

   ```
   "Column_name","Column_name","Column_name","Column_name","Column_name","Column_name"
   "data","data","data","data","data","data"
   ```

   where `Column_name` represents the name of a column in a table and `data` represents the data that is associated with a column.

   PDF You can view a PDF file with Acrobat reader. The information that you export to a PDF file is formatted into a table.

   Tip: When you export a table of data that contains many columns, the rows in that table might span multiple pages in the resulting PDF. For example, if a table contains 20 columns, then the row for a specific resource might be shown on more than one page (10 columns on the first page and 10 columns on the second page).

   To reduce the number of pages that a table spans, before you select the export action, hide the columns that you do not want to include in the PDF. Continue hiding columns until the table no longer spans multiple pages.

   HTML
   You can view an HTML file with a web browser. The information that you export to an HTML file is formatted into a table.
3. Optional: Depending on how your web browser is configured, you can specify the name and location of the generated file.
4. Click the save option to export the information to a file.

Related concepts:

"Viewing information about resources" on page 363
You can view detailed information about the resources that are monitored by Tivoli Storage Productivity Center. Resources include storage systems, volumes, pools, servers, hypervisors, fabrics, and switches. You can also view information about internal and related resources.

"Customizing lists of resources, jobs, and alerts"
The web-based GUI uses lists to show information about resources, jobs, and alerts. The information in each list is organized into columns and rows. You can customize these lists by filtering information, sorting rows, and showing, hiding, and reordering columns.

Customizing lists of resources, jobs, and alerts

The web-based GUI uses lists to show information about resources, jobs, and alerts. The information in each list is organized into columns and rows. You can customize these lists by filtering information, sorting rows, and showing, hiding, and reordering columns.

Related reference:

"How performance information is displayed" on page 388
The performance view is displayed when you view the performance of resources, which includes resources in threshold violations, alerts, and performance monitors. Information in the performance view is organized into two main sections: a chart or table and a legend.

Filtering lists

You can filter the items that are shown in a list. Use filtering when you want to limit the rows that are shown in a list based on values in a specific column or set of columns. For example, when viewing alerts, you can filter the list so that only alerts with a value of Warning in the Severity column are shown.

1. In the web-based GUI, go to a list of resources, jobs, or alerts that you want to filter.
2. Click the **Filter** icon to determine how to filter the list. The following options are available:

   Filter by Default Columns
   Select this option to apply the filter to the set of the default columns in a list. Each list contains a set of default columns on which you can filter.

   For example, in the lists of alerts on the **Home** > **Alerts** page, select this option to apply filter text to the following columns at the same time: `Condition`, `Severity`, `Alert Category`, `Resource`, `Internal Resource`, `Total Occurrences`, `Alert Name`.

   Filter by a Specific Column
   Select this option to apply filter text to a specific column in a list. For example, on the **Home** > **Alerts** page, select the **Condition** column if you want to apply filter text to values in that column only.

   Restriction: Some columns in lists cannot be used for filtering and are not shown under the **Filter by a Specific Column** option.

3. In **Filter** text box

, type the text for filtering a list. You can type letters, whole words, partial words, and numbers. When filtering on certain columns, you can select from a list of values rather than typing in filter text.

   For example, when filtering on the **Status** column, you can select one of the following values as the filter text: `Error, Unreachable, Warning, Normal, Unknown`.

4. Press Enter to apply the filter to the list. Only rows with values that match or partially match the filter text are shown.

   For example, on the **Home** > **Alerts** page, if you select to filter on the Severity column and type `critical` as the filter text, only alerts that have a severity of `critical` and `critical - acknowledged` are shown in the list.

   Tip: The name of a column to which a filter applies is shown next to the **Filter** icon .

5. Optional: To remove the filtering for a list, click **Reset**.

Sorting lists

Each list of resources in the web-based GUI has a column or set of columns that determines the order of its rows. You can sort the columns in a list to organize the rows according to your requirements.

For example, the list of ports on a fabric is ordered by switch name, blade slot, and port number, in that order. You can change the sort order of these columns to organize a list of resources according to your requirements.

Tip: A visual indicator is shown next to the primary column that determines the sort order of a list.

1. In the web-based GUI, go to a list of resources, jobs, or alerts that you want to sort.
2. Complete the following tasks to sort the rows in the list:
   - To change whether the rows in a list are sorted by ascending or descending order, click the name of the primary column in the heading row. Click the name of the column again to reverse the sort order.
   - To change the column that determines the sort order of resources, click the name of that column. For example, click the Name column to sort the list of resources alphabetically, or click the **Status** column to sort the list according to the status of resources.

     Tip: When you click a column other than the primary sort column, that column becomes the new primary sort column. However, the previous primary column maintains its sort order (ascending or descending).

     For example, if a list of resources is ordered by the Name column in ascending order and you click the Status column, the rows are sorted according to the statuses of resources. At the same time, the resources for each type of status are still sorted in ascending order by the Name column. There is a maximum of three columns that can be sorted upon at the same time.
3. Optional: To reset the order of the rows for a column to the default setting, right-click the heading row in the list and select **Reset Table Preferences**.

Showing, hiding, and reordering columns in lists

Each list in the web-based GUI has a set of columns that contain information about resources, jobs, or alerts. You can change the order of columns, show columns that are hidden, and hide columns that are shown in these lists.

1. In the web-based GUI, go to a list of resources, jobs, or alerts that you want to customize.
2. Complete the following tasks to customize the columns in the list:
   - To show or hide the columns in a list, right-click any of the column headings and select or clear the check box next to a column name.
   - To change the order of columns, click a column heading and drag it to a new position in the list.

   Tip: The changes that you make to column settings are saved from session to session.
3. Optional: To reset the order and display of columns to the default setting, right-click any of the column headings and select **Reset Table Preferences**.

Managing tasks

Tasks are used to provision storage and optimize resources in your storage environment. Use the Tasks page in the web-based GUI to manage all the tasks that are used by Tivoli Storage Productivity Center to provision storage and optimize resources. Use resource list and resource details pages to manage the tasks for specific resources and resource types.

Tasks are created when you take any of the following actions in the web-based GUI:

- Assign storage to servers or hypervisors by using the Provision Storage wizard.
- Move volumes to higher or lower tiers by using the Analyze Tiering wizard.
- Move or convert volumes by using the Transform Storage wizard.
- Distribute the workload of volumes across pools on the same tier by using the Balance Pools wizard.

You can use the web-based GUI to take the following actions:

- View the status of tasks in a central location.
- Identify all the tasks that are related to a specific resource or a type of resource.
- Identify all the tasks that are encountering problems when the tasks are running.
- Manage the tasks. The actions that are available depend on the type of task that you select. For example, you can run or rename a tiering-analysis task, or export provisioning task details to a text file.

Use the stand-alone GUI to take the following configuration actions that affect the tasks that are displayed in the web-based GUI:

- Change the number of task runs that are displayed.
- Change the retention settings for removed resources. The retention settings determine how long related tasks are shown in the web-based GUI after a resource is removed.

Tips:

- In the stand-alone GUI, the term *schedules* or *jobs* is used instead of *tasks*.
- A *task run* is a single invocation of a task.

Viewing tasks

Use the Tasks page in the web-based GUI to view all the tasks that are used by Tivoli Storage Productivity Center to provision storage and optimize resources. Use resource list and resource details pages to view the tasks for specific resources and resource types.

Configuring the display of tasks

You can use the stand-alone GUI to configure how tasks are displayed in the web-based GUI.

Setting the number of task runs that are displayed:

You can change the number of task runs that are shown in the web-based GUI by using the configuration settings in the stand-alone GUI. By default, the information for the last five runs of a task is shown in the web-based GUI.

1. Start the stand-alone GUI.
2. In the navigation tree, expand **Administrative Services** > **Configuration** and click **Log File and Cached Batch Report Retention**.
3. In the **Maximum number of runs to keep of each schedule** field, enter the number of task runs that you want to retain for display.
4. Save your changes.

After you update the value for the **Maximum number of runs to keep of each schedule** field in the stand-alone GUI, the list of task runs in the web-based GUI is automatically updated within a few minutes.

Setting the retention of removed resources:

Removed resources are resources that are no longer monitored by Tivoli Storage Productivity Center. When a resource is no longer monitored, to configure how long its related tasks are shown in the web-based GUI, use the stand-alone GUI.

For example, if a storage system is removed from Tivoli Storage Productivity Center, and the retention setting is set to two weeks, the tasks for that storage system are shown for an extra two weeks before they are removed from the Tasks page. The default setting for retaining task information about removed storage systems is two weeks.

1. Start the stand-alone GUI.
2. In the navigation tree, expand **Administrative Services** > **Configuration** and click **Removed Resource Retention**.
3. Select a resource and enter the number of days to retain information after the resource is no longer monitored by Tivoli Storage Productivity Center.
4. Save the changes.

After you change the retention settings in the stand-alone GUI, the list of tasks on the Tasks page in the web-based GUI is automatically updated within a few minutes.

Viewing all tasks

Use the Tasks page as a central location to view the tasks that Tivoli Storage Productivity Center is using to provision storage and optimize resources. You can view all the tasks for resources that are monitored by Tivoli Storage Productivity Center.

1. From the navigation pane in the web-based GUI, click **Home** > **Tasks**. The information about tasks is organized into columns. These columns include information about the status of tasks, the date and time of the last run for tasks, and the schedule for tasks.
2. Optional: View the status icons at the top of the page to view a summary of task statuses. This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.

   Tip: To view descriptions of the icons and columns on the Tasks page, click the **Help** icon in the upper-right corner of the page.

Viewing tasks for a type of resource

Use resource list pages to view only the tasks for a specific type of top-level resource. Each resource list page has a **Tasks** tab that lists only the tasks for that type of resource. For example, you can use the list page for storage systems to view only the tasks that are associated with storage systems.

1. From the navigation pane, select a type of top-level resource. For example, to view the tasks that are related to storage systems, click **Storage Resources** > **Storage Systems**.

   Exception: Provisioning tasks that provision network-attached storage (NAS) file shares to servers or hypervisors are not shown on the resource list pages for servers or hypervisors. You must use the Tasks page to view these tasks.
2. Click the **Tasks** tab to view a list of all the tasks for the monitored resources of that type.

   The information about tasks is organized into columns. These columns include information about the status of tasks, the most recent date and time when tasks ran, and the schedule for tasks.

   Tip: The status icon on the **Tasks** tab represents the most critical status of the tasks that are associated with the monitored resources. For example, if one of the tasks for storage systems failed, the following icon is shown on the tab:

   Tasks
3. Optional: View the status icons on the **Tasks** page for a summary of task statuses. This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.

   Tip: To view descriptions of the icons and columns on the **Tasks** page, click the **Help** icon .

Viewing tasks for a specific resource

Use resource detail pages to view only the tasks for a specific monitored resource. Each resource detail page has a **Tasks** link that lists only the tasks for the associated storage system, server, hypervisor, switch, or fabric.

1. From the navigation pane, select a type of top-level resource. For example, if you want to view the tasks for a storage system, click **Storage Resources** > **Storage Systems**.
2. Right-click a resource in the list and select **View Details**. The resource detail page is displayed. In the **General** section, the number next to **Tasks** represents the number of tasks that are associated with the resource. For example, **Tasks (3)** indicates that you can view three tasks.

The icon next to the **Tasks** link represents the most critical status of the tasks that are associated with the resource. For example, if a task failed, the following icon is shown:

Tasks (3)

3. Click **Tasks** in the **General** section to view the tasks for the resource. The information about tasks is organized into columns. These columns include information about the status of tasks, the most recent date and time when tasks ran, and the schedule for tasks.
4. Optional: View the status icons on the Tasks pane to view a summary of task statuses. This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.

   Tip: To view descriptions of the icons and columns on the Tasks pane, click the **Help** icon .

Viewing task details

Use task detail pages to view detailed information about the tasks that Tivoli Storage Productivity Center uses to optimize resources and provision storage. You can use this information to troubleshoot the reason why a task failed, analyze details of a task before you implement task recommendations, or schedule future task runs.

On all the task detail pages, you can view information such as the task status, the duration and completion date of task runs, and the actions that you can take to manage the task. For tasks such as tiering analysis, balance analysis, and transform plan, you can view information about the recommendations that are generated by the most recent or earlier task runs. For provisioning tasks, you can view the status of the actions that the provisioning task completes for each of the resources that are configured in the task.

1. From the navigation pane, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tasks that are related to storage systems, click **Storage Resources** > **Storage Systems** and click the **Tasks** tab.
2. Locate the task that you want to view in more detail.
3. Right-click the task row and select **View Details**. The task details page opens in a separate window. For example, the Tiering Analysis page opens if you select a tiering-analysis task row.
4. Optional: To view task logs, click **Open Logs**. You can use the information in the log file to troubleshoot any errors that might occur when a task is run.
5. Optional: For provisioning tasks, you can export information about the provisioning task to a text file for offline viewing. To export information, click the icon at the top of the Provisioning page.

Viewing task logs

Task logs include detailed information about the status, actions, and progress of a task. Status information includes informational, warning, and error messages that are related to each action in a task that is taken during processing. You can use this information to troubleshoot any errors that might occur when a task is run.

The following examples show the type of information that you can view in the task logs:

- View the number of recommendations that are generated by tiering-analysis, balance-analysis, or transform-plan task runs.

- View the options that you select when you use wizards, such as the Transform Storage, Analyze Tiering, or Balance Pools wizards. For example, when you use the Analyze Tiering wizard, a tiering-analysis task is generated and the options that you choose in the wizard are recorded in the task log.
- View error messages that are related to task processing.

1. From the navigation pane, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tasks that are related to storage systems, click **Storage Resources** > **Storage Systems** and click the **Tasks** tab.
2. Locate the task that has a log file that you want to view.
3. Right-click the task row and select **Open Logs**. The Logs page shows log entries for the most recent task run.
4. Optional: On the Logs page, to view the log for a previous task run, select a task run from the **Select a log** list.
5. Optional: To view only the actions in a task run that have a Warning or Error status, select an option from the **Show all** list. You can choose to view only actions that have the following statuses:
   - Only error entries
   - Only warning entries
   - Error and warning entries
6. Optional: To view an explanation of the message that is associated with an action, click the link in the **ID** column.

The information on the Logs page is automatically updated every 30 seconds. New entries are added to the end of a log. You can view the following information on the Logs page:

- The overall status of a task. The icon that is shown in the **Select a log** list represents the most critical status that was generated by an action in the task run.
- The status for each action in a task.
- The date and time when an action was completed. The date, time, and time zone of the action is shown in the Date and Time column.
- The ID of the message that is associated with an action. You can click the value in the ID column to view more information about a message.
- The description of an action.

Tip: You can use the navigation and search functions of the web browser to locate information on the page.

Managing tasks for provisioning

Provisioning tasks are used to create storage volumes and assign the volumes to servers or hypervisors. Provisioning tasks are also used to create network-attached storage (NAS) file shares. Use the Tasks page in the web-based GUI to manage all the tasks that are used by Tivoli Storage Productivity Center to provision storage. Use resource list and resource details pages to manage the provisioning tasks for specific resources and resource types.

Use the web-based GUI to complete the following provisioning actions:

- Create provisioning tasks by using the Provision Storage wizard

- Run provisioning tasks immediately
- Schedule provisioning tasks to run at a specified date and time
- Export information about provisioning tasks to text files
- Delete provisioning tasks

Related concepts:

"Provisioning storage" on page 71
The Provision Storage wizard simplifies the steps for submitting provisioning requests. When you use the Provision Storage wizard, you concern yourself only with which service class is needed for the new volumes or shares, and how much capacity you require.

Creating provisioning tasks

Provisioning tasks are created when you use the Provision Storage wizard to assign storage volumes or NAS file shares to servers or hypervisors.

To use the Provision Storage wizard, you must have Administrator privileges.

Related tasks:

"Opening the Provision Storage wizard" on page 418
The Provision Storage wizard guides you through the steps for provisioning volumes or shares.

"Collecting information about shares on storage systems" on page 306
To collect information about shares in file systems, you must modify the `TPCD.config` file. When you provision shares, you can then view such information as the name of the share, the name of the file system, and the protocols that are used to share files.

Running provisioning tasks

Use the Tasks page to run provisioning tasks. You can run provisioning tasks immediately or you can schedule tasks to run at a date and time that you specify. When you run a provisioning task, the configuration changes that are contained in the task are implemented. For example, storage volumes are created and assigned to servers.

To run provisioning tasks, you must be the owner of the task or have Administrator privileges.

Provisioning tasks can run only once. You can run the task when you complete the steps of the Provision Storage wizard or you can save the provisioning task and run the task later.

The **Run Now** and **Schedule** actions are available for a provisioning task when the task is not running, and was not previously run.

1. From the navigation pane, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the provisioning tasks that are related to servers, click **Server Resources** > **Servers**, and click the **Tasks** tab.
2. Locate the provisioning task that you want to run.
3. Depending on when you want the provisioning task to run, choose one of the following options:
   - To run the provisioning task immediately, right-click the task row and select **Run Now**. The task status changes to Running and the provisioning task runs immediately.

- To schedule the provisioning task to run at a specific date and time, complete the following steps:
  a. Right-click the task row and select **Schedule**.
  b. From the Schedule page, select a future date and time from the **Provision** lists, and click **Save**.

  You can also run or schedule the provisioning task from the Provisioning detail page. On the Provisioning detail page, click **Execute** to run the task immediately or click **Schedule** to schedule the task to run at a specific date and time.

The provisioning task runs immediately or on the scheduled date and time. To check the progress of the task, you can take the following actions:

- View the task status on the Tasks page.
- View detailed information about the progress of the task on the Provisioning detail page. You can view the task duration and the status of the actions that the provisioning task completes for each of the resources that are configured in the task.
- View the task logs from the Tasks page or from the Provisioning detail page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.

Exporting provisioning task information

Use the Provisioning detail page to export details of a provisioning task to a text file. The task details that are shown on the Provisioning page are saved to the generated text file.

If you do not have authority to provision storage, you can export the provisioning task details to a file and email the file to a user with Administrator privileges. The administrator can run or schedule the provisioning task if they approve the task details.

1. From the navigation pane, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the provisioning tasks that are related to servers, click **Server Resources** > **Servers** and click the **Tasks** tab.
2. Locate the provisioning task that you want to export.
3. Right-click the task row and select **View Details**.
4. From the Provisioning page, to start the export, click the icon in the upper-right corner of the page.
5. Save the generated text file. Depending on how your web browser is configured, you can specify the file name and the location where the file is saved.

The provisioning task information is exported to a text file. The generated file includes all the provisioning task details that are listed on the Provisioning page. Columns that are hidden on the Provisioning page are included in the generated file. The information in the file is organized according to a default order and is not affected by the column order that is defined on the Provisioning page.

Managing tasks for tiering storage, balancing pools, and transforming storage

To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the Tasks page in the web-based GUI to manage all the tasks that are used by Tivoli Storage Productivity Center to optimize storage. Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.

To schedule, run, or delete tiering storage, balancing pools, and transforming-storage tasks, you must have Administrator privileges.

To complete any of the optimization tasks, such as balancing pools, you must complete the following steps:

1. Create the analysis task. Use a Tivoli Storage Productivity Center wizard to create the optimization analysis task. For example, use the Balance Pools wizard to create a balance-analysis task.
2. Run the analysis task to generate the optimization recommendations. For example, run the balance-analysis task to generate the recommendations to move volumes.
3. Implement the recommendations to optimize storage. For example, implement the recommendations to balance pools by running the analysis-execution task.

Use the Tasks, resource list, and resource detail pages in the web-based GUI to take the following actions for optimization tasks:

- Run the optimization analysis immediately to generate optimization recommendations.
- Schedule the optimization analysis to run at a specified date and time.
- Implement optimization recommendations immediately.
- Schedule the implementation of optimization recommendations for a specified date and time.
- Pause, cancel, and resume the implementation of optimization recommendations.
- Delete optimization tasks.
- View task logs.

Creating analysis tasks

To create tasks for optimizing resources in your storage environment, use Tivoli Storage Productivity Center wizards.

To use the wizards, you must have Administrator privileges.

Storage optimization tasks are created when you take any of the following actions in the web-based GUI:

Tier storage
Move volumes to higher or lower tiers by using the Analyze Tiering wizard. A tiering-analysis task is created.

Balance pools
Distribute the workload of volumes across pools on the same tier by using the Balance Pools wizard. A balance-analysis task is created.

Transform storage
Move or convert volumes by using the Transform Storage wizard. A transform-plan task is created.

Related tasks:

"Optimizing storage tiering" on page 428
To optimize the placement of volumes on storage tiers, use the Analyze Tiering wizard.

"Balancing pools" on page 436
To distribute the workload of volumes across pools on the same tier, use the Balance Pools wizard. To improve the performance of the pools, recommendations are generated to move volumes to other pools on the same tier.

"Transforming and migrating volumes" on page 437
You can transform fully allocated volumes to compressed or thin-provisioned volumes, or transform compressed or thin-provisioned volumes to fully allocated volumes. You can move volumes to other pools or to pools that are enabled for Easy Tier.

Running analysis tasks

To run tiering analysis, balance analysis, and transform-plan tasks, use the Tasks page. When you run the analysis for any of these optimization tasks, recommendations are generated to optimize the storage resources that are configured in the task. For example, when you run a tiering-analysis task, recommendations are generated to tier volumes based on the criteria that you set in tiering policies.

To run tiering analysis, balance analysis, and transform-plan tasks, you must have Administrator privileges.

You can run tasks immediately or schedule tasks to run at a future date and time.

Tip: The **Run Now** action is available when the task is enabled and the task is not running.

1. From the navigation pane, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tiering-analysis tasks that are related to servers, click **Server Resources** > **Servers** and click the **Tasks** tab.
2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to run.
3. Right-click the task row and select **Run Now**. The task status changes to Running.
4. Optional: To schedule future task runs for tiering analysis or balance-analysis tasks, right-click the task row and select **Schedule Analysis**. On the Schedule Analysis page, select the time that the task is run and how often the task is run.

   Restriction: You cannot schedule future task runs for transform-plan tasks.
5. Optional: To enable or disable a tiering analysis or balance-analysis task, select **Enabled** or **Disabled** on the Schedule Analysis page. When the task is enabled, the task runs according to the defined schedule. If the task is disabled, you can modify the task schedule but the task cannot run at the scheduled time.

When the task run completes, you can view the recommendations that are generated on the task detail page. For example, you can view the recommendations that are generated for a tiering-analysis task on the Tiering Analysis page. To check the progress of the task run, you can take the following actions:

- View the task status on the Tasks page.

- View information about the progress of the task on the task detail page. You can view the duration of the task run, and the date and time that the task run completed.
- View the task logs from the Tasks page or from the task detail page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.

Implementing recommendations to optimize storage

Use the Tasks page and task detail pages to implement optimization recommendations to tier storage, balance pools, and transform storage. You can implement recommendations immediately or you can schedule the implementation of the recommendations for a future date and time.

Implementing recommendations immediately:

To implement optimization recommendations to tier storage, balance pools, and transform storage, use the task detail pages. For example, use the Balance Analysis page to implement the recommendations to balance pools.

To implement storage optimization recommendations, you must have Administrator privileges.

Tivoli Storage Productivity Center uses analysis-execution tasks to implement optimization recommendations. When you implement recommendations for a tiering analysis, balance analysis, or transform-plan task, the following process takes place:

- An analysis-execution task is automatically created so that the recommendations can be implemented.
- The analysis-execution task runs immediately and the optimization recommendations are implemented. For example, volumes are moved to higher or lower tiers when storage-tiering recommendations are implemented.

You can view analysis-execution tasks on the Tasks page. The Related Task column shows the related tiering analysis, balance analysis, or transform-plan task.

The **Execute** action is available under the following conditions:

- Recommendations were generated by the latest run of the tiering analysis, balance analysis, or transform-plan task.
- An analysis-execution task was not already run for the latest set of recommendations.
- An analysis-execution task is not running or paused.

1. From the navigation pane, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the analysis-execution tasks that are related to servers, click **Server Resources** > **Servers** and click the **Tasks** tab.
2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to implement.
3. Right-click the task row and select **View Details**. For example, the Balance Analysis page opens if you select a balance-analysis task.
4. Click **Execute**. An analysis-execution task is created and immediately implemented.

The analysis-execution task runs immediately. To check the progress of the task, you can take the following actions:

- View the task status on the Tasks page.
- View the status of individual task recommendations on the Analysis Execution page.
- View the task logs from the Tasks page or from the Analysis Execution page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.

Scheduling the implementation of recommendations:

To schedule the implementation of optimization recommendations to tier storage, balance pools, and transform storage, use the Tasks page.

To schedule the implementation of storage optimization recommendations, you must have Administrator privileges.

The **Schedule Execution** action is available under the following conditions:

- Recommendations were generated by the latest run of the tiering analysis, balance analysis, or transform-plan task.
- An analysis-execution task is not running or paused.
- An analysis-execution task was not already run for the latest set of analysis recommendations. For each run of the tiering analysis, balance analysis, or transform-plan task, you can complete only one run of the analysis-execution task.

When you define the schedule, an analysis-execution task is created for the implementation of the optimization recommendations.

Tip: The analysis-execution task implements the recommendations that were current when the analysis-execution task was defined. For example, you define a tiering-analysis task that runs daily. You review the recommendations every day and, on a Tuesday, you schedule an analysis-execution task to run on the following Friday. When the analysis-execution task runs on Friday, the recommendations that were current on Tuesday are implemented, even if more recent recommendations are available.

1. From the navigation pane, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tiering analysis tasks that are related to servers, click **Server Resources** > **Servers** and click the **Tasks** tab.
2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to schedule.
3. Right-click the task row and select **Schedule Execution**.
4. On the Schedule Execution page, select the date and time that the recommendations are implemented and click **Save**.

An analysis-execution task is created that you can monitor on the Tasks page. The name of the analysis-execution task is based on the name of the related tiering analysis, balance analysis, or transform-plan task. The Related Task column shows the related tiering analysis, balance analysis, or transform-plan task.

When the analysis-execution task runs on the scheduled date, the optimization recommendations are implemented. For example, volumes are moved to higher or lower tiers when storage-tiering recommendations are implemented.

Pausing, resuming, and canceling the implementation of recommendations

Tivoli Storage Productivity Center uses analysis-execution tasks to implement optimization recommendations to tier storage, balance pools, and transform storage. To pause, resume, and cancel analysis-execution tasks, use the Tasks page.

To pause, resume, or cancel analysis-execution tasks, you must have Administrator privileges.

1. From the navigation pane, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the analysis-execution tasks that are related to servers, click **Server Resources** > **Servers** and click the **Tasks** tab.
2. Locate the analysis-execution task that you want to pause, resume, or cancel.
3. Right-click the task row and select **Pause**, **Resume**, or **Cancel**.

   The **Pause** action is only available when the analysis-execution task is running.

   The **Resume** action is only available when the analysis-execution task is paused.

   The **Cancel** action is available when the task is running or paused.

When you pause, resume, or cancel analysis-execution tasks, the following processing occurs:

Pause The status of the analysis-execution task changes to Paused and the recommendations that are being implemented are suspended.

Resume
The status of the analysis-execution task changes to Running and the recommendations are resumed.

Cancel
The status of the analysis-execution task changes to Canceled and the recommendations that are being implemented are stopped. Any recommendations that are already implemented for this analysis-execution task are not rolled back. For example, a recommendation is generated to move a volume from `pool_a` to `pool_b`. If this recommendation is already implemented before the task is canceled, the volume remains in `pool_b`. Use the **Analysis Execution** page to determine which recommendations were implemented for the task before the task was canceled.

You cannot restart an analysis-execution task that is canceled. If you want to implement the same set of recommendations again, you must complete the following steps:

1. Delete the analysis-execution task that you canceled.
2. Use the **Schedule Execution** action for the tiering analysis, balance analysis, or transform-plan task to create an analysis-execution task.

Renaming tasks

You can change the name of a task by using the task detail page. For example, use the Provisioning page to change the name of a provisioning task. A task name is automatically generated when the task is created. If the original name is ambiguous or is not clear enough, you can change the task name to a unique task name that you provide.

You must have the following privileges to rename the different task types:

Provisioning
To rename provisioning tasks, you must be the owner of the task or have Administrator privileges.

Other task types
To rename all other tasks, such as tiering analysis, balance analysis, transform plan, and analysis-execution tasks, you must have Administrator privileges.

1. From the navigation pane, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view the tasks that are related to servers, click **Server Resources** > **Servers** and click the **Tasks** tab.
2. Locate the task that you want to rename.
3. Right-click the task row and select **View Details**. The task detail page opens in a separate window. For example, the Provisioning page opens if you select a provisioning task.
4. Type the new name in the task name field. The naming conventions for task names are:
   - The maximum length is 64 characters.
   - Uppercase and lowercase letters are allowed.
   - Spaces and numbers are allowed.
   - The following special characters are allowed:
     ! # % & * + - / = ? ^ _ { } ( )
5. Click **Rename**.

The task is renamed. You can view the renamed task on the Tasks page.

Deleting tasks

Use the Tasks page to delete tasks, such as provisioning, tiering analysis, balance analysis, transform plan, and analysis-execution tasks.

You must have the following privileges to delete the different task types:

Provisioning
To delete provisioning tasks, you must be the owner of the task or have Administrator privileges.

Other task types
To delete all other tasks, such as tiering analysis, balance analysis, transform plan, and analysis-execution tasks, you must have Administrator privileges.

The following conditions apply for the different task types:

Analysis-execution
You can delete analysis-execution tasks that were canceled, already run, or scheduled to run in the future. You cannot delete tasks that are running or paused.

Other task types
You can delete all other tasks if the tasks were already run or if the tasks are scheduled to run in the future. For example, you can delete provisioning, tiering analysis, balance analysis, and transform-plan tasks that were already run. You cannot delete tasks that are running.

1. From the navigation pane, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource. For example, to view tasks that are related to servers, click **Server Resources** > **Servers** and click the **Tasks** tab.
2. Locate the task that you want to delete.
3. Right-click the task row and select **Delete**. Alternatively, you can delete the task by clicking **Delete** on the task detail page. The Delete page opens in a separate window.
4. If you are deleting a provisioning task that was not yet run, enter the reason that you are deleting the task.
5. Click **Delete**. To exit the window without deleting the task, click **Cancel**.

Managing resources in a vSphere environment

You can use the vSphere Web Client extension for Tivoli Storage Productivity Center to manage and monitor your virtual environment and storage devices.

You can provision block and file storage to hypervisors from storage devices that are monitored by Tivoli Storage Productivity Center. You can monitor the progress of your storage provisioning activities in task information in the vSphere Web Client.

Provisioning block storage by using the vSphere Web Client extension for Tivoli Storage Productivity Center

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision block storage on Tivoli Storage Productivity Center. You can also select to have a data store created on the storage volume by the Block Storage Provisioning wizard.

Ensure that you complete the following tasks in Tivoli Storage Productivity Center:

- If you want to provision storage by using service classes to describe performance and storage requirements, ensure that the service classes are created in Tivoli Storage Productivity Center. Tivoli Storage Productivity Center provides a number of predefined service classes, which you can use, modify, or delete.
- If you want to provision storage by using capacity pools to specify your storage requirements, ensure that the capacity pools are created in Tivoli Storage Productivity Center.
- Before you provision storage to a hypervisor, ensure that the hypervisor is added into Tivoli Storage Productivity Center. If the hypervisor belongs to a cluster, all the hypervisors on that cluster must be managed by Tivoli Storage Productivity Center.
- Ensure that the hypervisor has fabric connectivity. If automatic zoning is enabled, the fabrics and switches must be managed by Tivoli Storage Productivity Center.

For more information about configuring for storage cloud provisioning, see the Tivoli Storage Productivity Center information center. Search for *Configuring for storage cloud provisioning*.

Start the vSphere Web Client.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters**.
2. Select the hypervisor for which you want to provision storage.
3. Click the **Actions** menu item, and select **All TPC Actions**.

4. To open the **Provision Block Storage** wizard, select **Provision Block Storage...**.
5. In the wizard, enter a size for the total usable storage space on the volume, and select the unit of measurement from the list.
6. Select the service class for the block storage that you want to associate with the volume. You can select only service classes with available capacity that you are authorized to use.
7. Select the capacity pool from which you want the volume to be provisioned. You can select only capacity pools with available capacity.
8. Optional: If you want to create a data store on the storage volume when it is provisioned, select the **Create datastore** check box, and enter a name for the data store. The minimum volume size for creating a data store is 2 GiB.
9. Click **OK** to provision the volume.
10. To monitor the progress of the provisioning task, click the **Monitor** tab, and click the **Tasks** tab.
11. You can also monitor the progress of the creation of the data store by clicking the **Monitor** tab and then clicking the **Tasks** tab. The data store creation task is called **Create VMFS datastore**.

Related concepts:

"Service classes" on page 66
A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.

"Capacity pools" on page 70
Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. Configure capacity pools to track the used and available space for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

Related tasks:

"Provisioning file storage by using the vSphere Web Client extension for Tivoli Storage Productivity Center" on page 461
Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision file storage on Tivoli Storage Productivity Center.

Block storage provisioning in the vSphere Web Client extension for Tivoli Storage Productivity Center

Use the Provision Block Storage wizard to enter properties for the volumes that you want to provision. Tivoli Storage Productivity Center uses this information to create the volumes on the selected capacity pool.

You can also select to have a data store created on the storage volume by the wizard.

Before you use the wizard, ensure that you complete the following tasks in Tivoli Storage Productivity Center:

- If you want to provision storage by using service classes to describe performance and storage requirements, ensure that the service classes are created in Tivoli Storage Productivity Center. Tivoli Storage Productivity Center provides a number of predefined service classes, which you can use, modify, or delete.

- If you want to provision storage by using capacity pools to specify your storage requirements, ensure that the capacity pools are created in Tivoli Storage Productivity Center.
- Before you provision storage to a hypervisor, ensure that the hypervisor is added into Tivoli Storage Productivity Center. If the hypervisor belongs to a cluster, all the hypervisors on that cluster must be managed by Tivoli Storage Productivity Center.
- Ensure that the hypervisor has fabric connectivity. If automatic zoning is enabled, the fabrics and switches must be managed by Tivoli Storage Productivity Center.

The entity that you select from the Hosts and Clusters inventory view determines the information that Tivoli Storage Productivity Center requires when you use the wizard to provision block storage for that entity.

You must specify the following information to provision volumes:

Size The total usable storage space on a volume. The default unit for volume size is GiB. If you select to create a datastore, the minimum volume size is 2 GiB.

Service class
The name of the service class that you want to associate with the volume. For block storage, a service class typically represents a particular quality of service, for example, Gold, Silver, or Bronze.

Restriction: You can select only service classes with available capacity that you are authorized to use.

Capacity pool
The name of the capacity pool from which you want the volume to be provisioned.

Restriction: You can select only capacity pools with available capacity that you are authorized to use.

Create datastore
Select this check box if you want to create a datastore on the storage volume when it is provisioned.

Restriction: The minimum volume size for creating a datastore is 2 GiB.

Datastore name
The name of the datastore. The maximum character length is 42. Brackets, [], are not allowed.

Related concepts:

"Service classes" on page 66
A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.

"Capacity pools" on page 70
Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. Configure capacity pools to track the used and available space for block and file storage on any set of storage resources. Provisioning requests can

also be restricted to resources in a capacity pool.

Provisioning file storage by using the vSphere Web Client extension for Tivoli Storage Productivity Center

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision file storage on Tivoli Storage Productivity Center.

Ensure that you complete the following tasks in Tivoli Storage Productivity Center:

- If you want to provision storage by using service classes to describe performance and storage requirements, ensure that the service classes are created in Tivoli Storage Productivity Center. Tivoli Storage Productivity Center provides a number of predefined service classes, which you can use, modify, or delete.
- If you want to provision storage by using capacity pools to specify your storage requirements, ensure that the capacity pools are created in Tivoli Storage Productivity Center.
- Before you provision storage to a hypervisor, ensure that the hypervisor is added into Tivoli Storage Productivity Center. If the hypervisor belongs to a cluster, all the hypervisors on that cluster must be managed by Tivoli Storage Productivity Center.
- Determine whether you must set access path host name replacements. For more information, see "Special consideration for provisioning by using the VMware vSphere Web Client extension" on page 69.

For more information, see "Cloud configuration" on page 65.

Start the vSphere Web Client.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters**.
2. Select the hypervisor for which you want to provision storage.
3. Click the **Actions** menu item, and select **All TPC Actions**.
4. To open the **Provision File Storage** wizard, select **Provision File Storage...**.
5. In the **Share Requirements** page of the wizard, enter a name for the share.
6. Enter a size for the total usable storage space on the share, and select the unit of measurement from the list.
7. Select the service class for the file storage that you want to associate with the share. You can select only service classes with available capacity that you are authorized to use.
8. Select the capacity pool from which you want the share to be provisioned. You can select only capacity pools with available capacity.
9. Enter the name of the data store that you want to associate with the file system.
10. Click **Next** to continue.
11. In the **Share Permissions** page of the wizard, specify the user and group with permission to mount the share on the hypervisor.
12. Select the **Read-only** check box if you want to limit the user or group to viewing files on the share.
13. Select the **Secure** check box to create the share with the secure flag set.
14. Click **OK** to provision the share.
15. To monitor the progress of the provisioning task, click the **Monitor** tab, and click the **Tasks** tab.

16. You can also monitor the progress of the creation of the data store by clicking the **Monitor** tab and then clicking the **Tasks** tab. The data store creation task is called **Create NAS datastore**.

Related concepts:

"Service classes" on page 66
A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.

"Capacity pools" on page 70
Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. Configure capacity pools to track the used and available space for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

Related tasks:

"Setting access path host replacements" on page 409
If you are going to use the VMware vSphere Web Client extension to provision file shares, verify whether you must provide Tivoli Storage Productivity Center with an access path host replacement.

"Provisioning block storage by using the vSphere Web Client extension for Tivoli Storage Productivity Center" on page 458
Use the vSphere Web Client extension for Tivoli Storage Productivity Center to provision block storage on Tivoli Storage Productivity Center. You can also select to have a data store created on the storage volume by the Block Storage Provisioning wizard.

File storage provisioning in the vSphere Web Client extension for Tivoli Storage Productivity Center

Use the Provision File Storage wizard to enter properties for the network-attached storage (NAS) shares that you want to provision. Tivoli Storage Productivity Center uses this information to create the shares on a file storage system, for example, IBM Scale Out Network Attached Storage.

Before you use the wizard, ensure that you complete the following tasks in Tivoli Storage Productivity Center:

- If you want to provision storage by using service classes to describe performance and storage requirements, ensure that the service classes are created in Tivoli Storage Productivity Center. Tivoli Storage Productivity Center provides a number of predefined service classes, which you can use, modify, or delete.
- If you want to provision storage by using capacity pools to specify your storage requirements, ensure that the capacity pools are created in Tivoli Storage Productivity Center.
- Before you provision storage to a hypervisor, ensure that the hypervisor is added into Tivoli Storage Productivity Center. If the hypervisor belongs to a cluster, all the hypervisors on that cluster must be managed by Tivoli Storage Productivity Center.
- Determine whether you must set access path host name replacements. For more information, see "Special consideration for provisioning by using the VMware vSphere Web Client extension" on page 464.

The entity that is selected from the Hosts and Clusters inventory view determines the information that Tivoli Storage Productivity Center requires to provision file storage for that entity.

You must specify share requirements and share permissions to provision file storage:

Share requirements

Share name
The name of the share that you want to associate with the file system.

Size The total usable storage space on a share. The default unit for share size is GiB.

Datastore name
The name of the datastore that you want to associate with the file system. The maximum character length is 42. Brackets, [], are not allowed.

Service class
The name of the service class that you want to associate with the share. For file storage, a service class typically represents a particular quality of service. For example, the `NormalIsolation` service class represents file systems where shared storage is allowed, while the `EnhancedIsolation` service class represents file systems where dedicated storage is required.

Restriction: You can select only service classes with available capacity that you are authorized to use.

Capacity pool
The name of the capacity pool from which you want the share to be provisioned.

Restriction: You can select only capacity pools with available capacity that you are authorized to use.

Share permissions

User The user name that has permission to mount the Network File System (NFS) share on a hypervisor. This user must be created in the file storage system with LDAP configured.

Group The group that has permission to mount the NFS share. This group must be created in the file storage system with LDAP configured.

NFS Client Permissions
Read-only
The user or group does not have permission to modify files on the NFS share.

Secure
The NFS share is created with the secure flag set. Only NFS connections that originate from secure IP ports are permitted.

Special consideration for provisioning by using the VMware vSphere Web Client extension

When a file share is provisioned by the VMware vSphere Web Client extension, an access path to the share is provided by Tivoli Storage Productivity Center. By default, the file share access path that is provided to the VMware vSphere Web Client extension includes the cluster name of the file storage system in which the share is created. A cluster name must be a valid host name that is registered with a Domain Name System (DNS) server. If the cluster name is not valid, the share is not accessible under the access path that is returned by Tivoli Storage Productivity Center.

To address this situation, a file storage system can set one of its custom tags to specify an access path host name replacement. A file storage service class can specify which of the file storage system custom tags contains the access path host name replacement.

If file shares are provisioned by using the VMware vSphere Web Client extension, determine whether you must set access path host name replacements.

Managing resources in the stand-alone GUI

You can use the stand-alone GUI to manage and monitor the resources and storage infrastructure in an enterprise environment.

In V5.2, many of the functions for managing and monitoring storage in the stand-alone GUI were moved to the web-based GUI. For more information about the storage-management tasks that you can complete in each GUI, see "Available functions in the interfaces" on page 34.

Related concepts:

Chapter 5, "Managing resources," on page 257
Use Tivoli Storage Productivity Center to monitor and manage the resources and storage infrastructure in an enterprise environment.

GUI overview

The stand-alone GUI is a Java based interface for managing storage resources that is organized into two main sections: a navigation tree and a content pane.

Navigation tree

The navigation tree of the stand-alone GUI is displayed on the left side of the main window. You can expand, select, and collapse the nodes in the tree to access storage-management functions in the interface.

Figure 18. Navigation tree 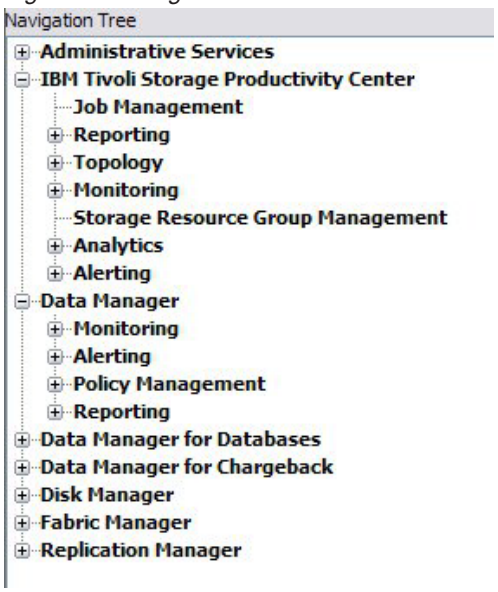	When you start Tivoli Storage Productivity Center, the navigation tree is expanded to show all the high-level functions in the product. To view the nodes within each high-level function, click the expand icon next to the node or double-click the name of the node. **Expanding and collapsing the navigation tree:** If the navigation tree gets too large, or if you want to return it to its original state, right-click the major nodes of the tree and select **Collapse Branch** or **Expand Branch**. For example, right-click the **IBM Tivoli Storage Productivity Center** node and select **Collapse Branch** to collapse all the nodes in that branch of tree. To expand all the nodes in the **IBM Tivoli Storage Productivity Center** branch, right click the node and select **Expand Branch**. **Printing the navigation tree:** Right-click a node and select **Print Branch** to print all the subnodes within that node.

Content pane

The content pane is displayed on the right side of the main window. When you select a node in the navigation tree, a corresponding panel is displayed in the content pane. Use the panels in the content pane to complete tasks such as configuring Tivoli Storage Productivity Center, creating data collection schedules, defining alerts, and viewing reports.

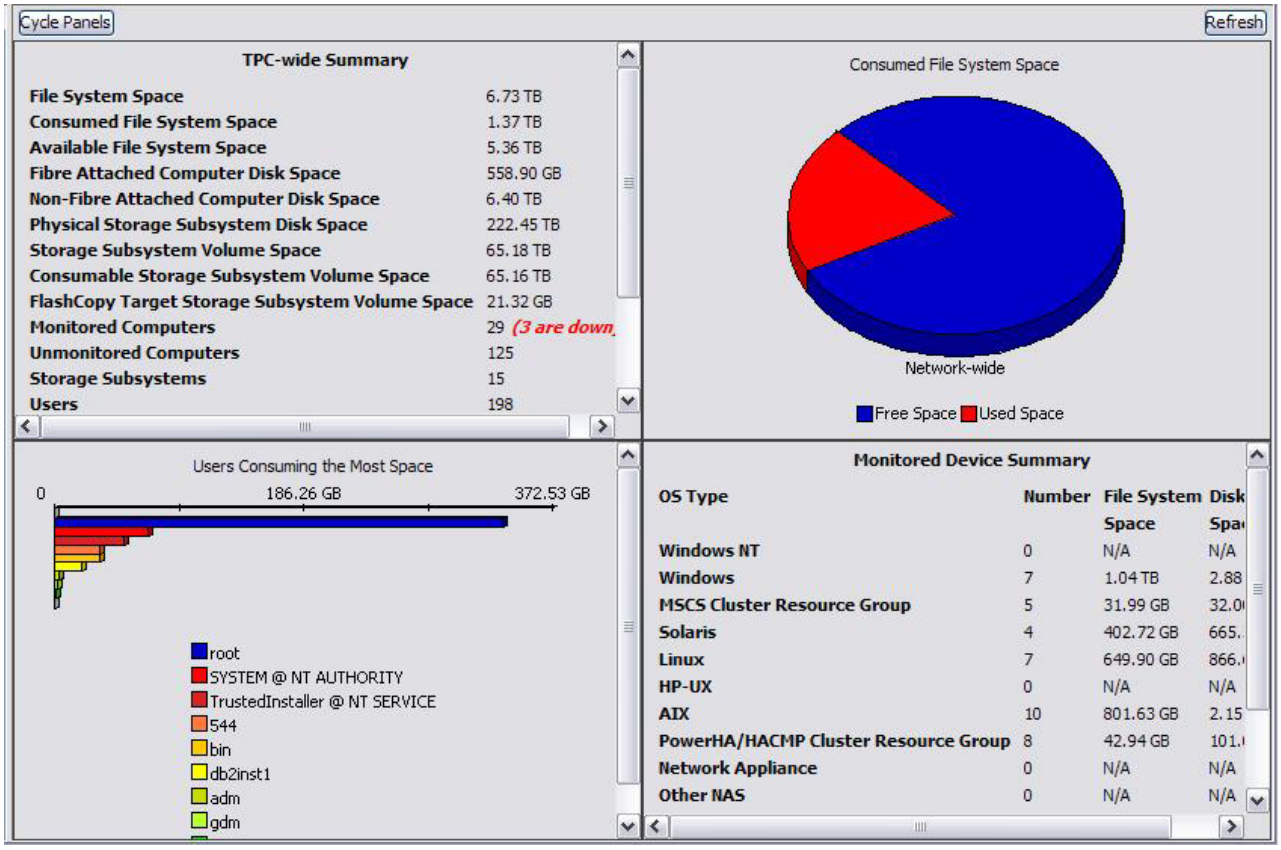

Figure 19. Content pane

Navigation tree

The navigation tree includes expandable and collapsible nodes for all the functions that are available in the current product license. These functions are organized into top-level nodes for Administrative Services, IBM Tivoli Storage Productivity Center, Data Manager, Disk Manager, Fabric Manager, and Replication Manager.

Administrative Services
Provides product configuration functions for Tivoli Storage Productivity Center.

This node is always present.

Tivoli Storage Productivity Center
Provides operational functions for Tivoli Storage Productivity Center.

This node is always present.

Data Manager
Provides operational functions for the Data Manager.

This node is present when Data Manager is installed and licensed.

For Data Manager functions, you must use a Storage Resource agent to collect data about a resource. For example, use scans and pings to collect information about computers. Use the web-based GUI to schedule probes and performance monitors to collect asset, status, and performance information.

Data Manager for Databases
Provides operational functions for Data Manager for Databases.

This node is present when Data Manager is installed and licensed.

Data Manager for Chargeback
Provides operational functions for Data Manager for Chargeback.

This node is present when Data Manager is installed and licensed.

Disk Manager
Provides operational functions for the Disk Manager.

This node is present when Disk Manager is installed and licensed. A subset of the functions is provided when you install the Data Server.

For Disk Manager functions, the subsystem CIMOMs and data collection are required (for example, through probes), except for resources that use the native API connection, for example: System Storage DS8000, SAN Volume Controller, and the XIV system.

Fabric Manager
Provides operational functions for the Fabric Manager.

This node is present when Fabric Manager is installed.

For Fabric Manager functions, the available agent types are CIM agents, out-of-band fabric agents, and Storage Resource agents. For Fabric Manager features such as performance monitoring, you must collect data from performance monitors. For a list of current agent types supported for the Switch Performance Management functionality and the Fabric Zone Configuration functionality, see "Fabric Manager" on page 1739.

Replication Manager
Set the triggering conditions and actions for alerts that are associated with events in IBM Tivoli Storage Productivity Center for Replication.

Dashboard

The dashboard provides a concise, yet detailed overview of the health of your storage environment. It quickly points out potential problem areas that need further investigation.

When you first log in to IBM Tivoli Storage Productivity Center you are asked if you want to collect statistics. You must answer yes in order for the dashboard information to be displayed. The dashboard then appears whenever you start the product or close all the active windows in a running session.

The dashboard contains four display areas and seven panels that you can cycle through to view your environment's storage information.

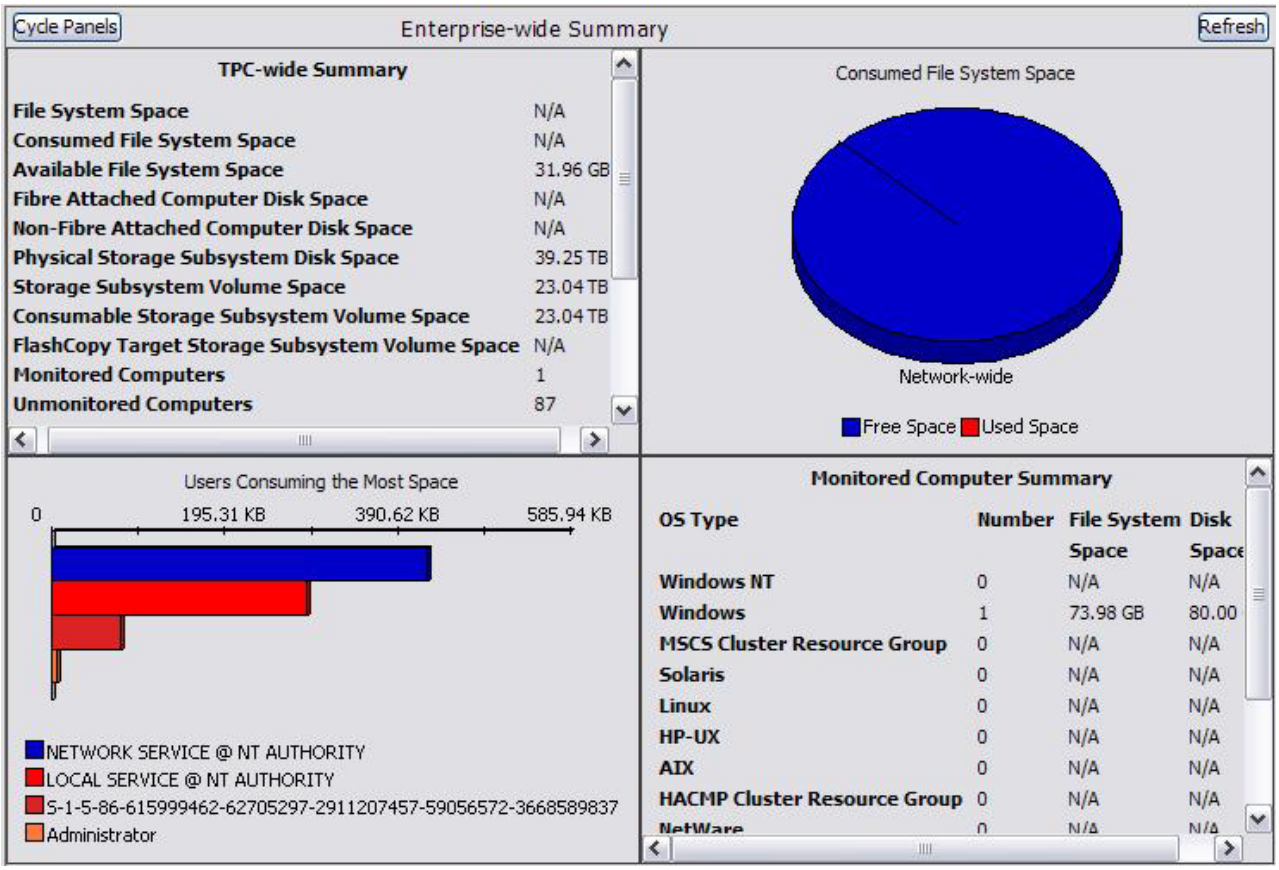

Using the dashboard:

Use the dashboard to monitor the overall status of your storage resources.

Purpose

When you first log on to Tivoli Storage Productivity Center or close all the active panels, the topic pane shows the "dashboard". The dashboard provides a concise, yet detailed overview of the status of your storage environment. Use it to quickly point out potential problem areas that need further investigation. The dashboard contains four displayable areas and eight panels that you can cycle through to view the storage information that has been collected about your environment.

Keep in mind the following considerations when viewing the information about the dashboard:

- The information in the dashboard is collected by data collection schedules and aggregated by the History Aggregator. To ensure that the latest information is displayed in the dashboard, run the history aggregator regularly . See the information center for instructions about how to create and schedule history aggregator jobs.
- The value in the **All** field on the Alerts Pending panel might be higher than the sum of the alerts shown in the other fields on that panel. This situation occurs because some alerts are not associated with a specific category, and are reflected in the value that is shown in the **All** field only.
- The names of some fields have been updated. For your reference, the original names of those fields are shown within parentheses ( ) next to their new names in this topic.

- Click **Refresh** to update the displayed panels. You must click **Refresh** to ensure that the latest information from the history aggregator is displayed in the dashboard.
- Click **Cycle Panels** to cycle through the panels in the dashboard and view information about your storage assets.

TPC-wide Summary panel

Use this panel to view an enterprise-wide summary of your storage through the following values:

File System Space (Capacity)
Total amount of unique file system storage space on computers where a Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems discovered on virtual machines where an agent is installed.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Consumed File System Space (Filesystem Used Space)
Total amount of unique file system storage space that is used or consumed by the computers where a Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems on discovered hypervisors that are assigned to virtual machines where an agent is installed.
- File systems on virtual machines where an agent is installed, but the hypervisors for the virtual machines have not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Available File System Space (Filesystem Free Space)
Total amount of unique file system storage space that is not used or available to the computers where a Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Fibre Attached Computer Disk Space (Computer Fibre Attached Disk Space)
Total amount of unique disk space on computers where a Storage Resource agent is installed and the disks are attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this field:

- Computer disks that are behind a fibre channel port that has not been identified by an agent.

- Computer disks that are not behind a fibre channel port.
- Computer disks discovered on virtual machines where an agent is installed.
- Computer disks that become missing after a probe.

Non-Fibre Attached Computer Disk Space (Computer Non-fibre Attached Disk Space) Total amount of unique disk space on computers where a Storage Resource agent is installed and the disks are not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this field:

- Computer disks that are behind a fibre channel port.
- Computer disks discovered on virtual machine where an agent is installed.
- Computer disks that become missing after a probe.

Physical Storage Subsystem Disk Space (Storage Subsystem Physical Disk Space) Total amount of unique internal physical disk space on the storage subsystem arrays that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this field:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller) that are attached from a storage subsystem.
- Storage subsystem disks that are on storage subsystems that have not been probed by the subordinate servers or the master server.

Storage Subsystem Volume Space (LUN Capacity)
Total amount of unique storage subsystem volume space and virtual storage volume space on the storage subsystem arrays that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this field:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller).
- Storage subsystem disks that are on storage subsystems that have not been probed by the subordinate servers or the master server.

Consumable Storage Subsystem Volume Space (Usable LUN Capacity)
Amount of file system storage space for a file system. Total amount of unique storage subsystem volume space and virtual storage volume space on monitored storage subsystem arrays that can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following columns in its calculation of the value for this field:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are used as FlashCopy targets.

- Storage subsystem volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller).

FlashCopy Target Storage Subsystem Volume Space (FlashCopy Target Capacity)
Total amount of unique storage subsystem volume space and virtual storage volume space on monitored storage subsystem arrays that is identified as FlashCopy target space. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this field:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are not used as FlashCopy targets.

Monitored Computers (Servers)
Total number of computers where a Storage Resource agent is installed and are detected by the Data server. Tivoli Storage Productivity Center does not include the following systems in its calculation of the value for this field:

- Systems where an agent is installed but was not started or was unable to communicate with the Data server.
- Systems where an agent is not installed.
- Systems where an agent is installed but that agent communicates with a different Tivoli Storage Productivity Center server.
- Systems that have been removed from the Administrative Services > Data Sources > Data/ or Storage Resource agents panel.
- Hypervisors

Unmonitored Computers (Servers)
Total number of computers found in a Windows domain that do not have a Storage Resource agent communicating with the Data server. Tivoli Storage Productivity Center does not include the following systems in its calculation of the value for this field:

- Systems that are not running a Windows operating system.
- Systems that are not in a Windows domain.
- Systems where an agent is installed and is communicating with a Data server.

Storage subsystems
Total number of storage subsystems that have been discovered and probed. Tivoli Storage Productivity Center does not include the following storage subsystems in its calculation of the value for this field:

- Storage Subsystems that have been discovered but not probed.
- Storage Subsystem that have not been discovered.
- Storage Subsystems that have been removed from the Disk Manager > Storage Subsystems panel.

Users Total number of user accounts discovered by Storage Resource agents.

Disks Total number of unique disks discovered by Storage Resource agent

probes, hypervisor probes, network storage filer probes, and storage subsystem probes. This value does not include disks that have become missing after a successful probe.

Storage Subsystem Volumes (LUNS)
Total number of unique storage volumes and virtual storage volumes discovered by storage subsystem probes. This value does not include storage volumes or virtual storage volumes that have become missing after a successful storage subsystem probe.

Filesystems
Total number of unique file systems discovered by probes. This value excludes file systems that have become missing after a successful probe.

Directories
Total number of directories discovered by scans. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this field:

- Directories within file systems that were discovered by a probe but were not scanned.
- Directories that have become missing after a successful probe.

Files Total number of files discovered by scans. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this field:

- Files within file systems that were discovered by a probe but were not scanned.
- Files that have become missing after a successful probe.

Rollup Summary panel

Use this panel to view a summary of information collected by the master server and by Tivoli Storage Productivity Center servers that are subordinate servers.

File System Space (Capacity)
Total amount of unique file system storage space on the computers that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems discovered on virtual machines where the Storage Resource agent is installed.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.
- File systems that have not been probed by subordinate servers or the master server.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computers that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems on discovered hypervisors that are assigned to virtual machines where the Storage Resource agent is installed.

- File systems on virtual machines where the Storage Resource agent is installed, but the hypervisors for the virtual machines have not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.
- File systems that have not been probed by subordinate servers or the master server.

Available File System Space (Filesystem Free Space)
Total amount of unique file system storage space that is not used or available to the systems that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.
- File systems that have not been probed by subordinate servers or the master server.

Physical Storage Subsystem Disk Space (Storage Subsystem Physical Disk Space)
Total amount of unique internal physical disk space on monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this field:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller) that are attached from a storage subsystem.

Storage Subsystem Volume Space (LUN Capacity)
Total amount of unique storage subsystem volume space and virtual storage volume space on monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this field:

- Storage subsystem volumes or virtual storage volumes that become missing storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller).

Consumable Storage Subsystem Volume Space (Usable LUN Capacity)
Total amount of unique storage subsystem volume space and virtual storage volume space on storage subsystem arrays that are monitored by subordinate servers and the master server and can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this field:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are used as FlashCopy targets.

- Storage subsystem volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller).
- Storage subsystem disks that are on storage subsystems that have not been probed by the subordinate servers or the master server.

Monitored TPC Servers
Total number of subordinate servers that are monitored by the master server.

Monitored Computers (Servers)
Total number of computers that are monitored by Tivoli Storage Productivity Center and are detected by the subordinate and the master server. Tivoli Storage Productivity Center does not include the following systems in its calculation of the value for this field:

- Systems where an agent is installed but was not started or was unable to communicate with the Data server.
- Systems where an agent is not installed.
- Systems where an agent is installed but that agent communicates with a Tivoli Storage Productivity Center server that is not a subordinate server or the master server.
- Systems that have been removed from the Administrative Services > Data Sources > Data/ or Storage Resource agents panel.

Storage subsystems
Total number of storage subsystems that have been discovered and probed by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following storage subsystems in its calculation of the value for this field:

- Storage Subsystems that have been discovered but not probed.
- Storage Subsystem that have not been discovered.
- Storage Subsystems that have been removed because of the storage subsystem settings on the Removed Resource Retention panel.

Consumed File System Space (Filesystem Used Space panel)

Use the chart on this panel to view the total amount of used file system storage space across the entire network.

Users Consuming the Most Space (chart) panel

Use this panel to view the following information:

- A list of the users consuming the most space within your environment
- A graphical view of the amount of space they are consuming

The chart on this panel shows information about users across the network. If multiple hosts contain the same user name, the value displayed for that user represents the space consumed on all the hosts. For example, if the user "root"

exists on three hosts that are monitored by Tivoli Storage Productivity Center, then the value displayed on this chart represents the space consumed by "root" on those three hosts.

Note: The space usage of administrative users within your environment is not included as part of this dashboard report.

Monitored Computer Summary panel

Use this panel to view the following information:

- The number of monitored computers running each Operating System type
- Total amount of file system and disk capacity per Operating System type

Filesystems with Least Free Space Percentage panel

Use this panel to view the monitored computers with the least percentage of free space, including the following values:

- Percent of free space on each computer
- Total file system capacity on each computer
- Name of the computers, file systems, and mount points

Users Consuming the Most Space Report panel

Use this panel to view the following information:

- A list of the users consuming the most space within your environment
- The total amount of storage each user is consuming

Alerts Pending panel

Use this panel to view the active alerts that are still pending.

Note: The number of alerts on this panel might be less than the number of alerts shown in the IBM Tivoli Storage Productivity Center > Alerting > Alert Log. Alerts for failed jobs (pings, probes, and scans) are not listed in separate rows on this panel, but are included in the total displayed in the **All** row.

These alerts are organized into the following categories:

- All
- Computer
- Disk
- Filesystem
- Directory
- User Quota
- OS User Group Quota
- Directed to *user*
- Storage Subsystem
- Switch
- Fabric
- Endpoint Device
- External
- Hypervisor

- Authentication Configuration
- Replication
- Configuration Analysis
- NAS Node
- NAS File Set
- Export

Working with the navigation tree and content pane

This section describes some methods for working with the nodes and windows in the navigation tree and content pane.

Right-clicking and pop-up menus:

Right-click nodes in the navigation tree to access menus that contain actions you can perform.

As you navigate through the IBM Tivoli Storage Productivity Center user interface, notice that you can right-click nodes in the navigation tree to access menus that contain actions related to that node. Use these actions to quickly access many product functions.

The following actions are available for the nodes in the navigation tree:

Option	Description
Create <object>	Create an object, including a group, data collection schedule, policy, batch report, or alert.
Delete	Delete the selected object definition, such as a data collection schedule, policy, batch report, group, or alert. This action is not available for definitions owned by TPCUser.
Rename	Rename the object definition. This action is not available for definitions owned by TPCUser.
Refresh	Refresh the list of objects in the navigation tree. The navigation tree is not updated automatically. You must click **Refresh** for a node to view the latest objects related to that node. For example, right-click the **Tivoli Storage Productivity Center** > **Monitoring** > **Probes** node and select **Refresh** to update the probe schedules that are listed in the navigation tree.
Job History	View the list of jobs related to a data collection or batch report schedule on the Job Management panel.
Run Now	Run a data collection or batch report schedule immediately.
Collapse Branch	Minimize the navigation tree under the selected node.
Expand Branch	Open the navigation tree under the selected node.
Print Branch	Print the contents of the navigation tree under the selected node.

Selecting panels in the content pane:

The content pane is displayed on the right side of the stand-alone GUI. When you click a node in the navigation tree, a corresponding panel is displayed in the content pane. Use the panels in the content pane to complete storage management tasks.

Each time you select a node in the navigation tree, the panel for the previously selected node is not deleted. Instead, it moves into the background and is accessible by clicking and on the toolbar. You can also select currently active panels from the **Window** > **Show Window List** in the menu bar.

In the **Window** menu, a **>** is shown next to the name of the active panel. Objects that you are editing (such as data collection schedules) are marked with a black arrow if they have no pending changes; a solid red arrow indicates objects that have pending changes.

Menu bar

Use the menu bar of the main window to access some of the main functions within IBM Tivoli Storage Productivity Center.

The following sections describe the options available from the menu bar.

File menu:

Use the options in the File menu to save objects and jobs, export report data, print, refresh alerts, close the current window, and exit the product session.

Menu Option	Description
Save	Save the changes for the current object (for example, groups, profiles, etc.) or job (for example, probes, scans, pings, etc.) that you are editing.
Save As	Save the changes for the current object or job that you are editing under a different name.
Export Data	Export the report data to other formats: • Formatted Report file (.txt) • Comma delimited file (.csv) • Comma delimited file with headers • HTML file This allows you to import the data to another application, such as Excel, if needed.
Print	Print the current data in the content pane. You can print tabular reports or report charts. You can also save the output to a PDF file or HTML file, and print those files at a later time.
Print Tree	Print the expanded contents of the navigation tree. You can also save output to a PDF or HTML file, and print those files at a later time.
Refresh Alerts	Refresh all the alerts from the database repository and update the GUI by highlighting the appropriate objects and showing all the new alerts in the alert logs.
Close	Close the current window that occupies the content pane.
Exit	End the current session.

View menu:

Use the options in the View menu to hide or show the navigation tree and highlight the node in the navigation tree for the currently open window.

Menu Option	Description
Tree	Hide and show the navigation tree. If you want to increase the size of the content pane to occupy the entire window, clear the check box for the **Tree** option. To redisplay the navigation tree next to the content pane in the main window, select the **Tree** option.
Current Page in Tree	Display the highlighted node on the navigation tree that was responsible for displaying the data on the content pane. Use this option to reorient the navigation tree with the window that is currently displayed in the content pane.

Connection menu:

Use the option in the Connection menu to end the current session and establish a new connection to the same IBM Tivoli Storage Productivity Center server or to a different server.

Menu Option	Description
New Connection	Initiate a new connection to a server from the current GUI session. For example, use this option to connect to the same server using a different user ID, or to a completely different server. All the current changes that are unsaved in the open windows will be lost when you establish a new connection.

Preferences menu:

Use the options in the Preferences menu to change the appearance of the user interface, edit the way the product manages open panels, determine what alerts are shown upon login, select what tab appears first when generating reports, and specify what unit of measurement to use when displaying storage sizes in reports.

Menu Option		Description
Look and Feel	**Windows**	Set the look and feel of the GUI to appear like Microsoft Windows GUI.
	CDE/Motif	Set the look and feel of the GUI to appear like the UNIX Motif interface.
	Metal	Set the look and feel of the GUI to appear like Metal, the cross-platform Java look and feel.

Menu Option		Description
Edit General	**Panel Retention**	The **Panel recall depth** value specifies the number of windows that are linked in the content pane for access via the and buttons on the toolbar. The default is 20. The **Panels kept in memory field** specifies the number of windows to store in memory, which helps to optimize the system performance. The default is 2. **Note:** Any open views in the topology viewer (for example, an L2:Storage Subsystem view, an L0:Computers view) are considered part of the topology viewer panel and are listed as tabs on the topology viewer panel. If the value in this field is a low number, such as the default of 2, then opening additional non-topology panels can exceed the number of panels that are kept in memory. This can result in the topology viewer panel refreshing and any tabbed views being closed. To avoid this situation, change the value in this field to a higher number, such as 10.
	On login, show	Control which alerts you will view upon logging in to a user interface session. You can select: • No Alerts • Alerts Directed to <Userid> • All Active Alerts
	Initial Reporting Tab to Display	Control what tab will be displayed when you first generate and view a report. You can select: • Selection tab • Report tab • Chart tab (if Available)

Menu Option	Description
Advanced Options	Select the unit of measurement that should be used to reflect the storage sizes included in IBM Tivoli Storage Productivity Center reports. This setting affects the results in online, exported, and Batch reports. You can choose from the following options: • **As Appropriate:** Choose this option to automatically determine the unit measurement (KB, MB, GB, TB) to display for storage sizes in a report. • **Kilobytes (KB):** Choose this option to display storage sizes in reports using the kilobyte unit of measurement. For example, if the storage capacity for a computer is 100 megabytes, it is shown as 102400.00 KB in a capacity report. • **Megabytes (MB):** Choose this option to display storage sizes in reports using the megabyte unit of measurement. For example, if the storage capacity for a computer is 1 gigabyte, it is shown as 1024 MB in a capacity report. • **Gigabytes (GB):** Choose this option to display storage sizes in reports using the gigabyte unit of measurement. For example, if the free space on a computer is 500 megabytes, it is shown as 0.49 GB in a free space report. • **Terabytes (TB):** Choose this option to display storage sizes in reports using the terabyte unit of measurement. For example, if the free space on a computer is 350 gigabytes, it would be shown as 0.34 TB in a free space report.

Window menu:

Use the options in the Windows menu to close all the open windows that do not have any changes, display the dashboard, and list all the open windows.

Menu Option	Description
Close all Windows Without Changes	Close all of the open windows that do not contain any modifications. Any windows that contain active changes to editable objects will not be closed.
Show Dashboard	Display the IBM Tivoli Storage Productivity Center dashboard.
Show Window List	Lists names of all open windows active in the content pane. Select one of these options to view the corresponding window. Beside each window name is an icon that indicates the status of the window: • **no icon** - No changes were made to the window • **solid red circle** - Changes were made, but have not been saved • **black arrow** (>) - Current window displayed. No changes were made • **solid red arrow** (>) - Current window displayed. Changes were made, but have not been saved

Help menu:

Use the options in the **Help** menu to access the online help, open the information center, open the product support page, open the demonstrations page, and display product information (such as the version number).

Menu Option	Description
Help Topics	Open the welcome page of the online help for IBM Tivoli Storage Productivity Center in a separate window. Use the online help to view information about how to use the product.
Help for Displayed Panel	Open the online help for the panel that is displayed. You can also press F1 to access the help for the panel.
Launch Information Center	Open the information center for IBM Tivoli Storage Productivity Center Information Center in a web browser. The IBM Tivoli Storage Productivity Center Information Center contains detailed instructions that explain how to install, upgrade, configure, maintain, and use the Tivoli Storage Productivity Center product. See "Web browser support" for information about web browser support for opening the IBM Tivoli Storage Productivity Center Information Center. An internet connection is required to access the information center. For information about web browser support, see "Web browser support".
Supported Products Data	Open the product support page in a web browser. Use this page to view a list of operating systems on which you can run Tivoli Storage Productivity Center and the resources that you can monitor. An internet connection is required to access the product support page. For information about web browser support, see "Web browser support".
Tivoli Storage Productivity Center Demos	Open a web page that contains flash demonstrations of Tivoli Storage Productivity Center. These demonstrations introduce the product, explain the business and storage administrator challenges that it can help you meet, and describe its features and functions. An internet connection is required to access the web page for flash demonstrations. For information about web browser support, see "Web browser support".
About	View the following information about Tivoli Storage Productivity Center: • Version information • Copyright information • Trademark information

Web browser support:

IBM Tivoli Storage Productivity Center starts a web browser when you access web pages from items in the **Help** menu or start another application by using its launch-in-context feature.

Note: The Tivoli Storage Productivity Center software is bundled with its own IBM Java JRE. When you use the Java Web Start method of remotely accessing the stand-alone GUI, you must use the IBM Java.

Before you start Tivoli Storage Productivity Center, ensure that you are using a supported web browser. For a list of web browsers that you can use with Tivoli Storage Productivity Center, see the support matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446. In the **Agents, Servers and GUI** column, click the version of Tivoli Storage Productivity Center that you are installing. On the next page, click **Web Browsers** to find the web browsers that you can use.

If you do not have a web browser that is configured for use with Tivoli Storage Productivity Center when you access a web page from its user interface, the Configure Element Launcher window opens. Use this window to configure a web browser for use with the product.

Tip: For information about the web browsers that you can use with the DS8000, see the DS8000 Information Center at http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp. Search for **Internet browser support**.

Note: The web GUI requires the use of cookies to manage user preferences, such as table customization and frequently viewed charts. Ensure that cookies are enabled to use these features.

The software to block pop-up windows must be turned off in Windows Internet Explorer for any application to start the graphical user interface (GUI) for the following storage systems:

- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified
- IBM SONAS

For example, if you are in Tivoli Storage Productivity Center and you want to start the Element Manager of the SAN Volume Controller GUI, you must configure Internet Explorer with the following settings.

1. Open Internet Explorer.
2. From the Internet Explorer toolbar, click **Tools** > **Pop-up Blocker** > **Turn Off Pop-up Blocker**.

If the Internet Explorer **Information Bar** displays a message that content is blocked because it was not signed by a valid security certificate, click the **Information Bar** and select **Show blocked content**.

Web browser support for Help menu items

Tivoli Storage Productivity Center provides a number of items in its **Help** menu that you can click to learn more about how to use the product. When you click one of these items, Tivoli Storage Productivity Center starts a web browser and displays the appropriate web page. You must have an internet connection on the system where you are running the Tivoli Storage Productivity Center user interface. You must also have a default browser that is configured to successfully access these web pages.

Toolbar

The toolbar on the Tivoli Storage Productivity Center main window provides one-click access to key commands.

Click the following buttons on the toolbar to access key commands in the product:

Button	Description
Back	Click this button to move backward through the active panels in the content pane.
Forward	Click this button to move forward through the active panels in the content pane.
Save	Click this button to save the current object that you are editing. The object changes are saved to the database repository.
Print	Click this button to print the current data or report chart that is displayed in the content pane.
Cancel edit session	Click this button to discard the contents of the content pane. If the content pane is in an edit session, you are notified if any changes were made.
Manage jobs, schedules, and runs	Access the Job Management panel. Use this panel to view and manage the schedules, runs, and jobs that are related to the monitored resources.
Stop current server request	Click this button to stop the current server request. This button becomes active when a request is made to the server and the user interface is waiting for a response. During the server request wait time, the cursor arrow is displayed as an hourglass. If you decide the current request is taking too long or you do not want to wait any longer, click this button to cancel that request and regain control of the GUI. **Note:** If you click this button, the current request continues to run, and the resources that are used by it on the server are not released until it completes. Frequent use of this button during a session might exhaust the resources on the server and cause the user interface to stop responding.

Data Manager

Data Manager is a comprehensive file and capacity management solution for heterogeneous storage environments. Data Manager includes enterprise-wide reporting and monitoring, policy-based management, and automated capacity provisioning for Direct Attached Storage (DAS), network attached storage (NAS), and SAN environments.

Data Manager helps you improve storage usage, plan for future capacity, and ensure availability by providing storage on demand for file systems. Use Data Manager to perform the following functions:

- Discover and monitor disks, partitions, shared directories, and servers.
- Monitor and report on capacity and utilization across systems to help you to identify trends and prevent problems.
- Provides a wide variety of standardized reports about file systems and storage infrastructure to track usage and availability.
- Provide file analysis across systems to help you to identify and reclaim space that is used by non-essential files.
- Provide policy-based management and automated capacity provisioning for file systems when user-defined thresholds are reached.

Using these functions, Data Manager helps you reduce storage costs by:

- Improving storage usage
- Enabling intelligent capacity planning
- Helping you manage more storage with the same staff
- Supporting application availability through computer uptime reporting and application database monitoring.

Information that is collected by Data Manager helps you understand what is really going on with data on your servers and in your storage environment. View when files are created, accessed, and modified, and by what group or user. This type of information helps system administrators map storage resources to the consumers of the resource. The ability to map storage consumption to storage hardware becomes increasingly important as the size of open systems environments increases.

In addition to understanding the current consumption and usage of data within the enterprise, Data Manager tracks the information over time. As this historical view of storage utilization show usage trends over time, the system administrator can also see a projected use of storage in the future. System administrators can prepare for the necessity to purchase more capacity in a planned proactive manner rather than reacting to out-of-space emergencies.

Use Data Manager policy functions to help you evaluate and control the usage and status of your enterprise storage management assets. Because Data Manager is policy-based, it has autonomic self-healing capabilities. Data Manager can detect potential problems and automatically adjusts based on the policies and actions you establish. Use the capability to provision storage based on storage policies to expand a file system, and to allocate storage to a volume.

For example, Data Manager can notify you when your system is running out of disk space or warn you of impending storage hardware failure. By alerting you to these and other issues that are related to your stored data, use Data Manager to prevent unnecessary system and application downtime.

Data Manager for Databases:

Data Manager for Databases provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance and assets in your relational databases. It can help you identify, evaluate, control and predict the storage needs of Relational Database Management Systems (RDBMSs), which include Oracle, Sybase SQL Server, Microsoft SQL Server, and UDB/DB2.

Data Manager for Databases is policy-based and through its autonomic self-healing capabilities it can detect potential problems and automatically make adjustments based on the policies and actions you have established. For example, it can notify you when your database tables are running out of storage space or warn you of a dropped tablespace. By alerting you to these and other issues related to your stored data, it enables you to prevent unnecessary system and application downtime.

The following table lists the RDBMS objects whose storage usage is monitored by Data Manager for Databases.

Oracle	SQL Server	Sybase	DB2/UDB
Instances	**Instances**	**Instances**	**Instances**
Databases	**Databases**	**Devices**	**Databases**
Tablespaces	**Datafiles**	**Databases**	**Tablespaces**
Datafiles	**Tables**	**Fragments**	**Containers**
Tables	**Indexes**	**Tables**	**Tables**
Indexes	**Log Files**	**Indexes**	**Indexes**
Extents			
Segments			
Redo Log files			
Control Files			
Archive Log Directories			

Note: Data Manager for Databases does not currently support the monitoring of clustered database applications.

Data Manager for Databases collects information about RDBMS storage assets and performs tasks against the RDBMS storage assets. Methods for managing and automating capacity utilization of your databases help you:

- Monitor storage assets associated with enterprise-wide databases and notification of potential problems before they occur
- Make intelligent capacity management decisions based on current and trended historical data
- See how much storage is being consumed by users, groups of users and OS's within the database
- Create policy-based management for databases when user-defined threshold is reached
- View storage utilization management from a database and application perspective

Data Manager for Chargeback:

Data Manager for Chargeback uses the storage usage information gathered by Data Manager and Data Manager for Databases to generate invoices that chargeback for storage usage.

With Data Manager for Chargeback you can collect usage information on a departmental, group, or user level. You can allocate costs by storage usage by user, disk capacity by computer, tablespace, or file system/physical device. You can create cost centers by creating user, computer, or tablespace groups, allowing organization to chargeback individuals or business units for their storage usage. By

understanding the costs associated with existing and future storage usage, you can improve the use of that storage and reduce the costs associated with its maintenances and upgrades.

Using Data Manager, you can run monitoring jobs that gather detailed statistics about the usage and trending of the storage consumed by the users and user groups within your organization. Using Data Manager for Databases, you can run monitoring jobs that gather detailed statistics about the usage and trending of the storage consumed within RDBMS Instances by the users and user groups within your organization.

In addition to providing invoicing for storage usage, Data Manager for Chargeback also integrates with the chargeback systems already implemented in your environment. It provides you with a higher level, application-specific CIMS output format which can be imported into CIMS applications. With Data Manager for Chargeback, you can export chargeback data for direct import into CIMS, giving you the ability to integrate your storage usage data with other enterprise chargeback information and processes.

Disk Manager

Disk Manager helps you manage SANs and heterogeneous storage from a single console.

With Disk Manager you can manage network storage components such as:

- IBM System Storage SAN Volume Controller
- IBM Storwize V7000
- Storwize V7000 Unified
- IBM TotalStorage Enterprise Storage Server (TotalStorage Enterprise Storage Server)
- Tivoli Storage Productivity Center disk systems (DS4000, DS5000, DS6000, and DS8000 series)
- IBM XIV Storage System
- IBM SONAS
- Other storage subsystems that support the SMI-S standards

For those devices that use the native interfaces, device discovery is performed when the device is added to Tivoli Storage Productivity Center.

For those devices that use the SMI-S standards, device discovery is done using the Service Location Protocol (SLP), as specified by SMI-S. Configuration of the discovered devices is possible in conjunction with CIM agents associated with those devices, using the standard mechanisms defined in SMI-S. Disk Manager also gathers events, and can launch an element manager specific to each device.

Disk Manager can:

- Collect and store performance data and provides alerts
- Provide graphical performance reports
- Help optimize storage allocation
- Provide volume contention analysis

After data collection is done, you can set thresholds to identify exceptions for certain performance metrics, depending on the type of device. Threshold checking is performed during data collection, and when performance is outside the specified bounds, alerts can be generated.

After performance data has been collected, you can configure Disk Manager to present graphical or text reports on the historical performance behavior of specified devices, or of specific internal components of those devices. The performance reports provide information about past period performance metrics or current performance in graphical form.

Note: Tivoli Storage Productivity Center supports discovery of and reporting on both FB and CKD volumes (and pools and arrays), but can create only FB volumes. Tivoli Storage Productivity Center does not support the creation of CKD volumes.

Managing TagmaStore CIM agents:

The TagmaStore CIM agents are provided by Hitachi Data Systems for the TagmaStore storage subsystem. The TagmaStore CIM agent collects information from the TagmaStore storage system.

IBM Tivoli Storage Productivity Center supports the Hitachi Data Systems TagmaStore CIM Agent. For information about the CIM agent level supported, see http://www.ibm.com/support/docview.wss?rs=40&q1=support+matrix&uid=swg21386446&loc=en_US&cs=utf-8&lang=en. Click the appropriate version under **Storage**. Tivoli Storage Productivity Center also supports the TagmaStore as a Storage Virtualizer. Tivoli Storage Productivity Center can display information for virtual disks and local disks.

Tivoli Storage Productivity Center cannot provide correlation information between the virtual storage used by TagmaStore and the volumes created on the storage due to an existing limitation of the CIM agent 5.8 from Hitachi Data Systems. However, Tivoli Storage Productivity Center reports correctly display the correlation between volumes created on the local storage and the local disks. This limitation has no impact on the topology but it does affect several reports that show the correlation:

- **Data Manager > Reporting > Asset > By Storage Subsystem > <HDS_device> > Disks > Volumes** (does not show the relation between the disk and Volume or volume in the tree)
- **Disk Manager > Reporting > Storage Subsystem > Volume to Backend Volume Assignment**
- **Disk Manager > Reporting > Storage Subsystem > Computer Views** (because Tivoli Storage Productivity Center cannot populate the information for volumes created on the virtual disks)

All volumes are created from a storage pool that is allocated from a primordial storage pool and an imported primordial pool. A volume cannot be created over both local and virtual extents.

For volume correlation, the host computer must have the Storage Resource agent deployed and the TagmaStore device must be in the same SAN fabric. There must also be a zone configured in the active zoneset between the ports of the host computer and the ports of the TagmaStore device. The Storage Resource agent must be configured for the fabric to which the host is connected.

For back-end correlation, the TagmaStore device ports and back-end subsystem ports must be in the same zone and the back-end subsystem has assigned storage volumes to all ports of the TagmaStore device.

Fabric Manager

Fabric Manager helps you manage the SAN fabric that connects the host systems and applications to the storage devices. It is a comprehensive management solution for multi-vendor SANs and includes automatic resource and topology discovery, monitoring, and alerts.

Fabric Manager is an enterprise scalable solution architected to ANSI SAN standards, allowing you to choose the appropriate products for your environment. It helps you:

- Simplify the task of SAN management and configuration
- Ensure SAN availability
- Improve SAN return on investment

Fabric Manager is able to do SAN fabric performance and capacity management reporting and monitoring. Zoning is one of many SAN security options. Zoning is the only configuration supported for this release. Along with subsystem based volume masking, it is almost universally employed to ensure that only systems that are authorized can access selected data. Zoning is provided by the switch at the port level, so that, for example, a host on port X can access the subsystem connected by port Y. When business needs change, often zones must change also.

The ability to do switch performance and capacity management reporting and monitoring can help you to determine if more bandwidth is needed. Wide-area and local-area IP networks, and storage area networks (SANs) move data from one place to another. The management of the bandwidth is needed to continually monitor for link utilization and errors. Data needs to be gathered to tune resources, balance workloads, and do capacity planning.

With Fabric Manager, the comprehensive management of multi-vendor SANs can help simplify the IT infrastructure by providing automatic resource and topology discovery, monitoring, and alerts. It brings all the sources of information about SAN topology and configuration into a single place, and creates topology mapping of the SANs. This topology offers both host-centric and device-centric views that can be displayed graphically. The SAN topology display tracks all topology and configuration changes through out-of-band and SMI-S monitoring. Without this type of centralized topology, information from a number of sources, such as Element Management tools, device logs, and SNMP traps would have to be continually monitored and manually correlated to determine the current SAN configuration and topology.

Fabric Manager provides you with the ability to view events happening in your SAN environment and records state changes. The events are displayed in a color-coded fashion and can be further customized to reflect organizational priorities. It will forward events signaling topology changes or updates to the IBM Tivoli Enterprise Console, another SNMP manager, or both. Source side filtering by the Fabric Manager enables you to control what events get sent to your centralized console. Filtering helps ensure that the administrators get the information they need and are not overwhelmed by extraneous information.

Discovering storage resources

Run discovery to collect initial information about the Windows servers, NAS filers, and IBM SAN File Systems in your environment that are not monitored by Tivoli Storage Productivity Center.

In Tivoli Storage Productivity Center V5.2 and later, the discovery of storage systems, switches, fabrics, and hypervisors is done automatically when you add resources in the web-based GUI. For more information about how to add and discover these resources, see "Adding resources" on page 257.

Discovering computers in the Windows domain, NAS, and SAN File System

This topic describes how a NAS discovery finds SAN File System clusters.

When you run a discovery to find Windows computers on a domain, it discovers only new Windows computers. For IBM Tivoli Storage Productivity Center to recognize a non-Windows computer, you must install an agent on that computer.

For a Windows domain, Tivoli Storage Productivity Center only performs a discovery on the first Windows system on the domain; not all the systems on the domain.

A NAS discovery finds NAS Filers, and a SAN File System discovery finds the SAN File System clusters.

For information about configuring your system for NAS and SAN File System discovery, see the section about the Manual NAS/NetWare Server Entry node in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide*.

Discovering Windows computers that do not have agents:

Use discovery schedules to detect Windows computers that do not have Storage Resource agents deployed on them.

To discover Windows computers that do not have agents in your Windows domain, complete the following steps:

1. In the navigation tree, expand **Administrative Services** > **Discovery**.
2. Right-click **Windows Domain, NAS, and SAN FS** and click **Run Now**.
3. When you identify the computers, you can deploy Storage Resource agent on the computers to get detailed information about the computers and file systems. Computers that are discovered by Tivoli Storage Productivity Center, but do not have Storage Resource agents deployed on them, are displayed on the Data Manager > Reporting > System-wide > Unmanaged Computers report.

Discovering fabrics with Storage Resource agents

When you deploy Storage Resource agents to computers, the agents automatically run discovery and provide information about the fabrics or switches to which the computers are attached. You do not need to run a specific discovery step for the fabrics connected to computers on which Storage Resource agent are deployed.

For example, if you have QLogic switches and have Storage Resource agents deployed on the computers that are attached to these switches, the switches are automatically discovered when you deploy the Storage Resource agents.

Discovering fabrics with the Storage Resource agent provides the following information:

- Fabric WWN, information about fabric-to-switch relationships, and key attributes
- Host and device information (collected from Storage Resource agents only), including HBA information
- Basic information to identify the host and devices in the fabric

After a discovery is run, you can run probes to collect topology and zoning information for the fabrics.

The vendors of fabric switches can help you determine whether to use CIM agents or Storage Resource agents to discover and monitor fabrics. For example, for Brocade switches, use CIM agents to collect all fabric data and enable performance monitoring. Use Storage Resource agents to collect information about HBAs. For Cisco and QLogic switches, you must have one or more Storage Resource agents on the fabric to collect full topology and zoning information.

See the *Planning for Fabric Manager* section in the *Tivoli Storage Productivity Center Information Center* or *Installation and Configuration Guide* for more information about using CIM and Storage Resource agents to collect information about fabrics.

Checking storage configurations

This section introduces methods that a system administrator can use to assess changes that have occurred in SAN configurations over time and to check whether SAN configurations comply with predefined policies.

Using configuration history

Capture and analyze historical data that identifies possible problems with a storage area network (SAN) configuration.

Note: The configuration history feature is available only with the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license.

The configuration history feature takes and displays snapshots of changes that occurred in your SAN configuration over a period that you specify. After you set the time period (how often to take snapshots and how long to store them), in a page similar to the topology viewer you can manipulate a snapshot selection panel to show changes that occurred between two or more points in the time period. System administrators can use the configuration history feature to:

Correlate performance statistics with configuration changes
For example, during the collection of performance statistics (including volume performance statistics) on an TotalStorage Enterprise Storage Server system you might delete a volume. While no new statistics are reported on that volume, the Tivoli Storage Productivity Center Performance Manager would have already collected partial statistical information before the volume was deleted. At the end of the data collection task, reporting of the partially collected statistics on the (now) deleted volume would require access to its properties which would not be available. The configuration history feature, with its ability to take and store snapshots of the configuration of a system, could provide access to the properties of the volume.

Analyze end-to-end performance
You want to know why performance changed on volume A during the last 24 hours. To learn why, it is useful to know what changes were made to

the configuration of the storage subsystem that might affect the performance of the volume, even if performance statistics were not recorded on some of those elements. For example, even if performance statistics on a per-rank basis are not collected, but the number of volumes that are allocated on a rank is increased from 1 to 100 over time, access to that configuration history information helps with analyzing the degraded performance of the volume over time.

Aid in planning and provisioning
The availability of configuration history can enhance the quality of both provisioning and planning. For example, historical data is useful when you provision storage in the web-based GUI.

To use the configuration history feature, complete the following steps:

1. In the Navigation tree pane, expand **Administrative Services** > **Configuration** > **Configuration History Settings**. The Configuration History Settings page displays for you to indicate how often to capture SAN configuration data and how long to retain it.
2. Perform the following to collect historical data:
   a. In the **Create snapshot every** field, click the check box to enable this option and type how often (in hours) you want the system to take snapshot views of the configuration.
   b. In the **Delete snapshots older than** field, click the check box to enable this option and type how long (in days) you want the snapshots to be stored.
   c. The page displays the total number of snapshots in the database and the date and time of when the latest snapshot was taken. To refresh this information, click **Update**.
   d. To optionally create and title a snapshot on demand, in the **Title this snapshot** field type a name for the on-demand snapshot and click **Create Snapshot Now**. If you do not want to title the on-demand snapshot, click **Create Snapshot Now**.
   e. To return to the default settings (the default settings are create snapshots every 12 hours and delete snapshots older than 14 days), click **Reset to defaults**.
   f. To save your settings, click **File** > **Save**.
3. In the Navigation tree pane, expand **Tivoli Storage Productivity Center** > **Analytics** > **Configuration History**, and click **Configuration History**. The software loads the snapshot data for the length of time that you specified. The Configuration History page (a variation of the topology viewer) displays the configuration entities and a floating snapshot selection panel. In the panel, you can define the periods against which the configuration is compared to determine whether changes occurred.

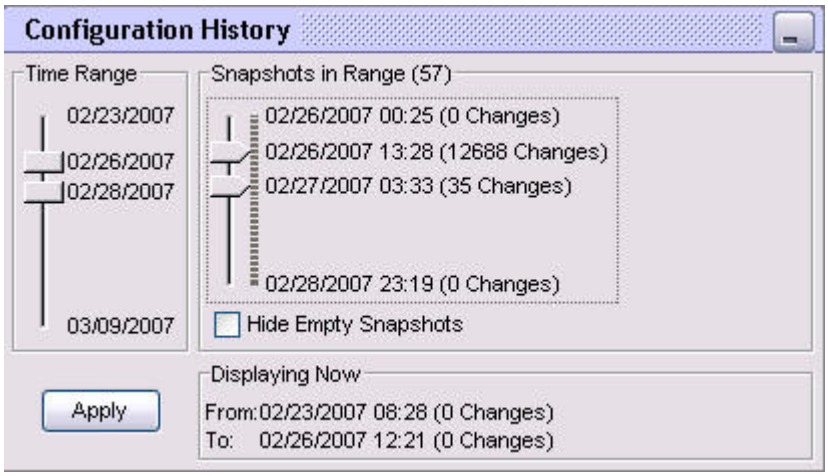

Figure 20. Snapshot selection panel. Use the thumb sliders to establish the time interval that you want to examine.

4. To define the time periods that you want to compare, complete the following steps:
   a. Using the mouse, drag the two "thumbs" in the **Time Range** slider to establish the time interval. The Time Range slider covers the range of time from the oldest snapshot in the system to the current time. It indicates the date as *mm/dd/yyyy*, where *mm* equals the month, *dd* equals the day, and *yyyy* equals the year.
   b. Drag the two thumbs in the **Snapshots in Range** slider to indicate the two snapshots to compare. With the Snapshots in Range slider, you can select any two snapshots from the time interval that is specified by the Time Range slider. The value in parentheses for the Snapshots in Range slider indicates the total snapshots in the currently selected time range. The Snapshots in Range slider has one tick mark for each snapshot from the time interval that you specified in the Time Range slider. Each snapshot in the Snapshots in Range slider is represented as time stamp *mm/dd/yyyy hh:mm*, where the first *mm* equals the month, *dd* equals the day, *yyyy* equals the year, *hh* equals the hour, and the second *mm* equals the minute. The value in parentheses for each snapshot indicates the number of changes that occurred between the current and the previous snapshot. Snapshots with zero changes are called empty snapshots. If you provided a title when you created an on-demand snapshot, the title displays after the time stamp. If you want to remove empty snapshots, click the check box to display a check mark in **Hide Empty Snapshots**. The **Displaying Now** box indicates the two snapshots that are currently active.
   c. Click **Apply**.
5. Determine the changes that occurred to the entities by examining the icons and the colors that are associated with them in the graphical and table views: For information about viewing the changes, see "Viewing configuration changes in the graphical view" on page 493 and "Viewing configuration changes in the table view" on page 495.

One single snapshot selection panel applies for all Configuration History views that are open at the same time. Any change that you make in this panel is applied to all of the Configuration History views.

Important: If you manually deleted a missing entity from the Tivoli Storage Productivity Center, the configuration history feature rediscovers and displays the

entity in the configuration history viewer as both Missing and in a Normal state. For example, Entity A (with ID 12) is missing and you decide to manually delete it by selecting the **Remove From Database** option from the pop-up menu for Entity A in the topology viewer or configuration history viewer. Later, the Tivoli Storage Productivity Center rediscovers Entity A and assigns it a new ID (ID 22). The configuration history feature treats the same Entity A (with ID 12 and ID 22) as two different entities and displays them in both Missing and Normal states in the configuration history viewer. This scenario can also occur if a missing entity is removed by the Tivoli Storage Productivity Center removed resource retention function and is rediscovered by the Tivoli Storage Productivity Center again at a later time.

Important: Name changes for volumes, storage pools, and managed disks are not highlighted as being changed in the Configuration History topology viewer.

Viewing configuration changes in the graphical view:

This section describes how to determine the changes that have occurred to a configuration over time by examining the icons and colors of the change overlay in the graphical view of the Configuration History page.

In the Configuration History page, a *change overlay* presents icons and colors that indicate changes in the configuration between the time that a snapshot was taken and the time that a later snapshot was taken:

- The icons display next to the name of the entity in the graphical view and in the change status column of the table view.
- The colors display as background colors for entities in the graphical view and for affected rows in the table view.

Table 74 describes the icons and colors of the change overlay.

Table 74. Icons and colors of the change overlay

Change overlay icon	Change overlay color indicator	Description
Yellow pencil	Blue background	Entity changed between the time that the snapshot was taken and the time that a later snapshot was taken.
Light gray circle	Dark gray background	Entity did not change between the time that the snapshot was taken and the time that a later snapshot was taken.
Green cross	Green background	Entity was created or added between the time that the snapshot was taken and the time that a later snapshot was taken.
Red minus sign	Red background	Entity was deleted or removed between the time that the snapshot was taken and the time that a later snapshot was taken.
Not applicable	Light gray background	Entity did not exist at the time that the snapshot was taken or at the time that a later snapshot was taken.

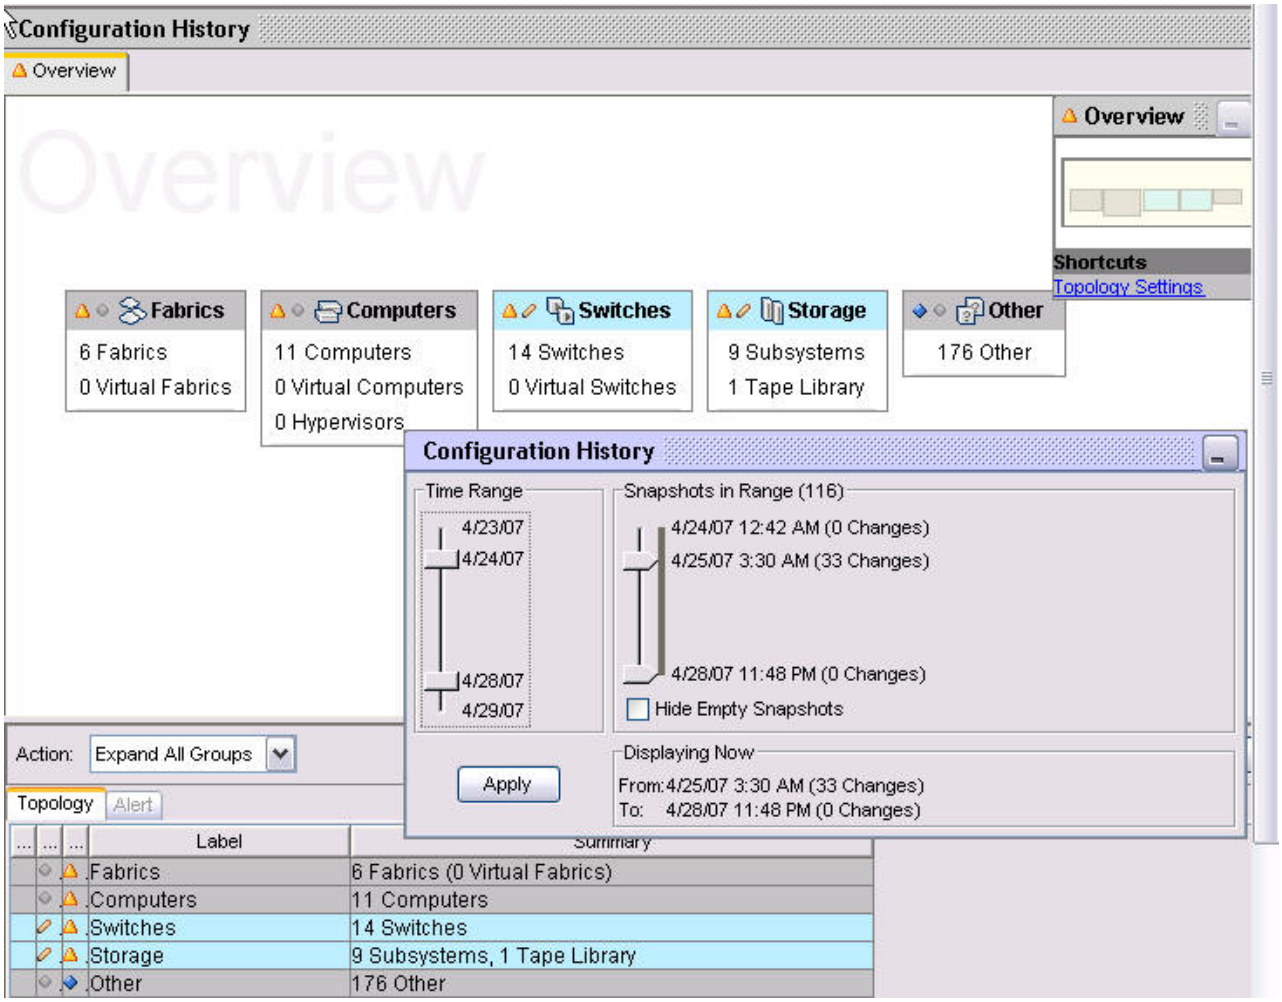

Figure 21. Icons and colors of the change overlay. In the graphical view, the pencil icon next to the switches and storage entities and the blue background color indicate that change occurred to these entities. The pencil icon and blue background also appears for these entities in the table view. In the snapshot selection panel, use the Time Range and Snapshots in Range sliders to determine when the change occurred.

To distinguish them from tabs in the topology viewer page, tabs in the Configuration History page (Overview, Computers, Fabrics, Storage, and Other) have a light gray background and are outlined in orange. The minimap in the Configuration History page uses the following colors to indicate the aggregated change status of groups:

- **Blue** - One or more entities in the group have changed. Note that the addition or removal of an entity is considered a change.
- **Gray** - All of the entities in the group are unchanged.

Entities in the graphical view may be *active* (they existed at one or both snapshots) or *inactive* (not yet created or deleted). *Active entities* act like they normally do in the topology viewer; when you select them all relevant information also appears in the table view. You can adjust a grouping of active entities, but you may not perform actions that change the database, such as pinning. *Inactive entities* do not exist in the selected snapshots, but exist in the other snapshots. They are shown in a light gray background and do not have a change icon associated with them. Inactive entities display to keep the topology layout stable and to make it easier to

follow what has changed (instead of having the entities flicker in and out of existence when you change the snapshot selection). Inactive entities are not listed in the table view.

An entity that is moved from one group to another group appears only once in the new group in the graphical view. For example, if the health status of a computer has changed from Normal to Warning, the Configuration History page displays the computer as changed in the Warning health group (and no longer displays the computer in the Normal health group).

Note: In the Configuration History view, the performance and alert overlays are disabled and the minimap's shortcut to the Data Path Explorer is not available.

Viewing configuration changes in the table view:

This section describes how to determine the changes that have occurred to a configuration over time by examining the icons and colors of the change overlay in the table view of the Configuration History page.

In the table view, a changed cell or tab is color coded. The change overlay uses the same colors and icons for the table view as for the graphical view (see Table 74 on page 493). The changed cell or tab displays both the old and new values as **Old Value** > **New Value**. To display a summary of change for the entire row, hover over the changed row. The summary displays the title of the column. It also displays the old and new value for each cell or tab that changed. For example, if a group contains one new entity but everything else in the group is unchanged, the group as a whole displays as “changed.” If you click the group, the one new entity displays as “created” and all other entities display as “unchanged.”

The Change column in the table view also indicates changed entities.

An object that is moved from one group to another group appears only once in the new group in the graphical view. For example, if the health status of a computer has changed from Normal to Warning, the Configuration History page displays the computer as changed in the Warning health group (and no longer displays the computer in the Normal health group).

Using configuration analysis

This topic describes how system administrators can determine whether an existing SAN configuration complies with predefined best practices (policies).

Note: The configuration analysis feature is available only with the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license

The configuration analysis feature allows a system administrator to select up to 16 policies and to specify the name of a target SAN (or zone set) for the analysis. The analysis tool reads the specified SAN data and the policy information for the specified policies from the IBM Tivoli Storage Productivity Center database. It outputs the policy violations as Tivoli Storage Productivity Center alerts. To view an alert, you can use one of three methods:

- Click the alert in the Configuration Analysis alert log
- Click the alert overlay icon ( ) next to the affected entity
- Click the Alert tab in the table view of the topology viewer

System administrators can use configuration analysis to learn of any best-practices violations their actions might have caused. For example, in provisioning a new storage subsystem, an administrator might inadvertently violate a zoning-related policy. If the zoning-related configuration changes persist in the database, configuration analysis can detect this policy violation.

To perform configuration analysis, complete the following steps:

1. Ensure that you have previously run discovery and probe jobs for the computers, fabrics, switches, storage, and other objects of interest.
2. Ensure that the Tivoli Storage Productivity Center Data Server and Device server are running. To do so, in the navigation tree pane, expand the following:
   - **Administrative Services** > **Services** > **Data Server** > **Server**
   - **Administrative Services** > **Services** > **Device Server** > **Device server**
3. Ensure that the following Tivoli Storage Productivity Center agents are installed and configured.

   Tip: The agent information that is required to evaluate a policy for violations must be collected and stored in the Tivoli Storage Productivity Center database.
   - You can have a combination of Storage Resource agent, out-of-band agent, and CIMOM fabric agents to collect all fabric topology information. This information includes switches, interconnections, and zoning, and is required to evaluate all configuration analysis policies.
   - You can have CIMOM agents for storage subsystems, which collect the storage subsystem information. This information includes subsystem details, storage volumes, storage ports, and storage volumes to hosts mappings, and is required to evaluate configuration analysis policies 2, 3, 5, 7 and 11. CIMOMs are not required to communicate with storage systems that use native interfaces. Tivoli Storage Productivity Center uses native interfaces to connect to the following storage systems:
     - DS8000
     - SAN Volume Controller
     - Storwize V3500
     - Storwize V3700
     - Storwize V7000
     - IBM Storwize V7000 Unified
     - The XIV system
     - IBM SONAS
   - Data agents or Storage Resource agents to collect the host information (operating system type and version). This information is needed for the evaluation of configuration analysis policies 3 and 9.
   - Storage Resource agents to collect the attribute information from the host information (HBA and host port details). This information is needed for the evaluation of configuration analysis policies 2, 3, 4, 8 and 9.
4. In the navigation tree pane, expand **IBM Tivoli Storage Productivity Center** > **Analytics**.
5. Right-click **Configuration Analysis**. A selection list displays.
6. Click **Create Analyzer**. The Create Analyzer window opens (see Figure 22 on page 497).

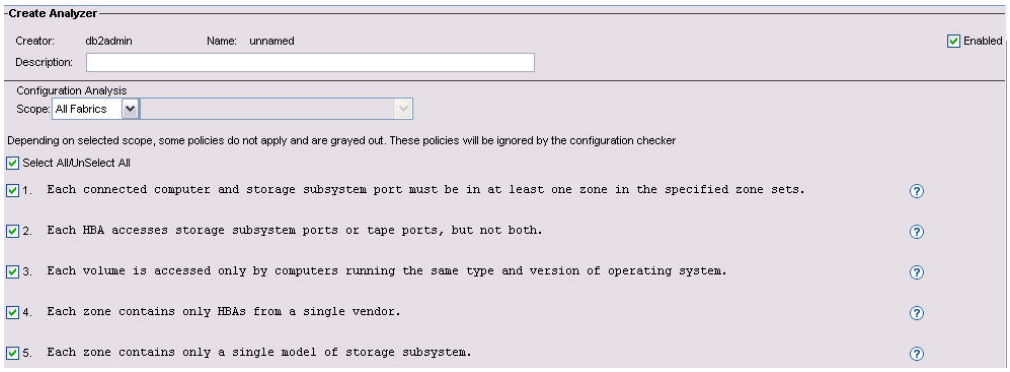

Figure 22. The Create Analyzer window. The figure shows a partial view of the window. Use the Create Analyzer window to select the scope and the policies of the configuration analysis.

7. To define the analysis, do the following steps:
   a. Select the check box for **Enabled**.
   b. In the **Description** field, type a brief description of the analysis job.
   c. To specify the scope of the SAN data to be checked, make a selection from the **Configuration Analysis Scope** list. Choices include:
      - All Fabrics
      - One fabric
      - One Zoneset

      If you select One fabric or One Zoneset, an adjacent list displays for you to click a specific fabric or zone set.
   d. Check the SAN data against up to 16 policies by performing one of the following:
      - Choose all of the policies by selecting the **Select All/UnSelect All** check box.
      - Choose one or more individual policies by selecting the check box for each policy.

      Note: For Policies 12 and 13, type values up to a maximum of 9999.

   The policies are described in the following list. You can also access each explanation by clicking ⓘ for that policy.

   Each connected computer and storage subsystem port must be in at least one zone in the specified zone sets.
   Determines whether an administrator forgot to zone a connected port. Putting connected ports into zones is useful for security and performance reasons. Ports are usually grouped into zones based on applications, server operating systems, or HBA vendors. The Fabric scope is not supported by this policy.

   Each HBA accesses storage subsystem ports or tape ports, but not both.
   Determines whether an HBA accesses both storage subsystem and tape ports. Because HBA buffer management is configured differently for storage subsystems and tape, it is not desirable to use the same HBA for both disk and tape traffic. The Fabric and Zone Set scopes are not supported by this policy because an HBA can be connected to multiple fabrics.

Each volume is accessed only by computers running the same type and version of operating system.
Determines whether computers that run different operating systems access the same storage volumes. Use of the same volumes by computers that run different operating systems may corrupt the data that is stored on the volumes. This applies, regardless of whether the computers are in the same zone.

Each zone contains only HBAs from a single vendor.
Determines whether HBAs from different vendor types are in their own zone. Receiving a registered state change notification (RSCN) may cause an HBA to lose a zoned device, preventing the HBA from seeing or communicating with other devices in the zone. To avoid losing a zoned device, keep HBAs from different vendor types in their own zone.

Each zone contains only a single model of storage subsystem.
Determines whether different storage subsystems are in the same zone. While no technical problem is associated with storage subsystems from different vendors and of different models being in the same zone, an administrator may find them more difficult to organize. When similar storage systems are in the same zone, an administrator can easily group them for different applications.

Each zone is part of a zone set.
Determines the presence of orphan zones. Orphan zones are not associated with any zone set. They are not useful because their definitions are not used and they take up switch resources.

Each host must be zoned so that it can access all of its assigned volumes.
Determines whether the zones that were configured by the storage administrator allow each computer to access all of the storage volumes that are assigned to it. The administrator specifies the storage subsystem ports through which the computer port accesses volumes, but may forget to configure zones that enable the ports to communicate during volume assignment. This policy also determines whether zoning makes assigned volumes inaccessible to the computer ports. The Fabric scope is not supported by this policy.

Each computer has only HBAs of the same model and firmware version.
Checks whether there is only one type of HBA in each computer. Using only one type of HBA minimizes configuration problems. The policy also checks whether firmware upgrades have been done properly for all HBAs in a computer. The Zone Set scope is not supported by this policy.

For each host type and operating system, every HBA of a given model must have the same firmware version.
Determines whether all firmware upgrades have been done for the HBAs in the operating system. The Zone Set scope is not supported by this policy.

Every SAN switch of a given model must have the same firmware version.
Determines whether firmware upgrades have been done for all switches of the same type. For example, if you have four identical models of SAN switches from the same vendor and you perform a firmware upgrade on one, it is best to perform the upgrade on all of the others. The Zone set scope is not supported by this policy.

Every storage subsystem of a given model must have the same firmware version.

Determines whether firmware upgrades have been done for all storage subsystems of the same type. For example, if you have four identical storage subsystems from the same vendor and you perform a firmware upgrade on one, it is best to perform the upgrade on all of the others. The Zone Set scope is not supported by this policy.

Each fabric may have a maximum of x zones.

Checks whether the number of zone definitions in the fabric is larger than the number that you entered. In large fabrics, too large a number of zone definitions can become a problem. Fabric zone definitions are controlled by one of the switches in that fabric, and limiting their number ensures that the switch's zoning tables do not run out of space. The Zone Set scope is not supported by this policy. You can enter up to a maximum of 9999 zones.

Each zone may have a maximum of x zone members.

Checks whether the number of zone members in a zone is larger than the number that you entered. In large fabrics, too large a number of zone members can become a problem. Fabric zone members are controlled by one of the switches in that fabric, and limiting their number ensures that the switch's zoning tables do not run out of space. You can enter up to a maximum of 9999 zones.

Replication Plan is intact with respect to the SRG(s) and the replication session associated during planning through SAN Planner.

Determines whether SRGs that have associated resiliency profiles are intact with respect to their members and their replication relationships.

Restriction: This policy is no longer valid and will be removed in a future release.

All the source volumes involved in Metro Mirror Failover/Failback sessions are conforming to 1:4 primary to secondary LSS for a failback direction scenario.

Metro Mirror (Failover/Failback) can be directional, and up to 4 LSS secondaries can be connected to a primary LSS.

Inter/intra site connectivity is valid for replication plan deployments.

Checks whether connectivity either direct or through fiber channel fabric exists for resiliency session deployments.

8. To specify how often to run the analysis, scroll to **Scheduling**.
9. In the **How often to run** area, schedule how often you want the job to run. Choose from the following actions:

Run Now

Run the job immediately.

Run Once at

Run the job once at a specified date and time. Use the list boxes to specify the month, day, and year. Type the hour and minutes over the existing defaults. Click AM or PM.

Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start

time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

Run Repeatedly
Run the job repeatedly according to a schedule that you specify:

- Indicate when to begin the job. Use the list boxes to specify the month, day, and year. Type the hour and minutes over the existing defaults. Click AM or PM.
- Indicate how often to run the job. Select one of the following:
  - In the **Repeat Every** field, click the desired number of minutes, hours, days, weeks, or months.
  - In the **Run on these days** field, click to display a check mark next to the desired days.

10. In the **How to handle time zones** area, indicate the time zone to be used when scheduling and running the policy. Click one of the following:

    Use the time zone that the server runs in
    Select this option to use the time zone of the location where the Data server resides.

    Use this time zone
    Select this option to indicate a specific time zone for when to run the job. From the list box, select the time zone.

11. From the menu bar, select **File** > **Save**. The Save As pop-up opens.
12. Type a name for the analyzer and click **OK**. The Notice pop-up displays **Analyzer submitted**.
13. Click **OK** to begin the job. To display the newly named analyzer job in the Navigation tree pane, expand **IBM Tivoli Storage Productivity Center** > **Analytics** > **Configuration Analysis**. The format of the name is *creator.analyzer*, where *creator* is the name that is displayed on the Create Analyzer window and *analyzer* is name that you assigned in step 12 (for example *db2admin.analyzer2*).
14. Expand the job to view information about the run. The information includes the number of times that the job was run, and the date and time of the run. The color of the type indicates the status:
    - Blue indicates that the run is in progress.
    - Green indicates that the run is finished.

    To refresh the status, right-click the job and click **Update Job Status**.
15. Click the run to view information about it. On a page titled with the name of the job and run number, the following information is displayed:

    Run Number of the run

    Status The status of the run (Success or Failed), not whether policy violations occurred

    Start Time
    The time at which the run started

    Finish Time
    The time at which the run ended

    # Jobs The number of jobs

    # failed
    The number of jobs that failed

Computer
The IP address of the computer on which the job was run

Status The status of the run

Start Time
The time at which the run started

Finish Time
The time at which the run ended

Log File Name
The complete path name of the log file

16. To view whether a policy violation occurred, click . The Job log file page displays the following information:
    - Scope of the run (all fabrics, one fabric, or one zone set)
    - Policies that were checked
    - Total violations
    - Breakdown of violations per policy
    - Whether the run completed successfully or the errors it encountered

    Note: To avoid overloading the Topology Viewer, the analysis job stores and displays the first 50 policy violations only. You must resolve these violations and run the job again to view any remaining violations, or you can change this policy violation limit by using the **setdscfg** command. The attribute to set is **MaxPolicyViolationsStored**.
17. To view one or more alerts generated by policy violations, in the Navigation tree expand **IBM Tivoli Storage Productivity Center** > **Alerting** > **Alert Log** > **Configuration Analysis**. The Alert History - All Policy Violations page displays a log of job runs that generated alerts for policy violations. A policy violation alert is generated for each policy that was violated during a run. A policy might be violated several times, but only one alert is generated. The text in the alert indicates the number of times the policy was violated. The Alert History - All Policy Violations page includes the following information:

Column	Description
	When clicked, this provides details about the alert.
	When clicked, this makes the Clear and Delete buttons active. When double-clicked, provides details about the alert.
Computer	Not applicable (N/A)
Configuration Analysis	The name of the configuration analysis job
Alert Type	The type of alert generated during the run
First Triggered	The date and time of when the first alert was generated
Last Triggered	The date and time of when the last alert was generated
# Times	The number of times that an alert was generated for a run
Alert Creator	The userid of the person who created the run
Alert Name	The official name of the alert

Column	Description
Script	Not applicable (N/A)

For more information about the alert log, see "Viewing alert logs" on page 356.

18. To view details about each alert, in the Alert History - All Policy Violations page click the run that you want, then click or double-click . The Detail for Alert page displays the following information:

Field	Description
Configuration Analysis	The name and scope of the analyzer job.
Run	The number of run in which the alert occurred.
Alert Type	The type of alert (Policy Violation).
State	The status of the alert. An *active* (unresolved) state displays a red exclamation mark ; an *acknowledged* (resolved) state displays a green check mark
Timestamp	The date and time that the alert was created. Expressed as the month, day, year, hour, minute, second, and AM or PM.
Alert Creator	The userid of the person who created the analyzer job.
Alert Name	The name of the alert.
Alert Text	A description of the policy that was violated and the number of times that it was violated during that analyzer run.

Resolving policy violations:
When you click the Alert tab in the topology viewer and expand the entries in the table view, information displays about each configuration analysis policy violation. To verify and fix the policy violation, use the information in the following table.

Configuration Analysis Policy	Enclosing Entity	Affected Entity	How to Verify and Fix the Policy Violation
Policy 1. Each connected computer and storage subsystem port must be in at least one zone in the specified zone sets.	Computer, Subsystem, or Other	Port	In the zone tab within the L2:Fabric view, verify that the port is not a member of the zone set. Fix the violation by adding the port to the appropriate zone within the zone set.
	Fabric	Zone set	

Configuration Analysis Policy	Enclosing Entity	Affected Entity	How to Verify and Fix the Policy Violation
Policy 2. Each HBA accesses storage subsystem ports or tape ports, but not both.	Computer	HBA	In the zone tab within the L2:Computer view, verify that HBA ports are accessing both the disk and tape subsystems. Fix the violation by configuring the computer access to disk and tape in two separate zone sets, and ensuring that only one of these is active at a given time.
Policy 3. Each volume is accessed only by computers running the same type and version of operating system.	Subsystem	Volume	In L1:Computers view, one can see that the computers accessing the volume do not have the same OS type and version. This violation can be fixed by mapping or assigning the volume only to the computers with same OS type and version.
	Computer	Port	
	Computer	Port	
Policy 4. Each zone contains only HBAs from a single vendor.	Fabric	Zone	In the zone tab within the L2:Fabric view, find the computers in the zone. Open the respective L2:Computer views and select the HBA tab in the table view to view the vendor of the HBAs on those computers. Fix the violation by editing the zone definition and ensuring that all the HBAs in the zone are from the same vendor.
Policy 5. Each zone contains only a single model of storage subsystem.	Fabric	Zone	In the L2:Fabric view, verify that the models of the subsystems in the zone are different. Fix the violation by editing the zone definition and ensuring that all the subsystems in the zone have the same model.
	Subsystem	Port	
Policy 6. Each zone is part of a zone set.	Fabric	Zone	In the zone tab within the L2:Fabric view, verify that the zone does not belong to any of the zone sets and is listed under Orphan Zones. Fix the violation by adding the zone to a zone set or by deleting the zone.
Policy 7. Each host must be zoned so that it can access all of its assigned volumes.	Computer or Other	Port	None of the zones in the L2:Computer or L2:Subsystem views has both the computer and subsystem port. Fix the violation by adding or editing the zone definition, and by ensuring that both the computer and subsystem port are in the same zone.
	Subsystem	Port	
	Subsystem	Volume	
Policy 8. Each computer has only HBAs of the same model and firmware version.	Computer	Computer	Open the L2:Computer view and select the HBA tab in the table view to verify that the model and firmware version for the HBAs on that computer are different. Fix the violation by installing the HBAs of the same model or upgrading the firmware version of the HBA on that computer.

Configuration Analysis Policy	Enclosing Entity	Affected Entity	How to Verify and Fix the Policy Violation
Policy 9. For each host type and operating system, every HBA of a given model must have the same firmware version.	Computer	HBA	Open the respective L2:Computer views and select the HBA tab in the table view to verify that the firmware and driver version for the HBAs on those computers are different. Fix the violation by upgrading the firmware and driver versions of the HBAs on those computers.
	Computer	HBA	
Policy 10. Every SAN switch of a given model must have the same firmware version.	Switch	Switch	Open the L1:Switches view and select the switch tab to verify that the firmware versions on those switches are different. Fix the violation by upgrading the firmware versions on those switches.
	Switch	Switch	
Policy 11. Every storage subsystem of a given model must have the same firmware version.	Subsystem	Subsystem	Open the L1:Subsystems view and select the subsystem tab to verify that the firmware versions on those subsystems are different. Fix the violation by upgrading the firmware versions on those subsystems.
	Subsystem	Subsystem	
Policy 12. Each fabric may have a maximum of *x* zones.	Fabric	Fabric	In the zone tab within the L2:Fabric view, verify that the number of zones present in the fabric is greater than *x* (the value entered by the user in this policy). Fix the violation by reducing the number of zones to a value less than *x*.
Policy 13. Each zone may have a maximum of *x* zone members.	Fabric	Zone	Search for the zone in the zone tab in the L2:Fabric view. Verify that the number of zone members present in that zone is greater than *x* (the value entered by the user in this policy). Fix the violation by reducing the number of zone members to a value less than *x*.
Policy 14. Replication Plan is intact with respect to the SRG(s) and the replication session associated during planning through SAN Planner. **Restriction:** This policy is no longer valid and will be removed in a future release.	Storage Resource Group (SRG)	Replication session	In the alert tab in the L0:Storage Resource Groups view, verify that the numbers of source and target volumes in the SRG are not equal or that the volumes are not in a replication relationship. Fix the violation by running the SAN Planner again and provisioning additional volumes or provisioning resiliency for the existing volumes.
	SRG	Volume	

Configuration Analysis Policy	Enclosing Entity	Affected Entity	How to Verify and Fix the Policy Violation
Policy 15. All the source volumes involved in Metro Mirror Failover/Failback sessions are conforming to 1:4 primary to secondary LSS for a failback direction scenario.	SRG	Logical Subsystem (LSS)	IBM Tivoli Storage Productivity Center does not have an LSS view, but in the L2:Storage view, you can verify the details of the LSSs responsible for the violation. This violation only occurs if the copy session is created incorrectly in IBM Tivoli Storage Productivity Center for Replication. Fix the violation by changing the session details in Tivoli Storage Productivity Center for Replication to conform with the policy.
Policy 16. Inter/intra site connectivity is valid for replication plan deployments.	SRG	Replication session	In the alert tab in the L0:Storage Resource Groups view, verify that there is no connectivity between source and target volumes. Fix the violation by manually providing a common fibre channel fabric or direct connection.

Enforce storage policies

Enforcing policies for storage usage can be critical to ensuring the appropriate and most cost-effective use of storage devices. Policies are enforced when conditions that you defined are detected by a monitoring job. Use quotas, constraints, and scheduled actions to manage and enforce your storage policies.

Use the **Policy Management** node in Data Manager and Data Manager for Databases to manage and enforce storage policies.

For example, to enforce a policy that limits the space a group of users can use on the network, use a Data Manager quota. That quota is triggered: 1) after a scan is run and 2) the scan detected that a group of users is using more storage space than was indicated in the quota.

The following table describes the jobs that you can use to enforce policies for storage usage in your environment:

Table 75. Descriptions of policy types and their triggering conditions

Policy Type / Manager	What can I do with this job?	What conditions can trigger an alert?
Quotas (Data Manager, Data Manager for Databases) • **Policy Management** > **Quotas** > **User** • **Policy Management** > **Quotas** > **OS User Group**	Use quotas to define limits on and monitor the amount of storage that a user or a group of users can use. You can set quotas to specify limits for storage usage at different levels. You can control precisely how much storage a user can use on specific storage entities: • Define quotas in Data Manager to set limits on the amount of storage that a user or a group of users can use at a network, file system, or computer level. • Define quotas in Data Manager for Databases to set limits on the amount of storage that a user or a group of users can use at a network, instance, database, or table space level. Monitoring job that detects conditions in quotas: scan	**User uses more than n**, where *n* represents a number that is measured in KBs, MBs, or GBs. **OS User Group uses more than n**, where *n* represents a number that is measured in KBs, MBs, or GBs.
Network Appliance Quotas (Data Manager) • **Policy Management** > **Network Appliance Quotas** > **Schedules** • **Policy Management** > **Network Appliance Quotas** > **Imported User Quotas** • **Policy Management** > **Network Appliance Quotas** > **Imported OS User Group Quotas**	Use Network Appliance Quotas to import, view, and report on the NetApp Quotas that are defined on NAS filers. Monitoring job that detects conditions in Network Appliance quotas: scan	**n% of hard limit**, where *n* represents a percentage of the hard limit that is defined in the NetApp Quota.
Constraints (Data Manager) **Policy Management** > **Constraints**	Use constraints: • To define the acceptable and unacceptable file types, file sizes, and file owners for computers in your environment. For example, you can use constraints to alert you when users store certain files (such as MP3 or .avi files) on a monitored computer. • To request an IBM Tivoli Storage Productivity Center archive and backup of the largest violating files that are identified by a constraint. IBM Tivoli Storage Manager protects your organization 's data from hardware failures and other errors by storing backup and archive copies of data on auxiliary storage. Monitoring job that detects conditions in constraints: scan	**Violating files use more than n**, where *n* represents a number that is measured in KBs, MBs, or GBs.

Table 75. Descriptions of policy types and their triggering conditions (continued)

Policy Type / Manager	What can I do with this job?	What conditions can trigger an alert?
Scheduled Actions (Data Manager, Data Manager for Databases) **Policy Management > Scheduled Actions > Scripts**	Use scheduled actions to schedule scripts to run against computers and computer groups. Scheduled actions are script-based. You can use any third-party tools for actions, such as recovery or provisioning. Monitoring job that detects conditions in scheduled actions: n/a	Not applicable. Through the Scheduled Action facility, you can run scripts according to the schedule that you define. You do not have to associate the script with an alert.
Archive/Backup (Data Manager) **Policy Management > Archive/Backup**	View and edit archive and backup jobs that were created based on files that are selected from reports in the Reporting facility. Monitoring job that detects conditions in archive/backups: scan	Not applicable. Any archive/backup jobs that are displayed under this node were created based on files that are shown in the following reports: • Usage > Files > Largest Files • Usage > Files > Most Obsolete Files • Usage > Files > Orphan Files • Usage > Files > Duplicate Files • Backup > Most At Risk Files • Usage Violations > Constraint Violations

Setting storage usage quotas

Use quotas to set limits on the amount of storage space that a user or group of users can consume. You can use Data Manager and Data Manager for Databases to set quotas at the file system, computer, network, database, table space, and instance levels.

Use quotas to set limits on the amount of storage that a user or a group of users can consume.

The following steps provide an overview of how quotas work:

1. A scan gathers extensive information about the storage resources within an enterprise. This information is stored in the database repository.
2. When a quota job runs, it checks the information in the repository to determine whether any quota violations exist.
3. If there are any violations, an alert is triggered and the appropriate action is taken.

Note: Because the information that is analyzed by quotas is provided by scans, schedule the quota to run after the scan to collect information. This ensures that quotas are using the latest information about an enterprise to determine violations.

You can define two different kinds of quotas: User and OS User Group.

User Quota
Select users and groups of users for a quota that are defined in Data Manager or Data Manager for Databases

OS User Group Quota
Select OS User Groups for a quota. OS User Groups are groups of users

that are defined in the operating system outside of Data Manager. You can create Groups that contain many OS User Groups. When you create a quota of this type, you can select individual OS User Groups and groups that consist of multiple OS User Groups.

By providing you with the ability to define quotas at different levels, you can control precisely how much storage a user or user group can consume on specific file systems, computers, network, databases, table spaces, and instances:

File system
Set storage usage limits on specific file systems and groups of file systems that are monitored by Data Manager. You are alerted if a user or group of users exceeds the storage usage limit that is defined for the specific file system or groups of file systems in the quota.

Computer
Set storage usage limits on specific computers and groups of computers that are monitored by Data Manager. You are alerted if a user or group of users exceeds the storage usage limit that is defined for the specific computers or groups of computers in the quota.

Database-Tablespace
Set storage usage limits on specific databases, table spaces, groups of databases, and groups of table spaces that are monitored by Data Manager for Databases. You are alerted if a user or group of users exceeds the storage usage limit that is defined for those storage resources in the quota.

Instance
Set storage usage limits on specific instances, computers, and groups of computers that are monitored by Data Manager for Databases. You are alerted if a user or group of users exceeds the storage usage limit that you defined for the specific instances, computers, and groups of computers in the quota.

Network
Set network-level storage usage limits for multiple computers, groups of computers, file systems, databases, and table spaces within your network that are monitored by Data Manager and Data Manager for Databases.

Creating a network-wide quota:

Learn how to create a quota that generates an alert when a user or group of users uses more than a specified amount of storage space across the file systems and computers on a network.

To create a network-wide quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management** > **Quotas**.
2. Expand the **User** or **OS User Group** node dependent upon which entity you want to apply a quota.
3. Select the level at which you want to apply the quota by right-clicking **Network**.
4. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
5. In the **Description** field, type a brief description of the quota.
6. Select a profile from the **Profile** selection list to control which storage usage statistics are used when quota violations are checked. You can select any of the GROUP_BY_OWNER type profiles shown. Make sure that the selected

profile is used in scans of all the file systems that are chosen on the **File systems** tab. This action ensures that the quota is fairly enforced for all users.

7. Select users or user groups against which you want to apply the quota:
   a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
   b. Click . The users or groups are moved to the **Current Selections** field.
8. Click the **File systems** tab; the **File systems** page opens.
9. Select file systems, file system groups, computers, and computers groups on which you want to limit the storage usage of users and user groups.
   a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups that are displayed in this field were created by using **Data Manager** > **Monitoring** > **Groups**.
   b. Click . The resources are moved to the **Current Selections** field.
10. Click the **When to Run** tab; the **When to Run** page opens.
11. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota one time at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This action ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

    Run Repeatedly
    Check the quota repeatedly according to a schedule you specify. You can:
    - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job is to begin running.
    - Indicate how often the job is to run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
    - Indicate which days on which run the job. A check mark is displayed next to the days you select.
12. In the **How to Handle Time Zones** area, indicate the time zone to be used when you schedule and run the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields to run it at these times.
    a. 9:00 AM of the time zone where the Data server is located
    b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job runs at 7:00 AM on an agent that is in the PST time zone.

    Use the timezone that the server runs in
    Select this option to use the time zone where the Data server is located.

Use this timezone
Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.

13. Click the **Alert** tab; the **Alert** page opens.
14. In the **Triggering Condition** area, specify the event that triggers the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
    c. In the **Value Units** field, select the unit of measurement. kilobytes, megabytes, or gigabytes.
15. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

SNMP trap
An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

TEC / OMNIbus Event
A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the `tivoliSRM.baroc` file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, or running programs.

Login notification
The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who receives alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who receives alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that is recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog is sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters that are listed in the **Specify Script** window are passed to the script when the alert is triggered.

Email An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit email** to change the actual email message.

16. Save the quota:
    a. Click **File** > **Save**. The Save As window opens.
    b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the navigation tree pane

Creating a network-wide quota for databases:

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of RDBMS storage space across the databases, tablespaces, groups of database-tablespaces, computers, and groups of computers on a network.

To create a network-wide quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager for Databases** > **Policy Management** > **Quotas**.
2. Select the level at which you want to apply the quota by right-clicking **Network**.
3. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
4. In the **Description** field, type a brief description of the quota.
5. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP_BY_OWNER type profiles shown. Make sure that the selected profile is used in scans of all the databases and tablespaces chosen on the **Databases-Tablespaces** tab. This ensures that the quota is fairly enforced for all users.
6. Select users or user groups against which you want to apply the quota:
    a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
    b. Click . The users or groups are moved to the **Current Selections** field.
7. Click the **Databases-Tablespaces** tab; the Databases-Tablespaces page opens.
8. Select the databases, tablespaces, database-tablespaces groups, computers, and computer groups on which you want to limit the storage usage of users and user groups.
    a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager for Databases** > **Monitoring** > **Groups**.
    b. Click . The resources are moved to the **Current Selections** field.
9. Click the **When to Run** tab; the When to Run page opens.
10. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

Run Once at
Check the quota one time at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This action ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

Run Repeatedly
Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job is to begin running.
- Indicate how often the job is to run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days you select.

11. In the **How to Handle Time Zones** area, indicate the time zone to be used when you schedule and run the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields to run it at these times.
    a. 9:00 AM of the time zone where the Data server is located
    b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job runs at 7:00 AM on an agent that is in the PST time zone.

    Use the timezone that the server runs in
    Select this option to use the time zone where the Data server is located.

    Use this timezone
    Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.
12. Click the **Alert** tab; the Alert page opens.
13. In the **Triggering Condition** area, specify the event that triggers the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than** *n*, where *n* represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
    c. In the **Value Units** field, select the unit of measurement. kilobytes, megabytes, or gigabytes.
14. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

TEC / OMNIbus Event
A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the `tivoliSRM.baroc` file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, or running programs.

Login notification
The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who receives alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who receives alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that is recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog is sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters that are listed in the **Specify Script** window are passed to the script when the alert is triggered.

Email An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit email** to change the actual email message.

15. Save the quota:
    a. Click **File** > **Save**. The Save As window opens.
    b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the navigation tree pane

Creating a computer quota:

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of storage space across computers and groups of computers.

To create a computer quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management** > **Quotas**.
2. In the **Description** field, type a brief description of the job.
3. Expand the **User** or **OS User Group** node depending upon which entity you want to apply a quota.

4. Select the level at which you want to apply the quota by right-clicking **Computer**.
5. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
6. In the **Description** field, type a brief description of the quota.
7. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP_BY_OWNER type profiles shown.
8. Select users or user groups against which you want to apply the quota:
   a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
   b. Click . The users or groups are moved to the **Current Selections** field.
9. Click the **Computers** tab; the Computers page opens.
10. Select computers and computers groups on which you want to limit the storage usage of users and user groups.
    a. In the **Available** field, click the computers and computers groups to which you want to apply the quota. The groups displayed in this field were created using **Data Manager** > **Monitoring** > **Groups**.
    b. Click . The resources are moved to the **Current Selections** field.
11. Click the **When to Run** tab; the When to Run page opens.
12. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota one time at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This action ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

    Run Repeatedly
    Check the quota repeatedly according to a schedule you specify. You can:
    - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job is to begin running.
    - Indicate how often the job is to run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
    - Indicate which days on which run the job. A check mark is displayed next to the days you select.
13. In the **How to Handle Time Zones** area, indicate the time zone to be used when you schedule and run the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields to run it at these times.
    a. 9:00 AM of the time zone where the Data server is located

b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job runs at 7:00 AM on an agent that is in the PST time zone.

Use the timezone that the server runs in
Select this option to use the time zone where the Data server is located.

Use this timezone
Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.

14. Click the **Alert** tab; the Alert page opens.
15. In the **Triggering Condition** area, specify the event that triggers the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
    c. In the **Value Units** field, select the unit of measurement. kilobytes, megabytes, or gigabytes.
16. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

SNMP trap
An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

TEC / OMNIbus Event
A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the `tivoliSRM.baroc` file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, or running programs.

Login notification
The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who receives alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who receives alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that is recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog is sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters that are listed in the **Specify Script** window are passed to the script when the alert is triggered.

Email An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit email** to change the actual email message.

17. Save the quota:
    a. Click **File** > **Save**. The Save As window opens.
    b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the navigation tree pane

Creating a file system quota:

Learn how to create a quota that generates an alert when a user or group of users uses more than a specified amount of storage space across file systems, groups of file systems, computers, and groups of computers.

To create a file system quota, complete the following steps.

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management** > **Quotas**.
2. Expand the **User** or **OS User Group** node dependent upon which entity you want to apply a quota.
3. Select the level at which you want to apply the quota by right-clicking **File system**.
4. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
5. In the **Description** field, type a brief description of the job.
6. Select a profile from the **Profile** selection list to control which storage usage statistics are used when quota violations are checked. You can select any of the GROUP_BY_OWNER type profiles shown. Make sure that the selected profile is used in scans of all the file systems that are chosen on the **File systems** tab. This action ensures that the quota is fairly enforced for all users.
7. Select users or user groups against which you want to apply the quota.
    a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
    b. Click . The users or groups are moved to the **Current Selections** field.
8. Click the **File systems** tab to open the page.
9. Select file systems, file system groups, computers, and computer groups on which you want to limit the storage usage of users and user groups.
    a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups that are displayed in this field were created by using **Data Manager** > **Monitoring** > **Groups**.
    b. Click . The resources are moved to the **Current Selections** field.
10. Click the **When to Run** tab to open the page.

11. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota one time at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This action ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

    Run Repeatedly
    Check the quota repeatedly according to a schedule you specify. You can:
    - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job is to begin running.
    - Indicate how often the job is to run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
    - Indicate which days on which run the job. A check mark is displayed next to the days you select.

12. In the **How to Handle Time Zones** area, indicate the time zone to be used when you schedule and run the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields to run it at these times.
    a. 9:00 AM of the time zone where the Data server is located
    b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job runs at 7:00 AM on an agent that is in the PST time zone.

    Use the timezone that the server runs in
    Select this option to use the time zone where the Data server is located.

    Use this timezone
    Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.

13. Click the **Alert** tab to open the page.
14. In the **Triggering Condition** area, specify the event that triggers the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than** *n*, where *n* represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
    c. In the **Value Units** field, select the unit of measurement. kilobytes, megabytes, or gigabytes.
15. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

SNMP trap
An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

TEC / OMNIbus Event
A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the `tivoliSRM.baroc` file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, or running programs.

Login notification
The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who receives alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who receives alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that is recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog is sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters that are listed in the **Specify Script** window are passed to the script when the alert is triggered.

Email An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit email** to change the actual email message.

16. Save the quota:
    a. Click **File** > **Save**. The Save As window opens.
    b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the navigation tree pane

Creating an instance quota:

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of RDBMS storage space across instances, computers, and groups of computers.

To create an instance quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager for Databases** > **Policy Management** > **Quotas**.

2. Select the level at which you want to apply the quota by right-clicking **Instance**.
3. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
4. In the **Description** field, type a brief description of the quota.
5. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP_BY_OWNER type profiles shown.
6. Select users or user groups against which you want to apply the quota:
   a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
   b. Click

. The users or groups are moved to the **Current Selections** field.
7. Click the **Instances** tab; the Instances page opens.
8. Select the instances, computers, and computer groups on which you want to limit the storage usage of users and user groups.
   a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager for Databases** > **Monitoring** > **Groups**.
   b. Click

. The resources are moved to the **Current Selections** field.
9. Click the **When to Run** tab; the When to Run page opens.
10. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota one time at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This action ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

    Run Repeatedly
    Check the quota repeatedly according to a schedule you specify. You can:
    - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job is to begin running.
    - Indicate how often the job is to run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
    - Indicate which days on which run the job. A check mark is displayed next to the days you select.
11. In the **How to Handle Time Zones** area, indicate the time zone to be used when you schedule and run the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields to run it at these times.
    a. 9:00 AM of the time zone where the Data server is located

b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job runs at 7:00 AM on an agent that is in the PST time zone.

Use the timezone that the server runs in
Select this option to use the time zone where the Data server is located.

Use this timezone
Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.

12. Click the **Alert** tab; the Alert page opens.
13. In the **Triggering Condition** area, specify the event that triggers the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
    c. In the **Value Units** field, select the unit of measurement. kilobytes, megabytes, or gigabytes.
14. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

SNMP trap
An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

TEC / OMNIbus Event
A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the `tivoliSRM.baroc` file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, or running programs.

Login notification
The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who receives alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who receives alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that is recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog is sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters that are listed in the **Specify Script** window are passed to the script when the alert is triggered.

Email An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit email** to change the actual email message.

15. Save the quota:
    a. Click **File** > **Save**. The Save As window opens.
    b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the navigation tree pane

Creating a database-tablespace quota:

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of RDBMS storage space across databases, tablespaces, groups of databases, and groups of tablespaces.

To create a databases-tablespaces quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager for Databases** > **Policy Management** > **Quotas**.
2. Select the level at which you want to apply the quota by right-clicking **Databases-Tablespaces**.
3. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
4. In the **Description** field, type a brief description of the quota.
5. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP_BY_OWNER type profiles shown. Make sure that the selected profile is used in scans of all the databases and tablespaces chosen on the **Databases-Tablespaces** tab. This ensures that the quota is fairly enforced for all users.
6. Select users or user groups against which you want to apply the quota:
    a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
    b. Click . The users or groups are moved to the **Current Selections** field.
7. Click the **Databases-Tablespaces** tab; the Databases-Tablespaces page opens.
8. Select the databases, tablespaces, database-tablespaces groups, computers, and computer groups on which you want to limit the storage usage of users and user groups.
    a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager for Databases** > **Monitoring** > **Groups**.
    b. Click . The resources are moved to the **Current Selections** field.
9. Click the **When to Run** tab; the When to Run page opens.

10. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota one time at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This action ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

    Run Repeatedly
    Check the quota repeatedly according to a schedule you specify. You can:

    - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job is to begin running.
    - Indicate how often the job is to run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
    - Indicate which days on which run the job. A check mark is displayed next to the days you select.

11. In the **How to Handle Time Zones** area, indicate the time zone to be used when you schedule and run the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields to run it at these times.
    a. 9:00 AM of the time zone where the Data server is located
    b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job runs at 7:00 AM on an agent that is in the PST time zone.

    Use the timezone that the server runs in
    Select this option to use the time zone where the Data server is located.

    Use this timezone
    Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.

12. Click the **Alert** tab; the Alert page opens.
13. In the **Triggering Condition** area, specify the event that triggers the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
    c. In the **Value Units** field, select the unit of measurement. kilobytes, megabytes, or gigabytes.
14. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

SNMP trap
An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

TEC / OMNIbus Event
A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the `tivoliSRM.baroc` file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, or running programs.

Login notification
The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who receives alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who receives alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that is recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog is sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters that are listed in the **Specify Script** window are passed to the script when the alert is triggered.

Email An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit email** to change the actual email message.

15. Save the quota:
    a. Click **File** > **Save**. The Save As window opens.
    b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the navigation tree pane

Editing a storage usage quota:

Learn how to change the settings for a quota, including the users or user groups to which the quota applies, the resources on which the quota is set, and the amount that is set as the limit for storage usage.

To edit a quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management** > **Quotas**.

2. Expand the **User** or **OS User Group** node that is dependent upon the type of quota you want to edit.
3. Expand the level of the quota you want to edit. For example, if you want to edit a computer quota, expand the **Computer** node.
4. Select the quota that you want to edit. Information about the quota you want to edit is displayed in the right-topic pane.
5. In the **Description** field, type a brief description of the quota.
6. Select a profile from the **Profile** selection list to control which storage usage statistics are used when quota violations are checked. You can select any of the GROUP_BY_OWNER type profiles shown.
7. Select more users or user groups against which you want to apply the quota:
   a. In the **Available** field, click the users or groups of users against which you want to apply the quota. The groups displayed in this field were created using **Data Manager** > **Monitoring** > **Groups**
   b. Click >>. The users or groups are moved to the **Current Selections** field.
8. Remove users or user groups previously included in the quota:
   a. In the **Current Selections** field, click the users or groups of users that you want to remove from the quota.
   b. Click <<. The users or groups are moved back to the **Available** field.
9. Click the **Computers** or **File systems** tab dependent on the quota you are editing; the Computers or File systems page opens.
10. Select more file systems, file system groups, computers, or computers groups on which you want to limit the storage usage of users and user groups.
    a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups that are displayed in this field were created by using **Data Manager** > **Monitoring** > **Groups**.
    b. Click >>. The resources are moved to the **Current Selections** field.
11. Remove any file systems, file system groups, computers, or computers groups that you do not want included in the quota.
    a. In the **Current Selections** field, click the storage resources to remove.
    b. Click <<. The resources are moved to the **Current Selections** field.
12. Click the **When to Run** tab to open the page.
13. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota one time at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This action ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

Run Repeatedly

Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job is to begin running.
- Indicate how often the job is to run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days you select.

14. In the **How to Handle Time Zones** area, indicate the time zone to be used when you schedule and run the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields to run it at these times.
    a. 9:00 AM of the time zone where the Data server is located
    b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job runs at 7:00 AM on an agent that is in the PST time zone.

    Use the timezone that the server runs in

    Select this option to use the time zone where the Data server is located.

    Use this timezone

    Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.
15. Click the **Alert** tab to open the page.
16. In the **Triggering Condition** area, specify the event that triggers the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
    c. In the **Value Units** field, select the unit of measurement. kilobytes, megabytes, or gigabytes.
17. Click **File** > **Save** to save your changes.

Editing a storage usage quota for databases:

Learn how to change the settings for a quota, including the users or user groups to which the quota applies, the resources on which the quota is set, and the amount set as the limit for storage usage.

To edit a quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager for Databases** > **Policy Management** > **Quotas**.
2. Expand the level of the quota you want to edit. For example, if you want to edit a instance quota, expand the **Instance** node.
3. Click the quota you want to edit. Information about the quota you want to edit appears in the right-hand content pane.
4. In the **Description** field, type a brief description of the quota.

5. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP_BY_OWNER type profiles shown.
6. Select additional users or user groups against which you want to apply the quota:
   a. In the **Available** field, click the users or groups of users against which you want to apply the quota. The groups displayed in this field were created using **Data Manager** > **Monitoring** > **Groups**
   b. Click . The users or groups are moved to the **Current Selections** field.
7. Remove users or user groups previously included in the quota:
   a. In the **Current Selections** field, click the users or groups of users that you want to remove from the quota.
   b. Click . The users or groups are moved back to the **Available** field.
8. Click the **Databases-Tablespaces** or **Instances** tab depending on quota you are editing; the Databases-Tablespaces or Instances page opens.
9. Select additional databases, tablespaces, instances, computers, or computers groups on which you want to limit the storage usage of users and user groups.
   a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager for Databases** > **Monitoring** > **Groups**.
   b. Click . The resources are moved to the **Current Selections** field.
10. Remove any databases, tablespaces, instances, computers, or computers groups that you do not want included in the quota.
    a. In the **Current Selections** field, click the storage resources to remove.
    b. Click . The resources are moved to the **Current Selections** field.
11. Click the **When to Run** tab; the When to Run page opens.
12. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota one time at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This action ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

    Run Repeatedly
    Check the quota repeatedly according to a schedule you specify. You can:

    - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job is to begin running.

- Indicate how often the job is to run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days you select.

13. In the **How to Handle Time Zones** area, indicate the time zone to be used when you schedule and run the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields to run it at these times.
    a. 9:00 AM of the time zone where the Data server is located
    b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job runs at 7:00 AM on an agent that is in the PST time zone.

    Use the timezone that the server runs in
    Select this option to use the time zone where the Data server is located.

    Use this timezone
    Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.
14. Click the **Alert** tab; the Alert page opens.
15. In the **Triggering Condition** area, specify the event that triggers the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than** *n*, where *n* represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
    c. In the **Value Units** field, select the unit of measurement. kilobytes, megabytes, or gigabytes.
16. Click **File** > **Save** to save your changes.

Storage usage for Network Appliance quotas

Use Network Appliance quotas to import quotas from Network Appliance file servrs and determine how and when you will be alerted to the hard limits defined in those quotas.

Restriction: You can only import quotas if the NetApp device is configured as a file server. It is not available if the NetApp device was added only as a CIMOM.

You can use the NAS support within IBM Tivoli Storage Productivity Center to work with quotas defined using the software provided with Network Appliance (NetApp) File Servers. Using the **Network Appliance** node under **Policy Management**, you can do the following:

- Use the **Schedules** node to define a Data Manager quota that will import NetApp quotas and alert you to when those quotas are close to being exceeded.
- Use the **Imported User Quotas** node to view the definitions of User quotas stored on the NetApp file servers against which you ran a Data Manager quota (defined in the **Schedules** node).
- Use the **Imported OS User Group Quotas** node to view the definitions of OS User Group quotas stored on the NetApp file servers against which you ran a Data Manager quota (defined in the Schedules node).

For example, you can define to be alerted when:

- A user or user group is close to reaching the hard limit defined within the corresponding NetApp Quota
- The space consumed on a QTree is approaching the hard limit defined within the corresponding NetApp Quota

Specifically, the Data Manager NAS component enables you to work with the following types of Network Appliance quotas:

- Quotas that limit the space that QTree directories can consume on a volume
- Quotas that limit the space consumed by specific users on a volume
- Quotas that limit the space consumed by specific user groups on a volume

Quick Steps for Working with Network Appliance quotas

The following describes the general steps you must perform to work with Network Appliance Quotas:

1. Run a Data Manager scan against the Network Appliance file servers whose quotas you want to check. This ensures that Data Manager is using the latest statistics when determining quota violations.
2. Expand **Policy Management** > **Network Appliance Quotas**.
3. Right click the **Schedules** node and select **Create NetApp Quota Job** to define a quota job. The NetApp Quota window appears.
4. Use the tabs in the NetApp Quota window to specify:
   - the Network Appliance file servers whose quotas you want to import
   - when and how often to run the quota
   - the condition (% of hard limit) that will trigger an alert for the Network Appliance quotas and how you will be notified of that alert. For example, if you enter 90% for the alert condition, Data Manager will notify you whenever any of the quotas you are importing are within 10% of being violated.
5. Run the quota job from the **Schedules** node to import the Network Appliance quota definitions from the target NAS filer. When quotas are imported, the following occurs:
   - any User or OS User Group quotas violations are shown in the Quota Violation reports
   - any directory quota (Qtree) violations and definitions are displayed in **Alert Log** > **> Directory**.

Creating a Network Appliance storage usage quota:

Learn how to create a Network Appliance storage usage quota that imports Network Appliance quotas from filers and alert you to violations of those quotas.

Complete the following steps to:

- define the NetApp filers whose quotas you want to import
- determine when and how often to run the quota that checks for user-defined violations (% of hard limit) of NetApp Quotas
- specify the condition (% of hard limit) that will trigger an alert violation for a quota and how you will be notified of that alert

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management** > **Network Appliance Quotas**.

2. Right click the **Schedules** node.
3. Select **New NetApp Quota Job** from the pop-up menu. The content pane displays the What to PROBE page.
4. In the **Description** field, type a brief description of the quota.
5. Select the NetApp filers against which you want to run the quota job:
   a. In the **Available** field, click the NetApp filers or NetApp filer groups against which you want to apply the quota.
   b. Click . The NetApp filers or NetApp filer groups are moved to the **Current Selections** field.
6. Click the **When to Run** tab; the When to Run page opens.
7. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

   Run Now
   Check the quota immediately.

   Run Once at
   Check the quota one time at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

   Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This action ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

   Run Repeatedly
   Check the quota repeatedly according to a schedule you specify. You can:
   - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job is to begin running.
   - Indicate how often the job is to run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
   - Indicate which days on which run the job. A check mark is displayed next to the days you select.
8. In the **How to Handle Time Zones** area, indicate the time zone to be used when you schedule and run the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields to run it at these times.
   a. 9:00 AM of the time zone where the Data server is located
   b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job runs at 7:00 AM on an agent that is in the PST time zone.

   Use the timezone that the server runs in
   Select this option to use the time zone where the Data server is located.

   Use this timezone
   Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.
9. Click the **Alert** tab; the Alert page opens.

10. In the **Triggering Condition** area, specify the event that triggers the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, enter the value that represents the percentage of a NetApp's hard limit at which the alert is triggered. The hard limit represents the most amount of space a user or directory (Qtree) can consume on a NAS filer. For example, the hard limit that is defined within a NetApp Quota for a user group is 150 MB, and you enter `90` in the **Value** field. The Data Manager quota reports a violation when that user group consumes 135 or more megabytes.
    c. In the **Value Units** field, the only value unit available is the % of Hard Limit.
11. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

    TEC / OMNIbus Event
    A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the `tivoliSRM.baroc` file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, or running programs.

    Login notification
    The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

    Login ID
    The ID of the user who receives alert notification upon logging in to the system.

    Windows Event Log, UNIX Syslog
    The ID of the user who receives alert notification upon logging in to the system.

    Event Type (Windows Event Log only)
    The type of event that is recorded to the OS log.

    Facility (UNIX Syslog only)
    The location where the UNIX Syslog is sent. You can select User or Local.

    Level (UNIX Syslog only)
    The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

    Run script
    A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters that are listed in the **Specify Script** window are passed to the script when the alert is triggered.

Email An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit email** to change the actual email message.

12. Save the quota:
    a. Click **File** > **Save**. The Save As window opens.
    b. In the **Specify NetApp Quota job name** field, type a name for the quota and click **OK**. The quota is displayed in the Navigation Tree pane

Viewing user quotas imported from Network Appliance filers:

Learn how to view the definitions of user quotas stored on the NetApp filers against which you ran a Data Manager quota.

Before you can import Network Appliance user quotas and view them through **Data Manager** > **Policy Management** > **Network Appliance Quotas** > **Imported User Quotas**, you must:

1. Run a scan against the Network Appliance filers whose user quotas you want to import.
2. Use **Policy Management** > **Network Appliance Quotas** > **Schedules** to define and run a quota checking job against the NetApp filers whose user quotas you want to import.

Complete the following steps to view imported Network Appliance user quotas:

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management** > **Network Appliance Quotas** > **Imported User Quotas**.
2. Click the NetApp Quota whose definition you want to view. The content pane displays information about the NetApp user quota.
3. In the Users page of the Edit *quota_name* window, view the following information about imported Network Appliance user quotas.

   Hard Limit
   The hard limit for used space that is defined within the NetApp quota.

   Current Selections
   The users against which the NetApp quota applies.
4. In the File Systems page of the Edit *quota_name* window, view the file systems to which the NetApp Quota applies.

Viewing OS user quotas imported from Network Appliance filers:

Learn how to view the definitions of OS User Quotas stored on the NetApp filers against which you ran a Data Manager quota.

Before you can import Network Appliance OS user quotas and view them through **Data Manager** > **Policy Management** > **Network Appliance Quotas** > **Imported OS User Group Quotas**, you must:

1. Run a scan against the Network Appliance filers whose OS user quotas you want to import.
2. Use **Policy Management** > **Network Appliance Quotas** > **Schedules** to define and run a quota checking job against the NetApp filers whose OS user quotas you want to import.

Complete the following steps to view imported Network Appliance OS user quotas:

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management** > **Network Appliance Quotas** > **Imported OS User Quotas**.
2. Click the NetApp Quota whose definition you want to view. The content pane displays information about the NetApp OS user quota.
3. In the OS Users page of the Edit *quota_name* window, view the following information about imported Network Appliance OS user quotas.

   Hard Limit
   The hard limit for used space that is defined within the NetApp quota.

   Current Selections
   The OS users against which the NetApp quota applies.
4. In the File Systems page of the Edit *quota_name* window, view the file systems to which the NetApp Quota applies.

Viewing directory (QTree) quotas and violations:

Learn how to view the definitions of user quotas stored on the Network Appliance filers against which you ran a Data Manager quota.

Before you can view QTree quotas and violations, perform the following steps:
1. Use **Policy Management** > **Network Appliance Quotas** > **Schedules** to define and run a quota checking job against the Network Appliance filers containing the QTree quotas you want to import.
2. Create and run a scan job that includes the following objects:
   - the NetApp filers containing the QTrees whose quotas you want to import
   - the directory group named TPCUser.NetApp QTree Group
3. Use **Policy Management** > **Network Appliance Quotas** > **Schedules** to rerun the quota checking job you defined in step 1.

Complete the following steps to view a QTree quota violation and definition:
1. In the Navigation tree pane, expand **Alerting** > **Alert Log**.
2. Click the **Directory** node. The Alert History — All Directories window appears. This window lists all the alerts triggered by directory and QTree quota violations.
3. To view details about a violation you can do either of the following:
   - Double click the entry representing the QTree violation to view details about that violation and the QTree quota definition.
   - Right click the entry representing the QTree violation and select **Edit Definition** from the pop-up menu to view the QTree quota definition.

     Note: You cannot edit a QTree quota definition using Data Manager. To edit a QTree quota definition you must use the software provided with your Network Appliance filer.

Setting file constraints

Use constraints to define the acceptable and unacceptable file types, file sizes, and file owners for a NAS filer or a set of NAS filers in your environment. You can also use constraints to request an IBM Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint.

Creating a file constraint:

Learn how to create a constraint that defines the acceptable and unacceptable file types, file sizes, and file owners for a computer or a set of computers in your environment.

To create a constraint, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management**.
2. Right-click **Constraints**, and click **Create Constraint**. The Create Constraint window opens.
3. In the **Description** field, type a brief description of the job.
4. Select file systems, file system groups, computers, and computer groups against which to apply the constraint.
   a. In the **Available** field, click the storage resources against which you want to apply the constraint. Select individual file systems and computers, and file system and computer groups that you previously defined by using the group facility. Cluster resource groups in a PowerHA SystemMirror for AIX or MSCS clustered environment are listed as computers.
   b. Click . The storage resources are moved to the **Current Selections** field.
5. Click the **File Types** tab; the File Types page opens. Select the file types that you want to allow or forbid.
6. Click one of the following selections.

   Forbid file matching these patterns
   Do NOT allow files of the specified type on the storage resources that are selected on the File Systems page.

   Allow ONLY file matching these patterns
   Allow files ONLY of the specified type on the storage resources that are selected on the File Systems page.
7. Click the file types that you would like to allow or forbid in the **Choose a file type** list box. The file types that are listed in this box are predefined in the system.
   a. In the **Choose a file type:** field, click the file types that you would like to allow or forbid.
   b. Click . The file types are moved to the **Forbidden files** field.
8. Optional: Create a custom file type:
   a. Enter the pattern in the **Or enter a pattern** field.
   b. Click . The file type is moved to the **Forbidden files** field.
9. Click the **Users** tab; the Users page opens. You can allow or forbid files that are owned by specific users
10. Click one of the following selections.

    Forbid file owned by selected users
    Do NOT allow files that are owned by the specified users.

    Allow ONLY file owned by the specified owner
    Allow files that are owned by the specified users only.
11. Select the users whose files you want to allow or forbid:
    a. In the **Available Users** field, click the users that you want to include.

b. Click >>. The users are moved to the **Forbidden Users** field.

12. Click the **Options** tab; the Options page opens. You can define alerts for files that are based on characteristics other than file type and file owner. You can also view a where clause that is used as a filter for locating files that violate the constraint definition. The following information is provided.

    Violating File Limits
    Indicate the maximum number of violating filenames that should be kept per agent.

    number of days since last access greater than
    Click this box and enter a length of time (in days) in which files have not been accessed. When a constraint job locates a file that has not been accessed in the length of time you indicated in this field, an alert is triggered.

    bigger than ##
    Check this box and enter a file size that should trigger an alert. When a constraint job locates a file of the size (or greater) that you indicated in this field, an alert is triggered. Select the unit of measurement for the value you entered in the list box. The valid choices are: bytes, **KB** (kilobytes)**MB** (megabytes), **GB** (gigabytes), and **TB** (terabytes).

    set-UID root
    Check this box trigger alerts on files that have the SET-UID root bit on (because this can be a security consideration).

    not backed up (Windows only) and number of days since last modification greater than
    Click this box and enter the number of days since a file was last modified and not backed up. When the constraint job locates a file that was modified and not backed up in the amount of days that you indicate in this field, an alert is triggered.

    File Filter text
    This section displays a *where* clause that will be used as a filter for locating files that violate the Constraint definition. The choices you make on this tab and the choices you made on the Users and File Types pages are reflected in the content of this *where* clause. To edit the *where* clause filter text, click **Edit Filter**. This will allow you to create new where conditions and groupings using the where clause interface.

13. Click the **Alert** tab; the Alert page opens. You can define the alerting criteria for a constraint.

14. In the **Triggering Condition** area, specify the event that triggers the constraint alert:

    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.

    b. In the **Value** field, type a numeric value that represents the amount of space a file must consume to violate the constraint.

    c. In the **Value Units** field, select the unit of measurement. kilobytes, megabytes, or gigabytes.

    Once Data Manager finds enough files that violate the condition (that exceed the total size amount you specified), the constraint is considered violated and

the alert is triggered. When the alert is triggered, the actions that you specify in the **Triggered Actions** section are taken.

15. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

    TEC / OMNIbus Event
    A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the `tivoliSRM.baroc` file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, or running programs.

    Login notification
    The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

    Login ID
    The ID of the user who receives alert notification upon logging in to the system.

    Windows Event Log, UNIX Syslog
    The ID of the user who receives alert notification upon logging in to the system.

    Event Type (Windows Event Log only)
    The type of event that is recorded to the OS log.

    Facility (UNIX Syslog only)
    The location where the UNIX Syslog is sent. You can select User or Local.

    Level (UNIX Syslog only)
    The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

    Run script
    A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters that are listed in the **Specify Script** window are passed to the script when the alert is triggered.

    Email An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit email** to change the actual email message.

16. Save the constraint:
    a. Click **File** > **Save**. The Save As window opens.
    b. In the **Specify Constraint name** field, type a name for the constraint and click **OK**. The constraint is displayed in the navigation tree pane. The constraint is automatically checked whenever an appropriate scan job is run.

Editing a constraint:

Edit a constraint to change the acceptable and unacceptable file types, file sizes, and file owners, add or remove users and user groups to which the constraint applies, edit the constraint options, and change the alert settings.

To edit a constraint, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management** > **Constraints**.
2. Click the constraint that you want to edit. Information about the constraint you want to edit displays in the topic pane.
3. Add or remove the file systems, file system groups, computers, and computer groups against which to apply the constraint.
   a. In the **Available** field, click the storage resources that you want add. Click . The users or groups are moved to the **Current Selections** field.
   b. In the **Current Selections** field, click the storage resources that you want remove from the constraint. Click . The users or groups are moved to the **Available** field.
4. Click the **File Types** tab; the File Types page opens. You can edit the file types that you want to allow or forbid.
5. Click one of the following selections.

   Forbid file matching these patterns
   Do NOT allow files of the specified type on the storage resources that are selected on the File Systems page.

   Allow ONLY file matching these patterns
   Allow files ONLY of the specified type on the storage resources that are selected on the File Systems page.
6. Click the file types that you would like to allow or forbid in the **Choose a file type** list box. The file types that are listed in this box are predefined in the system.
   a. In the **Choose a file type:** field, click the file types that you would like to allow or forbid.
   b. Click . The file types are moved to the **Forbidden files** field.
7. You can create a custom file type:
   a. Enter the pattern in the**Or enter a pattern** field.
   b. Click . The file type is moved to the **Forbidden files** field.
8. Click the **Users** tab; the Users page opens. You can allow or forbid files that are owned by specific users
9. Click one of the following selections.

   Forbid file owned by selected users
   Do NOT allow files that are owned by the specified users.

   Allow ONLY file owned by the specified owner
   Allow files that are owned by the specified users only.
10. Select the users whose files you want to allow or forbid:
    a. In the **Available Users** field, click the users that you want to include.
    b. Click . The users are moved to the **Forbidden Users** field.

11. Click the **Options** tab; the Options page opens. You can edit alerts for files that are based on characteristics other than file type and file owner. You can also view a where-clause that is used as a filter for locating files that violate the constraint definition. This page contains the following information:

    Violating File Limits
    Indicate the maximum number of violating filenames that should be kept per agent.

    number of days since last access greater than
    Click this box and enter a length of time (in days) in which files have not been accessed. When a constraint job locates a file that has not been accessed in the length of time you indicated in this field, an alert is triggered.

    bigger than ##
    Check this box and enter a file size that should trigger an alert. When a constraint job locates a file of the size (or greater) that you indicated in this field, an alert is triggered. Select the unit of measurement for the value you entered in the list box. The valid choices are: bytes, **KB** (kilobytes)**MB** (megabytes), **GB** (gigabytes), and **TB** (terabytes).

    set-UID root
    Check this box trigger alerts on files that have the SET-UID root bit on (because this can be a security consideration).

    not backed up (Windows only) and number of days since last modification greater than
    Click this box and enter the number of days since a file was last modified and not backed up. When the constraint job locates a file that was modified and not backed up in the amount of days that you indicate in this field, an alert is triggered.

    File Filter text
    This section displays a *where* clause that will be used as a filter for locating files that violate the Constraint definition. The choices you make on this tab and the choices you made on the Users and File Types pages are reflected in the content of this *where* clause. To edit the *where* clause filter text, click **Edit Filter**. This will allow you to create new where conditions and groupings using the where clause interface.

12. Click the **Alert** tab; the Alert page opens. Use the **Alert** tab to define the alerting criteria for a constraint.
13. In the **Triggering Condition** area, specify the event that triggers the constraint alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the amount of space a file must consume to violate the constraint.
    c. In the **Value Units** field, select the unit of measurement. kilobytes, megabytes, or gigabytes.

    Once Data Manager finds enough files that violate the condition (that exceed the total size amount you specified), the constraint is considered violated and the alert is triggered. When the alert is triggered, the actions that you specify in the **Triggered Actions** section are taken.

14. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

    TEC / OMNIbus Event
    A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the `tivoliSRM.baroc` file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, or running programs.

    Login notification
    The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

    Login ID
    The ID of the user who receives alert notification upon logging in to the system.

    Windows Event Log, UNIX Syslog
    The ID of the user who receives alert notification upon logging in to the system.

    Event Type (Windows Event Log only)
    The type of event that is recorded to the OS log.

    Facility (UNIX Syslog only)
    The location where the UNIX Syslog is sent. You can select User or Local.

    Level (UNIX Syslog only)
    The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

    Run script
    A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters that are listed in the **Specify Script** window are passed to the script when the alert is triggered.

    Email An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit email** to change the actual email message.

15. Click **File** > **Save** to save your changes.

Requesting an archive or backup:

Request anIBM Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint. Tivoli Storage Manager protects your data from hardware failures and other errors by storing backup and archive copies of data on auxiliary storage.

Data Manager provides a method within the alerting function of a constraint where you can automatically start a Tivoli Storage Manager archive or backup to run against the files that violate that constraint. For example, you can do the following actions:

- Define a constraint to identify the .vbs files within your organization that are larger than 2 megabytes. Other file types include .avi, .dll, .doc, .exe, .gif, .jpg, .mp3, .rtf, .txt., wsh, and .zip. You can also define a pattern to use when identifying files that need to be archived or backed up.
- Define a constraint that runs Tivoli Storage Manager to archive and then delete those files

Before you begin

Before you can use the Tivoli Storage Manager archive and backup functions with Data Manager constraints, you must ensure that the following conditions are met:

- Your user ID has Administrator authority within Tivoli Storage Productivity Center.
- Tivoli Storage Manager client software is installed on each of the hosts where a Storage Resource agent is used to perform archive and backup functions.
- (For SAN attached storage only) Tivoli Storage Manager client software must be installed on a host computer where a Storage Resource agent is configured to act as the scanning agent for NAS storage.
- A Tivoli Storage Manager client must be configured to run without a password. This condition is typically done by either including the password in the client options (dsm.opt) file, or using the options PASSWORDACCESS = GENERATE in the client options file.
- The Tivoli Storage Manager client must be version 5.1 or higher

Defining an archive or backup in a constraint:

Learn how to request an IBM Tivoli Storage Manager archive and backup of the largest violating files that are identified by a constraint. Request a Tivoli Storage Manager archive and backup on the alert page of a constraint definition.

Create a constraint.

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management**.
2. Right-click **Constraints**, and click **Create Constraint**. The Create Constraint window opens.
3. Select file systems, file system groups, computers, and computer groups against which to apply the constraint.
   a. In the **Available** field, click the storage resources against which you want to apply the constraint. Select individual file systems and computers, and file system and computer groups that you previously defined by using the group facility. Cluster resource groups in a IBM PowerHA SystemMirror for AIX or MSCS clustered environment are listed as computers.
   b. Click

. The users or groups are moved to the **Current Selections** field.
4. Click the **File Types** tab; the File Types page opens. You can select the file types that you want to allow or forbid.
5. Click one of the following selections.

Forbid file matching these patterns
Do NOT allow files of the specified type on the storage resources that are selected on the File Systems page.

Allow ONLY file matching these patterns
Allow files ONLY of the specified type on the storage resources that are selected on the File Systems page.

6. Click the file types that you would like to allow or forbid in the **Choose a file type** list box. The file types that are listed in this box are predefined in the system.
   a. In the **Choose a file type:** field, click the file types that you would like to allow or forbid.
   b. Click . The file types are moved to the **Forbidden files** field.
7. Optional: Create a custom file type:
   a. Enter the pattern in the **Or enter a pattern** field.
   b. Click . The file type is moved to the **Forbidden files** field.
8. Click the **Users** tab; the Users page opens. You can allow or forbid files that are owned by specific users
9. Click one of the following selections.

   Forbid file owned by selected users
   Do NOT allow files that are owned by the specified users.

   Allow ONLY file owned by the specified owner
   Allow files that are owned by the specified users only.
10. Select the users whose files you want to allow or forbid:
    a. In the **Available Users** field, click the users that you want to include.
    b. Click . The users are moved to the **Forbidden Users** field.
11. Click the **Options** tab; the Options page opens. You can define alerts for files that are based on characteristics other than file type and file owner. You can view a where clause that is used as a filter for locating files that violate the constraint definition. The following information is included.

    Violating File Limits
    Indicate the maximum number of violating filenames that should be kept per agent.

    number of days since last access greater than
    Click this box and enter a length of time (in days) in which files have not been accessed. When a constraint job locates a file that has not been accessed in the length of time you indicated in this field, an alert is triggered.

    bigger than ##
    Check this box and enter a file size that should trigger an alert. When a constraint job locates a file of the size (or greater) that you indicated in this field, an alert is triggered. Select the unit of measurement for the value you entered in the list box. The valid choices are: bytes, **KB** (kilobytes)**MB** (megabytes), **GB** (gigabytes), and **TB** (terabytes).

    set-UID root
    Check this box trigger alerts on files that have the SET-UID root bit on (because this can be a security consideration).

not backed up (Windows only) and number of days since last modification greater than

Click this box and enter the number of days since a file was last modified and not backed up. When the constraint job locates a file that was modified and not backed up in the amount of days that you indicate in this field, an alert is triggered.

File Filter text

This section displays a *where* clause that will be used as a filter for locating files that violate the Constraint definition. The choices you make on this tab and the choices you made on the Users and File Types pages are reflected in the content of this *where* clause. To edit the *where* clause filter text, click **Edit Filter**. This will allow you to create new where conditions and groupings using the where clause interface.

12. Click the **Alert** tab; the Alert page opens. You can define the alerting criteria for a constraint.
13. In the **Triggering Condition** area, specify the event that triggers the constraint alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the amount of space a file must consume to violate the constraint.
    c. In the **Value Units** field, select the unit of measurement. kilobytes, megabytes, or gigabytes.

    Once Data Manager finds enough files that violate the condition (that exceed the total size amount you specified), the constraint is considered violated and the alert is triggered. When the alert is triggered, the actions that you specify in the **Triggered Actions** section are taken.
14. In the **Triggered Actions** area, select **Archive/Backup** and click **Define**. The Archive / Backup Options window opens. Click one of the following selections.

    Archive

    Perform an archive on the files that violate a constraint. Archives are copies of data from online storage copies to auxiliary storage. If you select the Archive option, you can also check the **Delete After Successful Archival** box to automatically delete the files that violated the constraint after the files are successfully archived.

    Selective Backup

    Perform a selective backup of the files that violate a constraint. With a selective backup, you can back up specified files. These files are not excluded in the include-exclude list and meet the requirement for serialization in the backup copy group of the management class that is assigned to each file.

    Incremental Backup

    Perform an incremental backup of the files that violate a constraint. With an incremental backup, you can back up new or changed files or directories from a client domain or from specified directories or files. These directories or files are not excluded in the include-exclude list and meet the requirements for frequency, mode, and serialization as defined by a backup copy group of the management class that is assigned to each file.

Tivoli Storage Manager Command Options
Enter more options that are passed on to the Tivoli Storage Manager **dsmc** command. You can enter up to 512 bytes of text.

15. Save the constraint:
    a. Click **File** > **Save**. The Save As window opens.
    b. In the **Specify Constraint name** field, type a name for the constraint and click **OK**. The constraint is displayed in the navigation tree pane. The constraint is automatically checked whenever an appropriate scan job is run.

    For more information about each of the TSM actions you can select when you define triggered actions for a constraint, see the *IBM Tivoli Storage Manager Client User's Guide*

Scheduling script-based actions

Use scheduled actions to run user-defined scripts against selected computers and computer groups.

You can access the scheduled actions facility using either Data Manager or Data Manager for Chargeback. The actions that you schedule to run are based on user-defined scripts, allowing you to use any third-party tools for actions, such as recovery or provisioning. You can schedule scripts to run:

- Immediately
- Once at a specified time and date
- Repeatedly according to a schedule you specify

Select the option you want and then edit the time and date information accordingly. You can also select the time zone for the schedule you specify. You can either select the time zone that the Data server is in or any other time zone.

Note: Scheduled actions should not be confused with triggered actions. Triggered actions enable you to run scripts based upon conditions and alerts that are detected within your storage environment. Scheduled actions enable you to run scripts according to the schedule that you define, without having to associate the script with an alert.

Running scripts using scheduled actions:

Use the scheduled action facility to run a script according to a schedule that you define. The scheduled action facility is available in Data Manager and Data Manager for Databases.

Use scheduled actions to run scripts against selected computers and computer groups. The scheduled actions facility is script-based, so you can use any third-party tools for actions, such as recovery, or provisioning. You can schedule scripts to run:

- Immediately
- Once at a specified time and date
- Repeatedly according to a schedule you specify

Scheduled actions are different from triggered actions. Use triggered actions to run scripts based upon conditions and alerts that are detected within your storage environment. Use scheduled actions to run scripts according to the schedule that you define, without having to associate the script with an alert.

To create a scheduled action, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management** > **Scheduled Actions** or **Data Manager for Databases** > **Policy Management** > **Scheduled Actions**.
2. Right-click **Scripts**, and click **Create Script**. The Create Script panel opens.
3. In the **Description** field, type a brief description of the job.
4. Select computers and computer groups against which to run the action:
   a. In the **Available** field, click the computers and computer groups against which you want to run the action. You can select individual computers as well as computer groups that you have previously defined using the group facility. Cluster resource groups in a PowerHA SystemMirror for AIX or MSCS clustered environment are listed as computers.
   b. Click

. The computers and computer groups are moved to the **Current Selections** field.
5. Click the **Script Options** tab; the Script Options page opens. Use this page to select the script that you want to run against the selected computers and computer groups.
6. Select a script from **Script Name**. The scripts that appear in **Script Name** are stored in the \<*TPC_installation_directory*>\scripts directory on the Data server. The default installation directory where scripts are located is:
   - Windows: C:\Program Files\IBM\TPC\data\scripts
   - UNIX, Linux: /opt/IBM/TPC/data/scripts or /usrs/IBM/TPC/data/scripts

   Note: The scripts that appear in the Script Name list are stored on the computer where the Data server component is located. If you want to run a script that is stored on a Storage Resource agent, you must type the name of that script in the Script Name field.
7. Click the **When to Run** tab; the When to Run page opens. Use the When to Run page to specify when the scheduled action is run.
8. In the **How often to run** area, schedule how often you want to run the script that is associated with the action. You can choose from the following options:

   Run Now
   Run the script immediately.

   Run Once at
   Run the script one time at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

   Note: TheTivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This process ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

   Run Repeatedly
   Run the script repeatedly according to a schedule you specify. You can:
   - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job is to begin running.

- Indicate how often the job is to run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days that you select.

9. In the **How to Handle Time Zones** area, indicate the time zone to be used when you schedule and run the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields to run it at these times.
   a. 9:00 AM of the time zone where the Data server is located
   b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job runs at 7:00 AM on an agent that is in the PST time zone.

   Use the timezone that the server runs in
   Select this option to use the time zone where the Data server is located.

   Use this timezone
   Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.
10. Click the **Alert** tab; the Alert page opens. Use the Alert page to define the alerting criteria for a scheduled action.
11. In the **Triggering Condition** area, the only condition is **Script Failed**. An alert is automatically triggered if the scheduled action fails to run.
12. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
13. Save the scheduled action:
    a. Click **File** > **Save**. The Save As window opens.
    b. In the **Specify Script name** field, type a name for the job and click **OK**. The scheduled action is displayed in the navigation tree pane.

Editing a scheduled action:

Edit a scheduled action to change: the computers and groups against which the action is run, the script that is run, when the script is run, and the alerting conditions if the action fails to run.

To edit a scheduled action, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management** > **Scheduled Actions** > **Scripts** or **Data Manager for Databases** > **Policy Management** > **Scheduled Actions** > **Scripts**.
2. Select the scheduled action that you want to edit. Information about the action appears in the right-hand content pane.
3. In the **Description** field, type a brief description of the job.
4. Select computers and computer groups against which to run the action:
   a. In the **Available** field, click the computers and computer groups against which you want to run the action. Select individual computers, as well as computer groups that you have previously defined using the group facility. Note that cluster resource groups in a PowerHA SystemMirror for AIX or MSCS clustered environment are listed as computers.

b. Click . The computers and computer groups are moved to the **Current Selections** field.

5. Click the **Script Options** tab; the Script Options page opens. Use this page to change the script you want to run against the selected computers and computer groups.
6. Select a script from **Script Name**. The scripts that appear in **Script Name** are stored in the \<*TPC_installation_directory*>\scripts directory on the Data server. The default installation directory where scripts are located is:
   - Windows: C:\Program Files\IBM\TPC\data\scripts
   - UNIX, Linux: /opt/IBM/TPC/data/scripts or /usrs/IBM/TPC/data/scripts

   Note: The scripts that appear in the Script Name list are stored on the computer where the Data server component is located. If you want to run a script that is stored on a Storage Resource agent, you must type the name of that script in the Script Name field.
7. Click the **When to Run** tab; the When to Run page opens. Use this page to change when the scheduled action is set to run.
8. In the **How often to run** area, schedule how often you want to run the script that is associated with the action. You can choose from the following options:

   Run Now
   Run the script immediately.

   Run Once at
   Run the script one time at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

   Note: TheTivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This process ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

   Run Repeatedly
   Run the script repeatedly according to a schedule you specify. You can:
   - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job is to begin running.
   - Indicate how often the job is to run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
   - Indicate which days on which run the job. A check mark is displayed next to the days that you select.
9. In the **How to Handle Time Zones** area, indicate the time zone to be used when you schedule and run the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields to run it at these times.
   a. 9:00 AM of the time zone where the Data server is located
   b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job runs at 7:00 AM on an agent that is in the PST time zone.

Use the timezone that the server runs in
Select this option to use the time zone where the Data server is located.

Use this timezone
Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.

10. Click the **Alert** tab; the Alert page opens. Use the Alert tab to change the alerting criteria for a scheduled action.
11. In the **Triggering Condition** area, the only condition is **Script Failed**. This indicates that an alert will automatically be triggered if the scheduled action fails to run.
12. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

SNMP trap
An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

TEC / OMNIbus Event
A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the `tivoliSRM.baroc` file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, or running programs.

Login notification
The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who receives alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who receives alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that is recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog is sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters that are listed in the **Specify Script** window are passed to the script when the alert is triggered.

Email An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit email** to change the actual email message.

13. Click **File** > **Save** to save your changes.

Archive and backup functions

Data Manager provides a method within the alerting function of a constraint by which you can automatically invoke a IBM Tivoli Storage Manager archive or backup to run against the files that violate that constraint.

An archive is a "snapshot in time". Each time an archive is done, a new "snapshot" is taken. When an archive is performed, everything that has been specified is captured. So, for example, if an archive was performed on a C: drive, everything on C: will be archived. If, the next day, another archive of C: was performed, everything on the C: drive would be archived again. Files and file systems that have been archived as part of the Archive/Backup job can be deleted after archiving, but you must select **Delete After Successful Archival** specifically.

Lets say that, as a result of running a report, you find a large number of files that haven't been accessed in over two years. Looking at the cumulative sizes of these files you determine they are taking up space that would be better utilized elsewhere. You can archive and delete these files to reclaim that space. If, at some time in the future those files are needed, they can be retrieved from the archive. The result of the archive job is more free space without data loss.

A backup can be incremental (**Selective Backup**) or full. Usually a full backup is done first and subsequent backups are incremental. For example, say that a backup is performed on C: drive. All the files on the C: drive are then backed up, which is the full backup operation. The next day we need to back up C: drive again. Usually, this will be an incremental backup since a full backup was done the day before. In the incremental backup, only the files that have changed on the drive since the last backup will be backed up.

The archived or backed up files are sent to theTivoli Storage Manager server and are stored on that machine. Files are put into storage groups that are controlled through a Tivoli Storage Manager policy that the Tivoli Storage Manager server administrator usually creates.

Constraints

Data Manager provides a method within the alerting function of a constraint by which you can automatically invoke a Tivoli Storage Manager archive or backup to run against files that violate constraints you have defined. You can do the following:

- Define a constraint that will invoke Tivoli Storage Manager to archive and (if that option is selected) delete those files
- Define a constraint to run an incremental or full backup

See the scenario *Archive and delete files* to see an example of how constraints are developed.

Before you begin

Before you can use the Tivoli Storage Manager archive/backup functionality with Data Manager constraints, you must ensure the following:

- Your user ID has Administrator authority within IBM Tivoli Storage Productivity Center.
- Tivoli Storage Manager client software is installed on each of the hosts where a Storage Resource agent is used to perform archive and backup functions.
- (for SAN attached storage only) Tivoli Storage Manager client software must be installed on the Storage Resource agent host computer that is configured to act as the scanning agent for NAS storage.
- A Tivoli Storage Manager client must be configured to run without a password. This is typically done by either including the password in the client options (dsm.opt) file, or using the options PASSWORDACCESS= GENERATE in the client options file.
- The Tivoli Storage Manager client must be version 5.1 or higher.

Create an archive/backup job:

Create a IBM Tivoli Storage Manager archive/backup job for files in a report

You follow these general steps when you create an archive or backup.

1. Make sure you have followed the prerequisite steps
2. Decide whether to create a new constraint or use one of the default constraints
3. Generate the report using the constraint you decided to use
4. Create an archive or backup job from files in the report

To create an Archive/Backup job on files from a report, complete the following steps:

1. Generate and view a Largest Files, Most Obsolete Files, Orphan Files, Duplicate Files, Most At Risk, Constraint Violations report, or a report that you have created.
2. Click on the files you want to include in the archive or backup job. Use ctrl + click to select multiple, non-consecutive files; use shift + click to select multiple, consecutive files.
3. Right click on the reports and select **Create a new archive/backup job** from the pop-up menu. The **Create File Archive/Backup Action** window is displayed.
4. In the **Description** field, type a brief description of the job.
5. In the **Archive/Backup Options** tab, select the type of the job and its TSM command options, if any.

   Archive

   Use this option to archive the files that violate a constraint. Archives are copies of data from online storage copies to auxiliary storage.

   Delete After Successful Archival

   Automatically delete the files that violated the constraint after they are successfully archived. See the *IBM Tivoli Storage Manager Client User's Guide* for more information about Tivoli Storage Productivity Center archives.

   Note: Tivoli Storage Productivity Center administrators (Productivity Center Administrator role) can delete files even if they do not have the authority on the systems where the files are stored. This authority is provided by the Tivoli

Storage Productivity Center agents, which run as root on UNIX systems or with administrator privileges on Windows computers.

Selective Backup
Select this option to run a selective backup of the files that violate a constraint. You can back up specified files with a selective backup. These files are not excluded in the include-exclude list. They meet the requirement for serialization in the backup copy group of the management class that is assigned to each file. For more information about selective backups, see the *IBM Tivoli Storage Manager Client User's Guide*.

Incremental Backup
Select this option to run an incremental backup of the files that violate a constraint. You can do an incremental backup to back up new or changed files or directories from a client domain or from specified directories or files. These directories or files are not excluded in the include-exclude list. They meet the requirements for frequency, mode, and serialization as defined by a backup copy group of the management class that is assigned to each file. For more information about incremental backups, see the *IBM Tivoli Storage Manager Client User's Guide*.

IBM Tivoli Storage Manager Command Options
Enter more Tivoli Storage Manager command-line options in this field that is passed on to the Tivoli Storage Manager **dsmc** command. You can enter up to 512 bytes of text. See the *IBM Tivoli Storage Manager Client User's Guide* for more information about valid Tivoli Storage Manager command-line options.

6. Click the **Files** tab; the **Files** page opens. Use this page to view or remove the files included in an archive/backup job.
7. Click the **When to Run** tab. Use the **How often to run** area to specify when the archive or backup job runs.
8. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

   Run Now
   Check the quota immediately.

   Run Once at
   Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

   Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.
9. In the **How to Handle Time Zones** area, indicate the time zone to be used when you schedule and run the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields to run it at these times.
   a. 9:00 AM of the time zone where the Data server is located
   b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job runs at 7:00 AM on an agent that is in the PST time zone.

Use the timezone that the server runs in
Select this option to use the time zone where the Data server is located.

Use this timezone
Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.

10. Click the **Alert** tab; the Alert page opens. Use the Alert page to define the alerting criteria for an archive or backup job.
    a. In the **Triggering Condition** area, the only condition is **File Archive/Backup Failed**. This condition indicates that an alert is triggered automatically if the archive or backup fails to run.
    b. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

    TEC / OMNIbus Event
    A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the `tivoliSRM.baroc` file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, or running programs.

    Login notification
    The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

    Login ID
    The ID of the user who receives alert notification upon logging on to the system.

    Windows Event Log, UNIX Syslog
    The ID of the user who receives alert notification upon logging on to the system.

    Event Type (Windows Event Log only)
    The type of event that is recorded to the OS log.

    Facility (UNIX Syslog only)
    The location where the UNIX Syslog is sent. You can select User or Local.

    Level (UNIX Syslog only)
    The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

    Run script
    A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters that are listed in the **Specify Script** window are passed to the script when the alert is triggered.

Email An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit email** to change the actual email message.

11. Save the archive or backup:
    a. Click **File** > **Save**. The **Save As** window opens.
    b. In the **Specify File Archive/Backup Action name** field, type a name for the job and click **OK**. The new job is listed in **Data Manager** > **Policy Management** > **Archive/Backup** as **<creator>.<job name>**.

View and edit an Archive/Backup job:

Learn how to edit IBM Tivoli Storage Manager archive and backup jobs that were defined to run against the files selected from Data Manager reports or from reports you have generated yourself.

Each archive and backup that you have created appears as a node under **Archive/Backup**.

To view and edit a Tivoli Storage Manager archive or backup job, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management** > **Archive/Backup**. Select the archive or backup job that you want to edit. Information about the job appears in the content pane. All the user-defined information can be edited.
2. Change the description for the archive/backup job.
3. In the **Options** tab, you can select the Archive/Backup options. You can change whether the job is an archive or a backup, and add or change any additional Tivoli Storage Manager commands that should be run with the job.
4. In the **Files** tab the files that were included in the job are listed.
5. In the **When to Run** tab, select how often to run and how to handle time zones. You can change the **How often to run** and **How to handle time zones** settings.
6. In the **Alert** tab, **File Archive/Backup Failed** is the only **Triggering-Condition**. In **Triggered-Actions**, change the action you wish to occur if the job fails. You do not need to select any of the actions. Notification will be written to the error log by default.
7. Select **File** > **Save**. Name the Archive/Backup job. After you save a the job, an entry for it will appear under the **Data Manager** > **Policy Management** > **Archive/Backup** > **<Creator>.<Name>**.

Add additional files to an archive/backup job:

Add files to an existing IBM Tivoli Storage Manager archive or backup job from a report.

To add files to an existing Archive/Backup job, complete the following steps:

1. Add files to an existing Tivoli Storage Manager archive or backup job from a report
2. Click on the files you want to add to an archive or backup job. Use ctrl + click to select multiple, non-consecutive files; use shift + click to select multiple, consecutive files.
3. Right click on the report and select **Add to an existing archive/backup job** from the pop-up menu. The Archive/Backup Job Selection window appears.

4. Select the archive or backup job to which you want to add the files from the list box.
5. Click **OK** to add the files. The Archive/Backup - Files window appears and shows the file(s) you added to the selected job.

Working with storage resource groups

Use storage resource groups to organize logically related storage entities into named groups. For example, a storage resource group that represents a business critical email application might include all the hosts that participate in the email application cluster, the storage subsystems that provide storage to the application, and the switches and fabrics through which the application's data is configured to travel.

Storage resource group management

Use the storage resource group management list panel to view, create, and delete storage resource groups. A storage resource group contains a combination of storage subsystems, switches, fabrics, and other monitored resources. These resources can be a logical grouping that is based on a particular application or usage, such as resources that are related accounting or payroll.

This panel contains the following information about each storage resource group:

Storage resource group name
Specifies a user-defined name of the storage resource group.

State Specifies the state of the storage resource group:

Normal
The storage resource group is operating normally.

Warning
At least one part of the storage resource group is not operating or has serious problems.

Critical
The storage resource group is either not operating or has serious problems.

Missing
The storage resource group was previously discovered by IBM Tivoli Storage Productivity Center, but was not detected in the most recent discovery.

Unknown
The storage resource group was discovered but is not recognized by IBM Tivoli Storage Productivity Center.

Note: The state of a storage resource group might display as **Unknown** after it is created. In this case, the status of a new storage resource group might not yet be updated when the Storage Resource Group Management panel is displayed again. To resolve this issue, exit the Storage Resource Group Management panel, wait for the processing of the group to complete, and open the Storage Resource Group Management panel again to view the updated **State**.

Description (optional)
Displays the user-defined description for the storage resource group.

Create Creates a storage resource group that you selected.

Delete
Deletes a storage resource group that you selected.

Refresh Table
Refreshes the storage resource group management table with the most recent changes.

Create or Edit a Storage Resource Group

Use the Create or Edit Storage Resource group panel to create or modify storage resource groups.

To create a storage resource group:

- In the navigation tree, expand **IBM Tivoli Storage Productivity Center**, click **Storage Resource Group Management**, then click **Create** in the content pane.
- In the navigation tree, expand **IBM Tivoli Storage Productivity Center** > **Topology**, click a node that represents a storage resource, right-click a resource that is displayed in the content pane, and select **Add to new Storage Resource Group** from the menu.

To edit a storage resource group:

- In the navigation tree, expand **IBM Tivoli Storage Productivity Center**, click **Storage Resource Group Management**, then click the magnifying glass icon next to the storage resource group you want to edit.
- In the navigation tree, expand **IBM Tivoli Storage Productivity Center** > **Topology**, click **Storage Resource Groups**, right-click the storage resource group to edit in the content pane, and select **Launch Detail Panel** from the menu.

Field descriptions

Creator
Displays the user name of the creator.

Name Displays the name of the storage resource group or `unnamed`, if it is not yet named.

Description
Optional: Displays the user-defined description for the storage resource group.

Selected Elements
Lists the elements that are selected to be members of this storage resource group.

Add Adds one or more selected elements to the list.

The Storage resource group element selection panel is displayed.

Remove
Removes one or more selected elements that are shown in the **Selected Elements** list.

User-defined property 1 (UDP1)
Specifies any user-defined properties to be used by the topology viewer for custom groupings.

User-defined property 2 (UDP2)
Specifies any user-defined properties to be used by the topology viewer for custom groupings.

User-defined property 3 (UDP3)
Specifies any user-defined properties to be used by the topology viewer for custom groupings.

Storage Resource Group Element Selection:

Use the Storage Resource Group Element Selection panel to select the monitored resources that you want to include in a storage resource group. The resources that you can select are displayed in the topology viewer.

Available elements section

Displays the resources that you can include in a storage resource group. Use the topology viewer to select storage resources. For more information about the topology viewer, see Overview of Topology Viewer.

Topology viewer
Use the topology viewer to select the resources that you want to include in a storage resource group. To do select resources, complete the following steps.

1. Open a detailed view of a storage resource type on the Overview page in the topology viewer.
2. Expand the storage resource type to view the resources that are part of that type.
3. Select the resources that you want to include in the storage resource group.
4. Click >>.

Selected elements section

>> Click >> to move the selected resources into the **Selected Elements** pane.

<< Click << to remove the selected resources from the **Selected Elements** pane.

Selected Elements
Displays the resources that are included in a storage resource group. You can include the following resources in a storage resource group:

- Computer (virtual machine, physical or hypervisor)
- Computer disk (including virtual machine disks)
- Fabric
- Switch
- Other entities
- Storage subsystem
- Storage Subsystem disk
- Storage Subsystem volume
- Storage pool
- Other storage resource groups

OK Click **OK** when you are satisfied with all your selections. The Create Storage Resource Group panel displays with your selections.

Cancel
Click **Cancel** if you do not want to add resources to a storage resource group. No data is saved.

Working with storage resource group sessions:

This topic describes viewing and deleting replication sessions that are created in the SAN Planner wizard for a storage resource group.

You can delete the replication session from IBM Tivoli Storage Productivity Center and optionally also from IBM Tivoli Storage Productivity Center for Replication.

1. In the navigation tree pane, expand **Replication Manager** > **Storage Resource Group Sessions**. A table of the currently defined storage resource group sessions is displayed in the content pane.
2. To select a storage resource group session to delete, click a row in the table. You can use Ctrl+click and Shift+click to select multiple sessions simultaneously.
3. Click **Delete Storage Resource Group Session**. Tivoli Storage Productivity Center does not automatically delete the replication session from Tivoli Storage Productivity Center for Replication. A dialog box opens enabling you to specify how you want to delete the session. You can select from the following options:
   - Click **Yes** to remove the selected replication session from the Tivoli Storage Productivity Center database and the Tivoli Storage Productivity Center for Replication database.
   - Click **No** to remove the selected replication session from the Tivoli Storage Productivity Center database only.
   - Click **Cancel** to close the dialog box without deleting the selected replication session from the Tivoli Storage Productivity Center database or Tivoli Storage Productivity Center for Replication database.

Working with scripts

Use the triggered action facility to run scripts based upon conditions/alerts that are detected within your storage environment. Use the scheduled action facility to run scripts according to the schedule that you define, without having to associate the script with an alert.

Both the triggered action and scheduled action facilities are script-based, allowing you to use any third-party tools for actions, such as recovery and provisioning.

Running scripts with alerts

Run scripts in response to storage events and conditions that are associated with data collection, alerting, and policy management jobs.

Use the triggered action facility to define thresholds and constraints that result in notification or action. You can run a script when a job fails, an alert condition is detected in a storage resource, a constraint is violated (Data Manager only), or a quota is exceeded (Data Manager only). Triggered actions are script-based, so you can use any third-party tools for actions such as recovery or provisioning.

Complete the following steps to run a script in response to a condition in a job:

1. Create or edit a monitoring, alerting, or policy management job in the stand-alone GUI. You can define a condition that triggers a script for the following jobs:
   - Data Manager:
     - **Monitoring** > **<all_monitoring_jobs>**
     - **Alerting** > **<all_alert_jobs>**

- **Policy Management** > **Constraints**
- **Policy Management** > **Quotas**
- **Policy Management** > **Scheduled Actions**

• Data Manager for Databases:
- **Monitoring** > **<all_monitoring_jobs>**
- **Alerting** > **<all_alert_jobs>**
- **Policy Management** > **Quotas**
- **Policy Management** > **Scheduled Actions**

• Disk Manager:
- **Monitoring** > **Subsystem Performance Alert Configurations**
- **Alerting** > **Storage Subsystem Alerts**

• Fabric Manager:
- **Monitoring** > **Switch Performance Alert Configurations**
- **Alerting** > **<all_alert_jobs>**

2. Check **Run Script** on the **Alert** page for a job.
3. Click **Define**.
4. In the Specify Script window, enter the name of the script in the **Script Name** field. This script name is used to tell the server what script to send to the agent. If the script sent by the server is run by the agent, the name of the script is not the name that you specify in the **Script Name** field. Instead, the agent names the script as temporary file, runs it, then deletes it after the script is run. Because of this process, do not include code in a script that depends on the script name.
5. Specify where you want the script to run in the **Where to Run** field. You can select a specific agent, or the triggering computer. For example, if a file system alert is triggered on a specific computer, typically you want the script to run on the computer that triggered the alert. For some types of alerts such as Job Failed, the option to run the script on the triggering computer is not available.
6. The **Script Parameters** section shows the parameters that are passed to the script when it is run. Each triggering condition passes different parameters. The parameters are passed in the order that is specified in the GUI.
7. Click **OK**.

Running scripts using scheduled actions

Use the scheduled action facility to run a script according to a schedule that you define. The scheduled action facility is available in Data Manager and Data Manager for Databases.

Use scheduled actions to run scripts against selected computers and computer groups. The scheduled actions facility is script-based, so you can use any third-party tools for actions, such as recovery, or provisioning. You can schedule scripts to run:

- Immediately
- Once at a specified time and date
- Repeatedly according to a schedule you specify

Scheduled actions are different from triggered actions. Use triggered actions to run scripts based upon conditions and alerts that are detected within your storage environment. Use scheduled actions to run scripts according to the schedule that you define, without having to associate the script with an alert.

To create a scheduled action, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** > **Policy Management** > **Scheduled Actions** or **Data Manager for Databases** > **Policy Management** > **Scheduled Actions**.
2. Right-click **Scripts**, and click **Create Script**. The Create Script panel opens.
3. In the **Description** field, type a brief description of the job.
4. Select computers and computer groups against which to run the action:
   a. In the **Available** field, click the computers and computer groups against which you want to run the action. You can select individual computers as well as computer groups that you have previously defined using the group facility. Cluster resource groups in a PowerHA SystemMirror for AIX or MSCS clustered environment are listed as computers.
   b. Click

. The computers and computer groups are moved to the **Current Selections** field.
5. Click the **Script Options** tab; the Script Options page opens. Use this page to select the script that you want to run against the selected computers and computer groups.
6. Select a script from **Script Name**. The scripts that appear in **Script Name** are stored in the `\<TPC_installation_directory>\scripts` directory on the Data server. The default installation directory where scripts are located is:
   - Windows: `C:\Program Files\IBM\TPC\data\scripts`
   - UNIX, Linux: `/opt/IBM/TPC/data/scripts` or `/usrs/IBM/TPC/data/scripts`

   Note: The scripts that appear in the `Script Name` list are stored on the computer where the Data server component is located. If you want to run a script that is stored on a Storage Resource agent, you must type the name of that script in the `Script Name` field.
7. Click the **When to Run** tab; the When to Run page opens. Use the When to Run page to specify when the scheduled action is run.
8. In the **How often to run** area, schedule how often you want to run the script that is associated with the action. You can choose from the following options:

   Run Now
   Run the script immediately.

   Run Once at
   Run the script one time at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

   Note: TheTivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This process ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

   Run Repeatedly
   Run the script repeatedly according to a schedule you specify. You can:
   - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job is to begin running.

- Indicate how often the job is to run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days that you select.

9. In the **How to Handle Time Zones** area, indicate the time zone to be used when you schedule and run the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields to run it at these times.
   a. 9:00 AM of the time zone where the Data server is located
   b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job runs at 7:00 AM on an agent that is in the PST time zone.

   Use the timezone that the server runs in
   Select this option to use the time zone where the Data server is located.

   Use this timezone
   Select this option to indicate a specific time zone for when to run the job. When you select this option, a drop-down list is activated. Select the timezone from this drop-down list.
10. Click the **Alert** tab; the Alert page opens. Use the Alert page to define the alerting criteria for a scheduled action.
11. In the **Triggering Condition** area, the only condition is **Script Failed**. An alert is automatically triggered if the scheduled action fails to run.
12. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see "Triggered actions for alerts" on page 1682.
13. Save the scheduled action:
    a. Click **File** > **Save**. The Save As window opens.
    b. In the **Specify Script name** field, type a name for the job and click **OK**. The scheduled action is displayed in the navigation tree pane.

How scripts are run

Scripts can be run from the Data server or Tivoli Storage Productivity Centers.

During installation, IBM Tivoli Storage Productivity Center creates a directory on every server and agent computer called: \<*install_directory* >\scripts, where <*install_directory* > represents the directory where the agent or server is installed. The default installation directory is:

- (Windows Storage Resource agents) `C:\Program Files\IBM\TPC\agent\scripts`
- (UNIX, Linux Storage Resource agents) `/opt/IBM/TPC/agent/scripts` or `/usr/tivoli/ep/subagents/TPC/Data/scripts`
- (Windows Data server) `C:\Program Files\IBM\TPC\Data\scripts`
- (UNIX, Linux Data server) `/opt/IBM/TPC/Data/scripts` or `/usr/IBM/TPC/Data/scripts`

When a script is run:

1. The agent looks in the agent's local `\scripts` directory; if the agent can resolve that script name, it is executed from that location.
2. Otherwise (if the agent cannot resolve the script name) and the agent is configured to accept scripts from the server, the server looks in the local

\scripts directory. If the server can resolve that script name, it loads that script and sends it to the agent where you designated the script to be run.

When you run a script against a NAS filer, the script is run from the agent that is assigned to the file system where the triggering condition occurred.

3. The agent receives the script, names it as a temporary file, and runs it.
4. After the agent finishes running the script, the temporary file is deleted. Because the script is handled in this way, any script which is written to redirect any output must specify a fully qualified unique directory path for that output.

You run a script:

- In response to an alert
- For batch reports

There are two reasons why a script might not run when sent by the server:

- The script exists in the \scripts directory fro the agent. In this case, the agent ignores the script sent by the server and run its local version.
- If you did not check the **Agent may run scripts sent by server** option in the installation parameters screen, the agent ignores the script sent by the server.

  If you selected the **Agent may run scripts sent by server** option in the installation parameters screen, the agent runs scripts sent by the server.

  The advantage of selecting this option is you need only store one copy of the script in the \scripts directory for the server and do not have to keep a copy of the script on every agent computer. When a script must be run on a particular agent, the server accesses the script from the local \scripts directory and sends it to the agent to be run.

  Note:

  – You can change this option by editing the honorSentScripts parameter in the agent.config file. The agent.config file exists in every installation directory for the agent. See *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for information about the agent.config file.

  – After a Storage Resource agent is deployed, you can change the script configuration in the user interface. To change the script configuration, complete the following steps:

    1. From the navigation pane in the web-based GUI, click **Server Resources** > **Servers**.
    2. On the Servers page, right-click the server where the agent is deployed. Select **Modify Agents** > **Enable running scripts on agent** or **Modify Agents** > **Disable running scripts on agent** to enable or disable the running of scripts.

When the script is deployed to the agent, it is named as a temporary file. After the agent finishes running the script, the temporary file is deleted.

If the **Agent may run scripts sent by server** option is cleared, you must make sure that the script is stored in every \scripts directory for the agent. If the **Agent may run scripts sent by server** option is checked, the server deploys the script to the agent computer that has the low free space condition. In this case, you must ensure that the script exists in the \scripts directory for the server, and not on any of the \scripts directory for the agents.

Note: A good practice is to store the scripts you want to run globally (across all your agents) on the server. Store the scripts that you want to run locally on the agent computers where you want them to run.

How the server and agent resolve script names

An IBM Tivoli Storage Productivity Center server processes requests for agents running under UNIX, Linux, and Windows.

When associating and running scripts with triggered and scheduled actions, the Tivoli Storage Productivity Center server resolves the names of those scripts based on the following criteria:

- The OS under which an agent is running
- How the script name is specified on the Specify Script dialog box

Agent and Operating System

The operating system under which an agent runs determines the scripts that can be run by that agent:

UNIX, Linux operating systems
An agent running under a UNIX or Linux operating system will not run or receive (from the server) scripts that have an extension. Agents running under UNIX or Linux only accept and run script files that do not contain an extension within their name. The first line in a script file determines what interpreter is needed to run the script.

Windows operating systems
An agent running under a Windows operating system will only run or receive (from the server) scripts that have an extension. The extension of a script file determines what interpreter is needed to run the script.

Specifying a Script Name

Use the Specify Script dialog box when associating a script with a triggered or scheduled action. Enter the name of a script on the **Script Name** field to associate that script with the action you want to perform.

The name of the script you enter might or might not contain an extension (including an extension is optional). Depending on the agent that receives the script, the following occurs after a triggered or scheduled action that contains a script is run:

Agent running under UNIX or Linux operating systems
If the specified script name contains an extension, the agent ignores the extension and searches for a file of the same name (without an extension). For example, if you enter `backup.vbs` in the **Script Name** field, a UNIX or Linux agent will search for a file with the name `backup`. If the script exists, UNIX or Linux will read the first line of the script and use the appropriate interpreter to run the script. If the file does not exist, the triggered or scheduled action will not run the script.

Agent running under Windows operating systems
If the specified script name contains an extension, the agent will search for a file of the same name/extension. For example, if you enter `backup.vbs` in the **Script Name** field, an agent will search for a file with the name `backup.vbs`. If the script does exist, the agent will determine what interpreter is needed to run the script based on its extension. If the file does not exist, the triggered or scheduled action will not run the script. If

the specified script name does not contain an extension, an agent will look for the file name in the following order (based on its extension): `.com, .exe, .bat, .cmd, .vbs, .vbe, .js, .jse`

The process for script name resolving enables you to run scripts with the same name across multiple operating systems. To run a script on both a UNIX or Linux and Windows operating systems, you must have two versions of that file on either the server or the agent. For example, if you want to run a provisioning script against a UNIX or Linux and Windows agent, you must have two versions of that file (`provision.bat` and `provision`) stored in the `\scripts` directory on either the server or agent.

Using the topology viewer

This topic describes how to use the topology viewer to view detailed information about your storage environment and monitor entities and groups of entities.

You can view health status, performance, and other details about switches, computers, storage subsystem servers, fabrics, and other entities. To launch the topology viewer, expand **IBM Tivoli Storage Productivity Center Topology** in the Navigation Tree pane.

To display detailed information about an entity (for example, the host name of a computer, the used and unused capacity of a storage subsystem), right-click the entity and select **Launch Detail Panel**.

To select an entity in the graphical view, click the entity image. In the table view, click the name of the entity. To make multiple selections, press CTRL and click.

Note:

- The availability of certain action items depends on the entity class and view level.
- When you right-click an object in the graphical view, a context menu is displayed. The action items in this menu are identical to the action items displayed when you right-click the corresponding object in the table view and click the arrow in the **Action** field.
- If the discovery process is not completed and there are entities in the environment that have not been rendered into the current view, the view will build dynamically as entities are discovered.
- Some columns and fields displayed within the Topology Viewer might not be populated with data for the storage resources that you are monitoring. Columns and fields appear empty if you have not configured Tivoli Storage Productivity Center to discover or collect the information intended for those columns and fields. Use discovery, probe, ping, or scan schedules to ensure that you collect information for all the columns and fields that appear in the Topology Viewer.

The topology viewer

This topic describes how the topology viewer provides a graphical representation of the physical and logical resources that have been discovered in your storage environment by IBM Tivoli Storage Productivity Center.

The information displayed by the topology viewer is collected by discovery, probing, performance data collection, and alerting. The topology viewer consists of two views, a graphical view and a table view, organized vertically with the graphical view over of the table view. The table view shows the same information as the graphical view, but in a tabular format.

The topology viewer supports four levels of detail, from a high-level overview of the entire storage environment to detailed information about individual entities. This allows you to view as much or as little detail as you need.

Overview

A global, highly aggregated view of the entire storage environment. This is the default view when you expand **IBM Tivoli Storage Productivity Center** > **Topology**.

Groups level (L0)

This level focuses on groups of a particular type. The default groups are:

- **Computers**: All discovered computer groups in your storage environment
- **Fabrics**: All fabric groups in your storage environment. A *fabric* is a network of entities that are attached through one or more switches. The topology viewer also displays Virtual SANs (VSANs) and Logical SANs (LSANs).
- **Switches**: All switches in your environment.
- **Storage**: All storage subsystems in your environment.
- **Other**: All entities that are discovered in your environment but that do not fit in either the computer, fabric or storage classes.

The groups level is the default level upon launching the topology viewer by clicking either **Computers**, **Fabrics**, **Switches**, **Storage** or **Other** under the **Topology** node. The groups level shows one or more groups of entities that correspond to the topology class that you selected. Information in the graphical view at this level can include any individually pinned entities.

Group level (L1)

This level focuses on one selected group and its related entities (for example, a group of computers). At L1, individual entities can be shown by expanding the group box. You can also display lines representing connections between entities or between entities and groups.

Detail level (L2)

This level focuses on individual entities (for example, a single computer, switch, or subsystem) and the paths to associated logical and physical entities. This level shows details all the way to the port level.

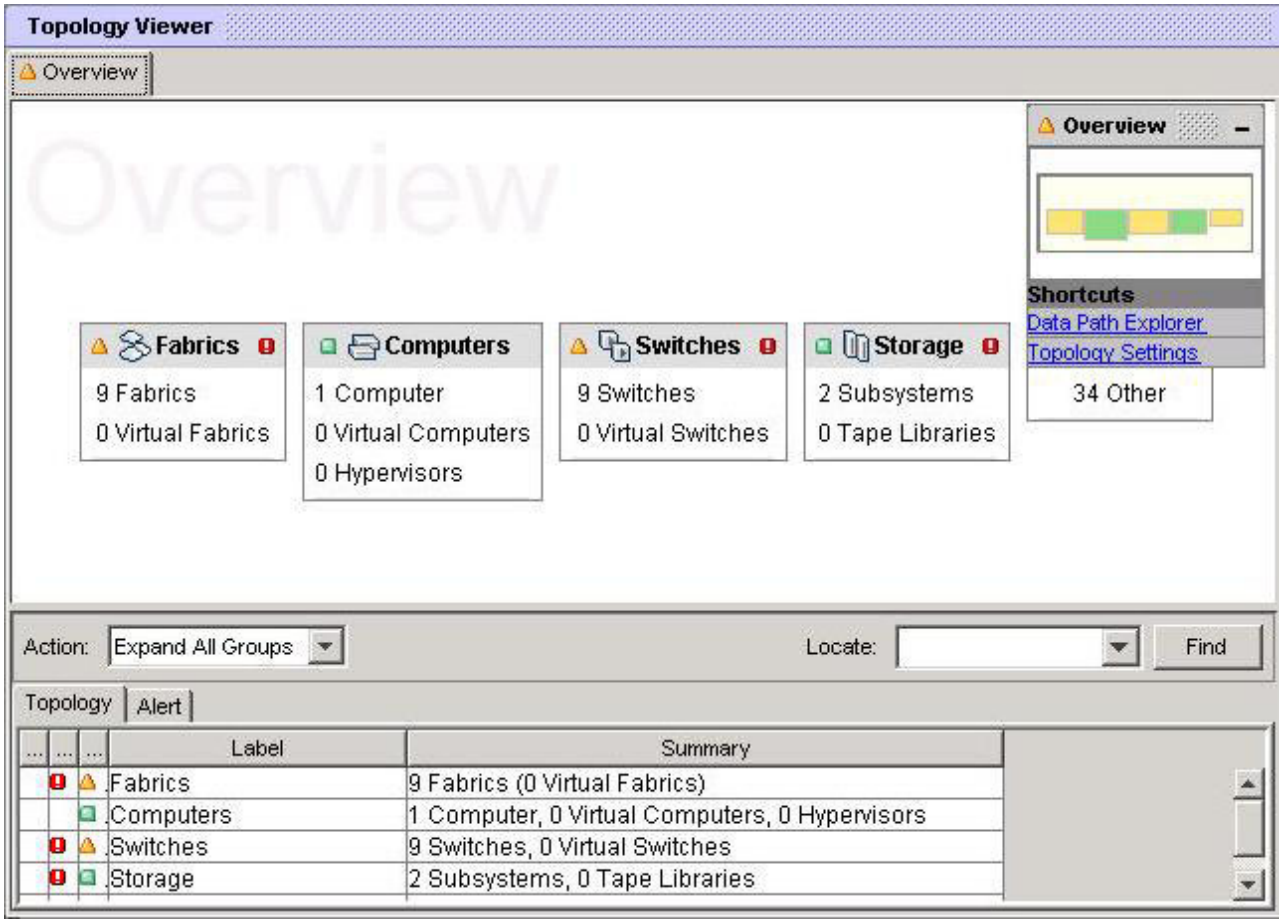

Figure 23. Topology overview

In addition, the topology viewer depicts the relationships among resources (for example, the disks comprising a particular storage subsystem). Detailed, tabular information (for example, attributes of a disk) is also provided. With all the information that the topology viewer provides, you can more quickly and easily monitor and troubleshoot your storage environment and perform critical storage-management tasks.

Topology viewer interface controls:

Progressive information disclosure reduces the visual complexity of a system. In the topology viewer, this means that you can identify and focus on entities in your storage environment that are of particular interest and hide other entities from view.

The tabs displayed in the tabular view represent classes of entities (a physical device or logical resource discovered byIBM Tivoli Storage Productivity Center). Individual entities (for example, a computer named *accounting1*) are listed in their corresponding tabs. Entity attributes (for example, the operating system of a computer) are organized into table columns. The classes entities are assigned to are Computers, Fabrics, Storage Subsystems, and Other.

Note: The Other class contains entities discovered by Tivoli Storage Productivity Center but for which the type cannot be determined.

Changes to the graphical view or the tabular view are automatically applied to the other view, keeping both views synchronized. For example, if an entity is selected in the graphical view, it is also highlighted in the table view.

As more of your resources are discovered and displayed in the topology viewer, it is not possible to see all the discovered resources in a single topology window. To make it easier for you to find the resources you want to view, Tivoli Storage Productivity Center provides a tool called the *minimap*. The minimap allows you to shift the focus of the topology viewer to different areas within a particular view. The minimap is a small window within the topology view.

Linux/AIX users: you must use SHIFT+ALT+mouse drag to reposition the minimap.

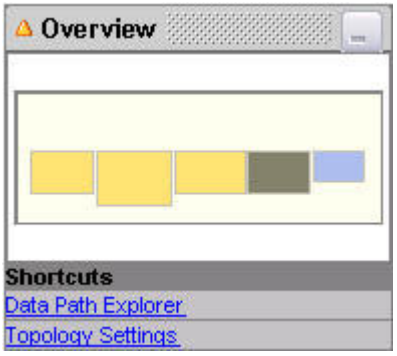

Figure 24. Minimap. You can click the shortcuts to launch Data Path Explorer or to view or change topology settings.

The minimap provides a miniature representation of the current topology view. Clicking on the minimap displays a multidirectional cursor that allows you to move the focus of the current topology view in any direction. As you drag the cursor around the minimap, the topology view will move accordingly. This makes navigating around complex topologies much easier.

The minimap uses the following colors to indicate the aggregated health status of groups:

- Green (normal): All entities in the group are operating normally.
- Yellow (warning): At least one entity is operating normally, and one or more entities have a health status of warning, critical, unknown, or missing.
- Red (critical): The status of all entities is either critical or warning.
- Gray (unknown): None of the entities in the group has an known health status.
- Black (missing): All of the entities of the group are missing.

In the Configuration History page, the minimap uses the following colors to indicate the aggregated change status of groups:

- Blue (changed): One or more entities in the group changed (an addition or deletion of an entity is considered a change).
- Gray (unchanged): All entities in the group are unchanged.

Included in the minimap are shortcuts that you can click to launch the Data Path Explorer or to view or change the topology settings:

- To launch the Data Path Explorer, select one or more hosts, subsystems, disks, volumes, or Mdisks and click the Data Path Explorer shortcut.

Note: The shortcut to Data Path Explorer is not available in the Configuration History topology view.

- To view or change the topology settings, click the Topology Settings shortcut. The Settings window displays for you to specify the type of overlays to display, the type of filtering to perform, whether connections between topology objects should display, and whether the Zone tab should display in the table view.

The graphical view:

This topic describes how the topology viewer groups entities that share certain characteristics. A collapsed group consisting of many entities or subgroups takes up much less space in the graphical view than the individual entities or subgroups would.

In the topology viewer, you can expand or collapse a group or subgroup by clicking the plus (+) character or minus (-) character, respectively, located in the group box.

Note: Pinned entities within a group remain visible, even when the group is collapsed.

By default, the topology viewer groups entities by class (computers, fabrics, switches, storage subsystems, and other entities). However, you can create additional groups (but not additional classes) using your own criteria. These custom groups can provide new perspectives on and important information about your storage environment. For example, you can create a special group belonging to the computer class that consists of computers that use the same operating system or storage subsystems of a particular model.

The following figure shows an expanded computer group representation in the graphical view:

Figure 25. Computer group

In the graphical view you can expand or collapse a group or subgroup by clicking the plus (+) character or minus (-) character, respectively, located in the group box.

The title bar of the group box shows the following:

- An icon indicating the health status of the group (for example, normal).
- An icon indicating the types of entities contained in the group.
- The class of group (for example, computers).
- The grouping criterion (for example, Windows).
- The total number of entities in the group. If filtering is turned on, the number in the brackets is the ratio of displayed entities to total entities. For example, 5/10 indicates that 5 entities are shown out of 10 total entities.

The topology viewer provides the following default groups. These groups are accessible, if they exist in your topology, in the overview or by expanding the primary default groups (**Computers**, **Fabrics**, **Switches**, **Storage** and **Other**) that are displayed when you first launch the topology viewer.

Table 76. Group locations in the topology viewer

Type of group	Description	Path to group in topology viewer
Computers	All the computers that are discovered in your environment or a group of computers associated with a particular fabric or switch	**Computers**
Fabrics	All the fabrics and virtual fabrics that are discovered in your environment	**Fabrics**
Storage	All the storage-related devices that are discovered in your environment	**Storage**
Other	All the entities that are discovered in an environment but that do not fit in either the computer, fabric, or storage categories	**Other**
Fabric	All entities (for example, switches) that belong to or are related to a fabric	**Fabrics** > **Fabric**
Virtual Fabric	All entities that belong to a virtual fabric	**Fabrics** > **Fabric**
Switches	A set of switches	**Fabrics** > **Fabric**
Subsystems	A set of subsystems	**Storage** > **Subsystems**
Device	Devices associated with a particular entity (for example, a computer)	**Computers** > **Computers**
Volumes	A set of volumes that are associated with an entity	Different types of volumes groups exist at the L2 level of each of the four main categories: **Computers**, **Fabrics**, **Storage**, **Other**
HBA	A set of HBAs that are associated with an entity	**Computers** > **Computers** > **Computer**
Disks	A set of disks that are associated with an entity	**Storage** > **Subsystems** > **Subsystem**
Pools	A set of pools that are associated with an entity	**Storage** > **Subsystems** > **Subsystem**
FC Ports	A set of FC ports that belong to an entity	**Fabrics** > **Fabric** > **Switch**
Switch	All ports on a switch. Due to the large number of ports on some switches, ports are shown as FC port groups within the switch	**Fabrics** > **Fabric** > **Switch**

As shown by the tabs of the graphical panel (**L0:Computers**, **L1:Computers**, and **L2:Computers**), other views have been opened but are hidden by the overview.

The topology tab in the table view displays information about the entities in the groups. You can use the minimap to shift the focus in the view.

Right-click the object and select **Open Detail View** to view entities in a more detailed view. For example, at a high level an object representing a storage subsystem is depicted as an icon with a label. In a more detailed view, the same storage subsystem can be rendered as a box containing sub-entities representing disks, with the corresponding table view showing attributes of each individual disk. At yet another detail level, the ports of the storage subsystem might be shown as well as status information for each port.

GUI Topology Viewer has new L0 and L1 Switches views:

The IBM Tivoli Storage Productivity Center GUI has a Switches node under the Topology node. The Switches node opens the new L0 Switches view, from which you can also access the new L1 Switches view.

Overview
The Overview panel now shows a box with the number of switches present in your environment.

L0 Switches
The L0 Switches view shows all switches in your environment.

If you right-click on a Switches group, then click Group By, you will get a dialog box for group settings. You can choose to group by a single group or you can group by multiple categories:

- ByHealth
- ByUDP1
- ByUDP2
- ByUDP3
- ByFabric
- ByGroupSize

Click OK to close the dialog box and update the group settings.

L1 Switches
The L1 Switches view shows information on a group of switches in your environment that you selected to display. This also includes connectivity information: what entities this group is connected to.

L2 Switch
This view shows one switch, with all detailed port and connectivity information. You can access this view by double-clicking a single switch icon in either the L0 or L1 Switches view, or through the other Topology Viewer views.

The table view:

This topic describes how tabs in the table view provide detailed information about the entities (for example, computers, subsystems, and fabrics) that are relevant in the current view.

The following is a list of all the tabs available in the topology viewer. Not all tabs are visible in all views.

Topology
Detailed information about all the entities in the topology grouped by class of the entity (computers, fabric, storage, and other).

Computer
Detailed information about any computer in the computer group or groups in the current topology view.

Fabric Detailed information about fabric groups and virtual fabric groups in the current topology view.

Switch
Detailed information about switch groups and switch entities in the current topology view.

Subsystem
Detailed information about volume groups and volume entities in the current topology view.

Disk Detailed information about disk groups and disk entities in the current topology view.

Pool Detailed information about pool groups and pool entities in the current topology view.

Note: One of the columns in the tabular view refers to whether a storage pool is primordial. A primordial pool represents unallocated storage capacity on a storage device. Storage capacity can be allocated from primordial pools to create storage pools. This means that primordial pools are disk/device sources for allocation of storage pools.

Volume
Detailed information about volume groups and volume entities in the current topology view.

Other/Unknown
Detailed information about objects in your environment other than computers, switches, and subsystems in the current topology view.

FC Port
Detailed information about port groups and port entities in the current topology view.

Connection
Detailed information about all relationships between objects in the current topology view.

Alert Detailed information about all alerts that are associated with entities in the current view of the topology.

Zone Detailed information about any of the following groups or entities in a zone. By default, the **Zone** tab is disabled.

- Fabric groups
- Zone-set groups
- Zone groups
- Computer entities
- Switch entities
- Subsystem entities

Cluster
Detailed information about clusters in the current topology view.

Node Detailed information about nodes in the current topology view.

File Pool
Detailed information about file pools in the current topology view.

Filesystem
Detailed information about file systems in the current topology view.

Network Shared Disk
Detailed information in Network Shared Disks in the current topology view.

Export or Share
Detailed information about exports or shares in the current topology view.

File Set
Detailed information about files sets in the current topology view.

Identifying storage resources using icons:

This topic describes how the topology viewer displays entities and groups of entities with an appropriate icon based on the entity class.

The storage resource icons can help you quickly and easily identify storage resources. For example, an entity of class "computer" is displayed in the graphical view with an icon that resembles a computer.

The following table shows the icons for individual resources and for resource groups.

Topology Object	Icon	Group Icon
Computer		
Computer HBA		
Fabric		
Switch		
Switch port		
Storage subsystem		
Storage subsystem port		
Storage virtualizer (for example, SAN Volume Controller, HDS Tagmastore)		
Database volume		

Topology Object	Icon	Group Icon
File system volume		
Storage volume		
Other volume		
Storage disk		
Storage backend disk		n/a
Storage pool		
Unknown		

Note: *Unknown* entities are entities that are discovered by IBM Tivoli Storage Productivity Center but that do not have an agent installed or the agent is not working.

Identifying icons for storage resources that contain full disk encryption disks

The icons in the following table identify whether a storage system contains encrypted or encryptable disks. The group icon for encrypted storage disks, pools, and volumes is the same as the group icons that are shown in the first table.

For more information on supported systems that provide full disk encryption, see "Supported storage systems providing full disk encryption and solid-state drives" on page 1663.

Topology Object	Encrypted Icon	Encryptable Icon
Storage disk		
Storage pool		
Storage volume		

Identifying icons for storage resources that contain solid state drives

A solid-state drive is a storage device that stores persistent data on solid-state flash memory. The icons in the following table identify whether a storage system

contains solid-state drives. The group icons for storage disk, pool, and volume in the following table are the same as the group icons that are shown in the first table.

For more information on supported systems that provide solid-state drives, see "Supported storage systems providing full disk encryption and solid-state drives" on page 1663.

Topology Object	Icon
Solid-state drive storage disk	
Solid-state drive storage pool	

The topology viewer and virtual SANs:

Cisco switches support virtual SANs (VSANs), which is the logical partitioning of a fabric into multiple fabrics. The overall network is referred to as the physical infrastructure, and the logical fabrics are the VSANs.

SNMP queries to a Cisco switch collect information about the Cisco physical infrastructure and about any VSANs in that infrastructure. Storage Resource agents can also collect information about the VSANs but are unable to collect information about the physical infrastructure. The status of Cisco physical infrastructures that are shown in the topology view is determined based on the status that is propagated from the VSANs in the infrastructure and the status of the infrastructure itself. If the status of the physical infrastructure is missing but the status of all or some of the VSANs in the infrastructure is normal, infrastructure information can no longer be collected. Typically, it indicated indicates that the out-of-band agent is no longer working properly or that SNMP queries to the Cisco switches are no longer configured because the switch was removed from the list of agents in the **Out of Band Fabric** node.

The consequences of this situation are illustrated in the following scenario: A fabric includes a Cisco switch with a VSAN on it. An out-of-band agent on the switch collects physical infrastructure and VSAN information. If the out-of-band agent is removed or Fabric Manager is no longer able to communicate with the agent, IBM Tivoli Storage Productivity Center still detects the VSAN through Storage Resource agents, but physical infrastructure information can no longer be collected. As a result, in the fabric topology view, the status of the physical infrastructure box is shown as missing, while the status of one or more VSANs in the infrastructure shows as normal. To correct this situation, ensure that the out-of-band agent is configured for the Cisco switch.

The topology viewer and VMware:

This topic describes what the topology viewer displays for VMware support.

The Overview view of the topology viewer displays a count of hypervisors and virtual machines in the Computers summary.

The L0: Computers view displays the hypervisors and virtual machines in addition to other computers and includes the grouping By Virtualizer.

The L1: Computer view for a hypervisor displays volumes that were assigned to the hypervisor and the connectivity to fabrics.

The L1: Computer view for a virtual machine displays volumes that were assigned to the hypervisor and have been mapped to the VMware computer (virtual machine) as raw disks.

The L2: Computer view for a hypervisor includes a mapping of virtual machines and shows the virtual machine disk to physical disk relationship.

The L2: Computer view for a virtual machine displays the physical disk (what the virtual machine recognizes as a physical disk) to virtual machine disk relationship and hypervisor connectivity.

Note: IBM Tivoli Storage Productivity Center supports the mapping of storage between the storage system volume and ESX Server corresponding disk. This mapping is supported for the ESX Server 3.5 or VirtualCenter 2.5 or later. Previous versions of ESX Server or VirtualCenter are not supported for the mapping of storage to the storage system.

Viewing large environments in the Topology Viewer:

In environments where you are monitoring a large number of storage entities, IBM Tivoli Storage Productivity Center might not be able to display all of those entities within the Topology Viewer. To avoid this issue, there are some tasks you can perform to improve overall scalability and performance of the user interface.

If the Topology Viewer is unable to display all the entities in your environment, the user interface generates javacore and heapdump files in its installation directory. It is recommended that you check this directory to determine if this display problem is occurring in your environment. To do this, perform the following steps:

1. View the contents of the installation directory where Tivoli Storage Productivity Center's user interface is installed. The default installation directories are as follows:
   - Windows: C:\Program Files\IBM\TPC\gui\
   - UNIX/Linux: /<usr or opt>/IBM/TPC/gui/
2. Determine if javacore and heapdump files appear in the installation directory. You can identify these files by searching for files that begin with the word **javacore** or **heapdump** and have a **.txt** or **.phd** extension. For example: javacore.20080605.151906.8088.txt and heapdump.20080605.151903.8088.phd.

If these files appear in the installation directory, it might indicate that the display problem is occurring when you use the Topology Viewer.

There are two methods for working around this issue:

- Reducing the number of open topology views and disable unneeded overlays (such as the alert, performance, and zoning overlays). The number of topology views that are open and overlays that are enabled can affect how many storage entities are able to be displayed in the Topology Viewer.
- Increasing the memory allocation for IBM Tivoli Storage Productivity Center's user interface component. If reducing the number of open topology views and disabling unneeded overlays does not correct the problem, you can increase the amount of memory that is allocated to the user interface to help resolve the issue.

Reducing the number of open topology views and disabling unneeded overlays:

Perform these tasks to help improve the scalability of the Topology Viewer.

The number of topology views that are open can affect how many storage entities are able to be displayed within those views. The more views (tabs) that are open in the Topology Viewer, the more memory is used by the user interface. If the user interface requires more memory than is available to display all the entities within the open views of the Topology Viewer, some of the entities are not displayed. To workaround this problem, close any views that you no longer need to view and refresh the content of the views that you leave open.

The number of overlays that are enabled in the Topology Viewer can also affect how many storage entities are able to be displayed. The topology viewer uses color and icons to indicate the status of objects in the topology and tabular views. Overlays are colors and icons that used to indicate health, performance, change, and alerts. To disable overlays that you do not need, complete the following steps:

1. Right-click on any blank space within a topology view.
2. Select **Global Settings...** from the pop-up menu.
3. Uncheck the overlays you do not need to view. For example, clear the **Performance** and **Alert** check boxes.
4. Click **OK** to accept the changes. The content of the open views within the Topology Viewer are automatically refreshed according to your changes.

Increasing the memory allocation for the user interface:

Perform this task to allocate more system memory to the IBM Tivoli Storage Productivity Center user interface.

Increasing the amount of allocated memory might enable the user interface to correctly display a large number of storage entities in the Topology Viewer. This task is generally recommended and encouraged in larger environments.

To do this, perform the following steps:

1. Close any current sessions of the Tivoli Storage Productivity Center user interface.
2. Locate the batch or script file that is used to start the user interface. This file is located in the Tivoli Storage Productivity Center installation directory on the computer where the user interface is installed. The default file names and installation directories are as follows:
   - Windows: C:\Program Files\IBM\TPC\gui\tpc.bat
   - UNIX/Linux: /<usr or opt>/IBM/TPC/gui/TPCD.sh
3. Open the tpc.bat or TPCD.sh file in a text editor and locate both occurrences of the parameter that defines how much memory is allocated to the GUI component:
   - Windows parameter: **-mx256M**. The default setting for Windows is 256 megabytes.
   - UNIX/Linux parameter: **-xmx512M**. The default setting for UNIX and Linux is 512 megabytes.
4. Increase the value for all occurrences of this parameter. For example:
   - Windows: change **-mx256M** to **-mx1024M**
   - UNIX/LINUX: change **-xmx512M** to **-xmx1024M**

Note: Make sure that the computer on which the user interface is installed has enough memory to support the setting that you define.

5. Save the batch or script file and restart the user interface.
6. Check the Topology Viewer to ensure that all the monitored storage entities are displayed as a result of increasing the amount of memory allocated to the user interface. If you still experience this display problem, repeat steps 1-5 and further increase the value for the parameter in the startup file.

Contact an IBM service representative if changing the memory allocation for the user interface does not help resolve the issue.

Arranging entities in a group

This topic describes how to arrange entities in an expanded group.

You can arrange entities within the group box by health status or by entity label if an expanded group contains many entities. With either type of arrangement, you can specify ascending or descending order. To arrange entities, follow this procedure:

1. Right-click the graphical view and select **Arrange By**.
2. Select one of the following:
   - **Label (ascending)**
   - **Label (descending)**
   - **Health (ascending)**
   - **Health (descending)**

Checking the status of entities and groups

This topic describes how the topology viewer uses icons to indicate the health, performance, and operational status for entities.

To quickly check the health, performance, operational, change, and alert status of entities or groups of entities, look for the following icons associated with a particular entity or group in the topology. Position the cursor over these icons to see the health string in text format:

Table 77. Topology status indicators and their icons

Indicator	Icon
Health - Normal	
Health - Warning	
Health - Critical	
Health - Unknown	
Health - Missing	
Operational - Running	
Operational - Upgrading	
Operational - Missing	

Table 77. Topology status indicators and their icons (continued)

Indicator	Icon
Performance - Normal	
Performance - Unknown	
Performance - Warning	
Performance - Critical	
Performance - Missing	
Change - Objects did not exist at the two specified time periods	Objects in light gray with no icon
Change - Objects changed between the two specified time periods	
Change - Objects did not change between the two specified time periods	
Change - Objects were created between the two specified time periods	
Change - Objects were deleted between the two specified time periods	
Alert - Critical, Warning, or Informational	

The operational status icons (running, upgrading, missing) are displayed by default and cannot be hidden. However, you can specify whether to display health, performance, or alert icons. To display these icons, follow this procedure:

1. Right-click a blank space in the background of the graphical panel and select **Global Settings**.
2. In the Settings window, select **Health**, **Performance**, or **Alert**. Click **OK**.

Health status:

This topic describes how the topology viewer displays the health status of entities.

The health of an entity falls into one of the following categories:

- Normal: The entity is operating normally.
- Warning: At least one part of the entity is not operating or has serious problems.
- Critical: The entity is either not operating or has serious problems.
- Unknown: The entity was discovered but is not recognized by IBM Tivoli Storage Productivity Center.
- Missing: The entity was previously discovered by Tivoli Storage Productivity Center, but was not detected in the most recent discovery.

The health of a group falls into one of the following four categories:

- Normal: All the entities in the group are operating normally.
- Warning: At least one entity in the group has normal status.
- Critical: All entities are either in critical or warning states.

- Unknown: None of the entities in the group have a known health status.

Health indicators for some entities, such as ports, are only normal and critical. In certain situations, the health of an entity might not be available. This happens, for example, when the entity is identified as missing. In the graphical view, missing entities are shown with a semitransparent light gray layer over an existing icon.

Level-indicator buttons on the graphical view (for example, **L0:Computers**) display icons for health.

Multiple states from single entities are aggregated into group states or states of entities, using the following rules:

- If all entities are in the same state, the aggregated state will represent this state.
- You will get an aggregated critical state (red) if:
  - One entity is in the missing state and at least one other entity shows a critical state.
  - One entity is in the critical state and at least one other entity shows a missing state.
- You will get an aggregated normal state if:
  - One entity is in the normal state and all others are in normal or undefined state.
  - One entity is in the undefined state and all others are in normal or undefined state.
- All other combinations will aggregate to warning state (yellow).
- States displayed within the tabs of the topology view will be aggregated according to the focus of the specific entity class.
  - For Computers and Storage, only the states of these entities are used for aggregation.
  - For Fabric, only the switch status is used for aggregation.
  - For level 1 and 0, the states of all displayed entities are used for aggregation.

See "How status is determined for individual resources" on page 577 for more information about how the status of individual entities are displayed in the topology viewer.

At the Overview level, the health status indicated by the icon on **Overview** for the graphical view indicates the aggregated status of all the groups inside the graphical view (**Computers**, **Fabrics**, **Switches**, **Storage** and **Other**). At the L0 level, the health status indicated by the icon on the **L0** tab of the graphical view indicates the aggregated status of all the groups inside that particular L0 graphical view (for example, the health status for the L0:Computers view would aggregate the health status of all discovered computers, including those that are missing, computers whose health status is critical, computers whose health status is warning). At the L1 level, the health status indicated by the icon on the **L1** tab of the graphical view indicates the status of the particular group that is the focus of the view. (The status of the group is based on the aggregated status of the entities within the group.) The health status for the group is also indicated in the table view. At the L2 level, the health status indicated by the icon on the **L2** tab of the graphical view is the same as the status of the particular entity that is the focus of the view. The health status for the entity is also indicated in the table view.

Note:

- At the L1:Fabric level the health status indicated by the icon on the **L1:Fabric** tab of the graphical view might not be the same as the overall health of the group, based upon the aggregated status of all the entities within the group.
- At the L2:Subsystem level, the status of the volume group is the aggregated status of all the volumes that comprise the group and not the status of the host that owns those volumes.
- The number of computers that are displayed in the Computers view of the topology viewer might be more than the number of Tivoli Storage Productivity Center agents that are deployed in the environment. These additional nodes are displayed with an "unknown" status. For example, if Tivoli Storage Productivity Center determines that a node in a fabric is an HBA, that node is displayed in the Computers view (rather than in the Other view). This situation occurs because during a probe, Storage Resource agents can identify nodes in the topology viewer as HBAs in computers.

  To update nodes that have an unknown status in the Computers view, deploy an Storage Resource agent on the computer that has the HBA whose WWN is shown in the topology viewer.

How status is determined for individual resources:

The topology viewer shows the status for individual resources and for groups of resources. The status of an individual resource is determined by its consolidated status and the propagated status of its internal resources.

The following statuses are possible for individual resources:

- Normal: The resource is operating normally.
- Warning: At least one part of the resource is not operating or has potential problems.
- Missing: The resource was previously discovered by Tivoli Storage Productivity Center, but was not detected in the most recent discovery.
- Critical: The resource is either not operating or has serious problems.
- Unknown: The resource was discovered but is not recognized by Tivoli Storage Productivity Center.

Tivoli Storage Productivity Center determines the status of a resource by using its consolidated status and the propagated status of its internal components:

- A consolidated status is the status that is reported by the top-level resource for its own hardware. For example, if the cooling fans in a storage system are stopped and the internal temperature is too high, that status is reported to Tivoli Storage Productivity Center.
- A propagated status is the combined statuses of the monitored, internal resources for a top-level resource. For example, if a Warning status is detected on a volume in a storage system, that Warning status is used to determine the overall status of the storage system.

The statuses of the following internal resources are used to calculate the status of top-level resources.

Table 78. Internal resources that are used to determine the status of top-level resources

Top-level resource	Internal resources that are used to determine the status of a top-level resource
Fabric	Blades, and ports on blades Ports that are not blades Switches
Hypervisor	Controllers Ports on controllers Virtual machines
Server	Controllers, and ports on controllers Volume groups Monitored directories
Storage system	Clusters Controllers Disks External disks Modules Network shared disks Nodes Pools, and the internal resources that are associated with pools: • Volumes • RAID arrays • Managed disks Ports
Switch	Blades Logical switches Ports

Internal resources for a top-level resource might have different statuses. Tivoli Storage Productivity Center uses the most severe status of an internal resource when determining the overall status for a top-level resource. For example, in a storage system, a port might have a Critical status, a pool might have a Warning status, and a number of controllers might have an Unknown status. In this case, the overall status for the storage system is Critical because it is the most severe status that was detected on the internal resources.

Statuses for internal resources are propagated in the following way to determine the status of a top-level resource:

- If an internal resource has a Critical status, then the related top-level resource has a Critical status.
- If an internal resource has a Missing status, and no other internal resources have a Critical status, the related top-level resource has a Missing status.
- If an internal resource has a Warning status, and no other internal resources have a Critical or Missing status, the related top-level resource has a Warning status.
- If an internal resource has a Normal status, and no other internal resources have a Critical, Missing, or Warning status, the related top-level resources has a Normal status.
- If an internal resource has an Unknown status, and no other internal resources have a Critical, Missing, Warning, or Normal status, the related top-level resource has an Unknown status.

The following table shows some of the possible combinations of statuses and their resulting, propagated status.

Table 79. Propagation of the statuses for internal resources in the stand-alone GUI

Critical	Missing	Warning	Normal	Unknown	Resulting status for top-level resource
				X	Unknown
			X		Normal
			X	X	Normal
		X			Warning
		X		X	Warning
		X	X	X	Warning
	X				Missing
	X			X	Missing
	X		X	X	Missing
	X	X	X	X	Missing
X					Critical
X				X	Critical
X			X	X	Critical
X		X	X	X	Critical
X	X	X	X	X	Critical

Operational status:

This topic describes how the topology viewer has three categories of operational status for entities.

The operational status of most entities falls into one of the following three categories:

Table 80. Categories of operational status for storage entities

Status	Description
Running	The entity is running.
Upgrading	The entity is upgrading.
Missing	The entity was not discovered during the last refresh cycle.

However, some entities display additional operational status. For example, switches can display a status of **OK**, **Error**, **Stopped**, and other types of status.

Note: When IBM Tivoli Storage Productivity Center no longer detects an entity that was previously displayed in the Topology viewer, the operational status of that entity remains as the last known status before it went missing. This means that the operational status of an entity is not cleared or reset if that entity is not found in subsequent data collection jobs.

Performance status:

This topic describes performance status details collected for entities.

Performance information is collected for entities in the L2:Switch and L2:Subsystem views. You can see the performance indicators for individual switches and subsystems in these two views by turning on the performance feature in the Global Settings window. This feature is not enabled by default.

The performance status of an entity falls into one of the following five categories:

- Normal: The entity is performing at the expected level.
- Warning: At least one part of the entity is not operating at the expected performance level.
- Critical: The entity is not operating at the expected performance level.
- Unknown: The performance of the entity is not known.
- Missing: The entity was not found during the last refresh cycle and the performance status is not available.

The performance status does not get carried up to the L0 and L1 views. It is only available in the L2:Switch and L2:Subsystem views.

Performance status of an entity is not propagated upwards to its parent entity. For example, a SAN Volume Controller can cross several critical or warning stress thresholds (and show as yellow or red in the topology viewer) even though the underlying pools, disks or volumes do not have a degraded performance status.

After performance has been enabled, you will see a performance status indicator for any entity in these views that has performance data available. Position the cursor over a performance icon to see a list of performance metrics for that entity.

To enable the collection of performance information, perform the following steps:

1. Right-click a blank space in the background of the graphical panel and select **Global Settings**.
2. In the Settings window, select **Performance** and click **OK**.

Change status:

This topic describes how the Configuration History topology view displays the status of entities over time.

The Configuration History feature uses the icons and colors of the change overlay to indicate historical changes in a SAN configuration's entities. The same icons and colors are used in the graphical view, table view, and snapshot selection panel. The following table lists the five possible states of an entity and shows the associated icons and colors.

State of an entity	Change overlay icon	Change overlay color indicator
Entity changed between the time that the snapshot was taken and the time that a later snapshot was taken.	Yellow pencil	Blue background
Entity did not change between the time that the snapshot was taken and the time that a later snapshot was taken.	Light gray circle	Dark gray background
Entity was created or added between the time that the snapshot was taken and the time that a later snapshot was taken.	Green cross	Green background
Entity was deleted or removed between the time that the snapshot was taken and the time that a later snapshot was taken.	Red minus sign	Red background
Entity did not exist at the time that the snapshot was taken or at the time that a later snapshot was taken.	Not applicable	Light gray background

Level-indicator buttons in the graphical view (for example, **L0:Computers**) display the change icons.

For groups and other collections, the change overlay provides only aggregated status. That is, the change status for a group of entities is either “changed” or “unchanged.” Details for the group as a whole are not provided. Aggregated status displays in the minimap. In the Configuration History page, the minimap uses the following colors to indicate the aggregated change status of groups:

- Blue (changed): One or more entities in the group changed (an addition or deletion of an entity is considered a change).
- Gray (unchanged): All entities in the group are unchanged.

An object that is moved from one group to another group appears only once in the new group in the graphical view. For example, if the health status of a computer has changed from Normal to Warning, the Configuration History page displays the computer as changed in the Warning health group (and no longer displays the computer in the Normal health group).

Important: If you manually deleted a missing entity from the Tivoli Storage Productivity Center , the configuration history function rediscovers and displays the entity in the configuration history viewer as both Missing and in a Normal

state. For example, Entity A (with ID 12) is missing and you decide to manually delete it by using the Remove From Database option from the pop-up menu for Entity A in the topology viewer or configuration history viewer. Later, the Tivoli Storage Productivity Center rediscovers Entity A and assigns it a new ID (ID 22). The configuration history function treats the same Entity A (with ID 12 and ID 22) as two different entities and displays them in both Missing and Normal states in the configuration history viewer. This scenario can also occur if a missing entity is removed by the Tivoli Storage Productivity Center removed resource retention function and is rediscovered by theTivoli Storage Productivity Center again at a later time.

To display a summary of change for the entire row, hover over the changed row. The summary displays the title of the column. It also displays the old and new value for each cell or tab that changed. For example, if a group contains one new entity but everything else in the group is unchanged, the group as a whole displays as "changed." If you click the group, the one new entity displays as "created" and all other entities display as "unchanged."

The Change column in the table view also indicates changed entities.

Alert status:

This topic describes how the topology viewer displays alert status for all entities in the current view.

The alert status of an entity falls into one of the following categories:

- Warning: At least one part of the entity is not operating or has serious problems.
- Critical: The entity is either not operating or has serious problems.
- Informational: Related data about the entity is supplied to aid you in decision making.

When enabled, in the topology viewer the alert overlay is located next to entity icons, near the pinning indicator. The icon for the alert overlay is a red exclamation mark for the presence of alerts. In the title bar of groups and tabs, the alert icon is similarly located near the entity label. The alert overlay is disabled in the Configuration History page.

To find alerts that are associated with an entity in the table view, click the alert overlay indicator to go to the corresponding alert on the Alert tab. Or, right-click the entity to display a pop-up menu and click **Show Alerts in Tabular View** to go to the Alert tab and any corresponding alerts that are associated with that entity. The corresponding alerts are highlighted in purple. When you select an alert in the Alert tab, the entities that are associated with that alert are highlighted in purple in the graphical view. Note that more than one entity may become highlighted. An alert icon for a collapsed group indicates that alerts exist within the collapsed group.

The alert overlay provides rollup, similar to pinning. If alerts are inside a collapsed group, the number of alerts within that group are aggregated and indicated in the collapsed group. All alerts are aggregated into one count, independent of type or entity involved. Hover over the alert indicator for a tab or group to determine the total number of alerts that are rolled up into that indicator and the time stamp of the newest alert.

In the table view, click the Alert tab to view all alerts associated with entities in the current view of the topology. The Alert tab contains alerts as a hierarchical list. For configuration analysis alerts, the highest level provides information about which policy was violated during a particular analyzer run. The next level indicates the number of times that the policy was violated. This list contains one node for each policy violation. The Alert tab only lists the policy violations that contain at least one of the affected entities that are displayed in the graphical view. The remaining policy violations are filtered. The only exception is the Overview panel where the Alert tab lists all of the policy violations. Click each violation to view a list of affected entities. The parent entities for these affected entities are displayed in a separate column. You can resort the alert list by clicking the headers in the table. For instance, by clicking on the header for the entity the alert table is reorganized into a hierarchical list sorted by entity, then by alert type. If no alert was generated for an entity, the Alert column is blank.

To find an entity associated with an alert in the graphical view, click the alert on the Alert tab. The corresponding entity and any associated entities are highlighted (but not selected) in the graphical view. The highlighting is purple.

For more information about viewing alerts, see “Viewing all alerts” on page 356.

Creating custom groups

You can create groups that are based on criteria that is useful and relevant to you.

Custom groups are visible in both the graphical view and the table view. The grouping criterion is displayed in parentheses in the title bar of the group box.

For example, you can use the topology viewer to group computers by the type of operating system. You can also specify user-defined properties (UDPs) to create subgroups for each operating system-based group of computers. You can also group devices by location. To create these types of groups, use UDPs (for example, a text string of "Lab A"). To set UDPs use the pop-up menu item, **Launch Detail Panel**.

If you specify grouping by UDP, the groups are sorted in alphanumeric order in the graphical view. If the grouping is by health status, groups are sorted in the following order missing, critical, warning, normal, and unknown. You can also specify the arrangement of entities within a group box that is based on health status and label.

To group entities, follow this procedure:

1. Right-click a group and select **Group By**.
2. Select a group-by option. Options differ depending upon the class of the group.

   Single Group
   All entities are included in one group.

   Health
   Groups are based on health status (for example, normal or warning).

   Model The model of the storage subsystems server (for example, DS8000). This option is available for subsystem groups only.

   OS Type
   Groups are based on operating systems. This option is available for computer groups only.

RAID level
Groups are based on RAID level. This option is available only for pool groups at the **L2:Subsystems** level.

Fabric This option is available for switch groups only.

User-defined properties (UDPs)
UDPs are characteristics that you define for an entity. For example, you can create a UDP for several entities that uses the text string `new installation`. All of the entities that the UDP `new installation` are grouped together. The topology viewer provides three UDPs so that you can group resources in three different ways, depending on your needs. UDPs are not available for fabric groups.

Defining properties for grouping entities

This topic describes how User-defined properties (UDPs) are useful for describing specific properties of an entity or a set of entities.

You can also use UDPs to group entities with similar characteristics. The topology viewer lets you define three UDPs for an entity. UDPs are not available for fabric groups. To define one or more UDPs for an entity, follow this procedure:

1. Right-click an entity and select **Launch Detail Panel**.
2. Type a property in the **User Defined Property 1** field (for example, a specific operating system, location of computers, and so on).
3. Optional: Type additional properties in the **User Defined Property 2** and **User Defined Property 3** fields, and then click **Set**.
4. Type a description of the element manager and click **OK**.
5. Click **Launch**.

Displaying data paths from one endpoint in the topology to another

This topic describes how to view connectivity for selected entities in the topology using Data Path Explorer.

Data Path Explorer Overview:

Data Path Explorer displays the connectivity between a host disk or computer and a storage subsystem. The Data Path Explorer also provides SAN Volume Controller support. Using Data Path Explorer, you can follow the connectivity from a computer to SAN Volume Controller, and from a SAN Volume Controller to a storage subsystem.

When you view a single data path, you can monitor performance status and pinpoint weaknesses. You can monitor without having to traverse through the many entities that are typically seen in Topology Viewer.

Data Path Explorer shows only those entities that are part of the selected data path, and only certain types of entities. The types include host, switch, host bus adapter (HBA), port, storage subsystem, and storage volume.

Data Path Explorer functions like the rest of the Topology Viewer, with key variations.

Note: The shortcut to Data Path Explorer is not available in the Configuration History topology view.

Before Using Data Path Explorer

Different amounts of information can be shown by the Data Path Explorer, depending on which agents are deployed throughout the fabric.

To use Data Path Explorer for debugging performance problems, complete the following steps to ensure that information in the view is current.

1. In the Topology Viewer, right-click and select **Global Settings** from the menu.
2. Select the **Performance** check box to enable the performance overlay.
3. Click **OK**.
4. Ensure that I/O is running on the host system.
5. Create a threshold alert. For information about threshold alerts, see "Creating a performance threshold and viewing violations" on page 100.
6. Create and run performance monitors on the switches and storage subsystems that you want to view.
7. Run a discovery for the CIM agents and Storage Resource agents.
8. Run a probe to collect data about the storage subsystems, fabrics, and Storage Resource agents.
9. Run a performance monitor to collect performance data for the storage subsystems and switches that you want to view. For information about how to create and run performance monitors, see "Creating performance monitors" on page 300.

For information about configuring the data path explorer, see the topic about planning for the data path explorer in the *IBM Tivoli Storage Productivity Center Information Center*.

Graphical View

When you open a data path view, three tiles display in the graphical view: Initiator Entities, Fabrics, and Target Entities. These tiles contain all the entities in all the data paths that are related to the selected host or storage entity.

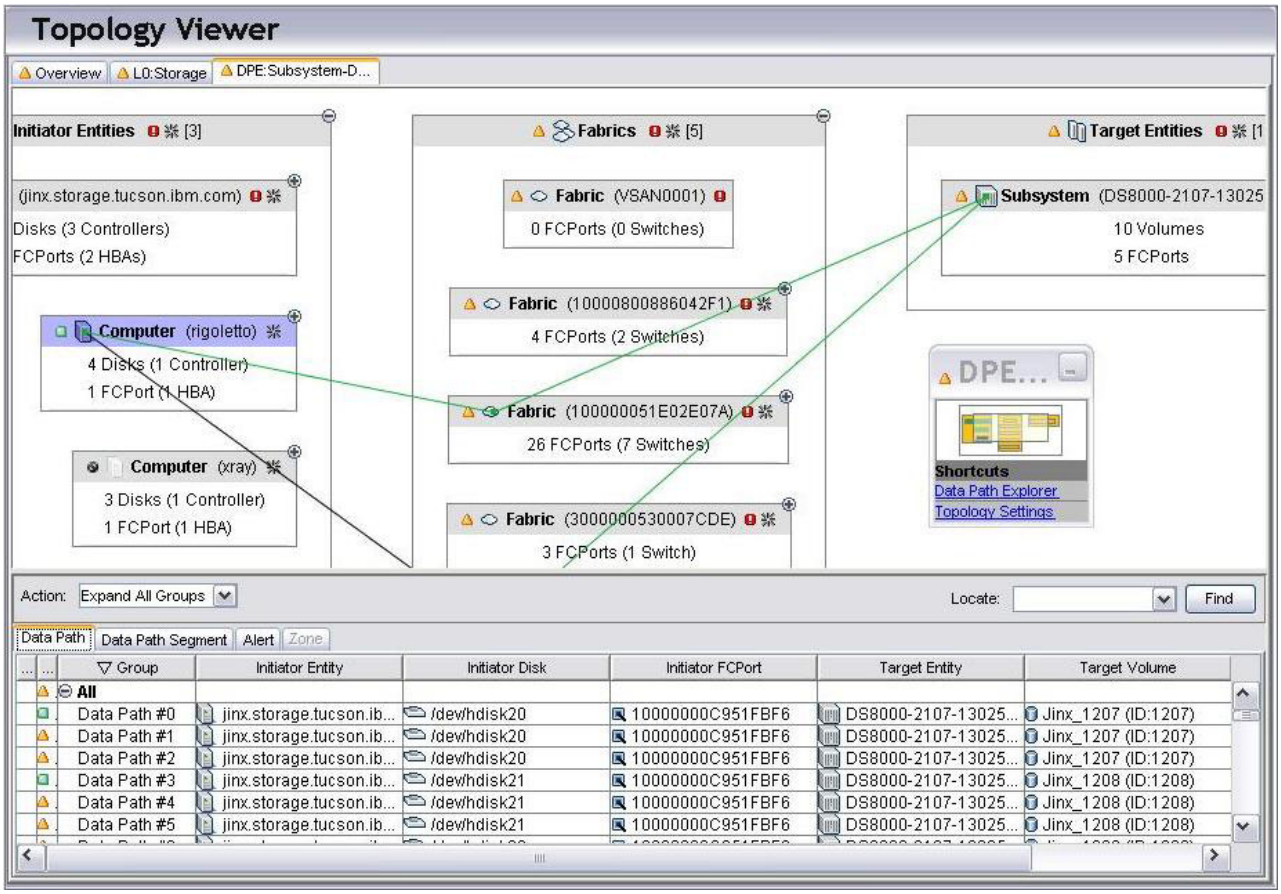

Figure 26. Data Path Explorer view

- Selecting an entity in any of the tile groups at any level displays connector lines. These lines show the data paths that the entity is involved in.
- Each tile is like an active topology viewer tab. It allows grouping, regrouping, expanding, collapsing of groups, and so forth. Each tile can be expanded to view the sub level details such as the disks, ports, and volumes that are involved in a data path.

Initiator Entities

The Initiator Entities tile contains the computers and related volumes, disks, and HBA ports that have valid data paths to the storage system. Expanded computer tiles display each computer, its health status icons, and a count of the entities within the computer that are part of the valid data paths. Expanded computer tiles show the controllers and disks, HBA ports, and volumes on the computer that are part of valid data paths. The initiator entities tile can also contain SAN Volume Controllers. A SAN Volume Controller is a storage virtualization device and can act as an initiator entity. Expanded SAN Volume Controller tiles show the SAN Volume Controller managed disks and ports that are part of valid data paths.

Fabrics

The Fabrics tile contains the fabrics entities, mainly switches and ports, through which valid data paths for the selected endpoint traverse. A count of the switches and switch ports with health status icons and labels are shown in the group title. Expanded Fabric tiles display a view of the

switches to the port level. Only switches and ports with valid data paths that traverse through them are displayed.

Target Entities
The Target Entities tile contains the endpoint storage systems in data paths for the selected endpoint. Expanded Target tiles display each storage subsystem, its status icons, a count of entities, and the number of subsystem front-end ports and volumes that are involved. When these tiles are expanded, the associated port and volume tiles display. These tiles can also be expanded to show individual entities.

Table View

The data path table displays under the graphical view. Each row in the table view shows a valid data path for the selected endpoint.

Display a single data path in the graphical view by selecting a row in the data path table.

The data path table has two tabs: Data Path and Data Path Segment.

Data Path tab
This tab is the default view. It shows each data path that is associated with the entity for which the data path was opened.

- Select an entity in the graphical view to display individual data paths in the table.
- The health status that is shown for each data path is aggregated.
- The data path table can be sorted by any of the column headers.

See “Columns in Data Path tab” on page 589.

Data Path Segment tab
This tab displays the segments in a data path. One segment is listed per row.

There are two ways to populate the Data Path Segment table:

- Select a data path in the Data Path tab. The Segment tab then lists the segments in the selected data path.
- Click an entity in the graphical view. The Segment tab lists the segments in all of the data paths that the entity is involved in. These segments are identified by data path number (0, 1, 4, and so on).

See “Columns in Data Path Segment tab” on page 590.

Related reference:

“Open a Data Path View” on page 588
Data Path Explorer provides an end-to-end topology view of the storage environment from a host disk to a storage volume. It displays the different data paths (SAN access paths or I/O paths) for each host disk.

“Displaying an entity's data paths” on page 589
You can view the data paths associated with any entity by selecting the entity in the graphical view. Lines connect the related entities in the data path or paths.

“Example use cases for Data Path Explorer” on page 591
The following examples show some of the use scenarios for Data Path Explorer:

Open a Data Path View:

Data Path Explorer provides an end-to-end topology view of the storage environment from a host disk to a storage volume. It displays the different data paths (SAN access paths or I/O paths) for each host disk.

You can open a data path view on any of the following endpoints:

host The host (computer) must have at least one disk that is assigned from a subsystem.

subsystem
The subsystem must have at least one volume that is assigned to a host.

volume
The volume must be assigned to a computer, and the fabric must be zoned so that the computer can see the volume.

disk The disk must be from a subsystem. The disk can be a host disk or SAN Volume Controller managed disk.

Open a data path explorer view in one of the following ways:

- Right-click an endpoint, and then select Open Data path view.
- Select an endpoint, and then select Open Data path view from the **Action** menu.
- Select an endpoint, and then select the **Data Path Explorer** shortcut in the minimap
- For multiple entities: Select one or more endpoints of the same type (for example, one or more storage volumes), and then select **Open Data Path View**. The data paths display for all of the endpoints selected. If you select a SAN Volume Controller entity, all other entities you select must be SAN Volume Controller entities.

Note:

1. Opening the **Data Path Explorer** view for a subsystem (including a SAN Volume Controller) that has several associated data paths can use extensive system resources. Select volumes or disks instead, and then open the **Data Path Explorer** view.
2. If the endpoint is a SAN Volume Controller, two data path views are opened. The SAN Volume Controller can act as both an initiator and target entity. The first view displays the data paths that are associated with the SAN Volume Controller acting as a storage subsystem or a target entity. The second view displays the data paths that are associated with the SAN Volume Controller acting as an initiator entity.

When a data path view is opened, three tiles display: the Initiator Entities, Fabric, and Target Entities tiles. These tiles contain all the associated entities in the data paths. Select an entity within one of the tiles to see the data paths that are associated with it. Further expanding the tiles might be necessary.

If you select an entity or a group of entities, that do not have a valid data path, a dialog box displays with the following message: "No available Data Path for the selected entities."

The Data Path Table lists each data path (one per row).

To return to a normal topology view, right-click an entity in the graphical or table view and select **Open Detail View** from the pop-up menu.

Displaying an entity's data paths:

You can view the data paths associated with any entity by selecting the entity in the graphical view. Lines connect the related entities in the data path or paths.

- In the table view-Data Path tab, the selected entity's data paths are highlighted.
- The table view-Data Path Segment tab displays the data paths belonging to the selected entity and the individual segments that make up the data path.

Multiple paths

Tivoli Storage Productivity Center shows the multiple paths from the disks to the storage subsystem volumes in the Data Path Explorer view. IBM Subsystem Device Driver (SDD) software allows the operating system to use multiple paths (that is, different initiator and target port pairs) through the SAN for storage access. On the L2 Computer View and the Data Path Explorer view, the disks with multiple paths managed by SDD are grouped in a controller representing the SDD driver.

For a Windows computer with multiple paths managed by SDD to subsystem volumes, a controller representing the SDD driver is displayed, and this controller contains the disks managed by the SDD.

For a non-Windows computer with multiple paths managed by SDD to subsystem volumes, Data Path Explorer displays a controller representing the SDD driver, and this controller contains the vpath disks. A separate disk (for example, hdisk on AIX) is created by the operating system on the non-Windows computer for each path to the storage volume. These disks are grouped by the controller that contains the initiator port present in the path.

The Linux SDD driver is no longer available. Starting with Red Hat Enterprise Linux Version 5 and SUSE Linux Enterprise Server Version 11, only the DM_Multipath is available. The SDD driver is no longer supported on HP, starting with HP-UX 11i v3 with Itanium®. HP-UX 11i v3 provides a native multipathing feature which is supported by Tivoli Storage Productivity Center.

Related reference:

"Columns in Data Path tab"
When an entity in the Data Path Explorer graphical view is selected, its data paths display in the Data Path tab.

"Columns in Data Path Segment tab" on page 590
The Data Path Explorer Segment tab shows each segment in a data path.

Columns in Data Path tab:

When an entity in the Data Path Explorer graphical view is selected, its data paths display in the Data Path tab.

The Data Path tab columns display the following information:

Located

This column is associated with the Locate function in the Topology Viewer. If search text is present in a row, the Locate column displays a magnifier icon.

Health (status)
The aggregated status of the data path.

Group The number of the data path.

Initiator Entity
The name of the computer or SAN Volume Controller in the data path.

Initiator Disk
The computer disk or SAN Volume Controller managed disk in the data path.

Initiator FCPort
The computer Fibre Channel port in the data path. Data in this column comes from the HBA port entity in the Computer tile. If a SAN Volume Controller is the initiator entity, this column contains the SAN Volume Controller Fibre Channel port in the data path.

Target Entity
The storage subsystem in the data path.

Target Volume
The storage volume in the data path. An asterisk indicates that this volume is a space-efficient volume or virtual disk (VDisk).

Target FCPort
The storage Fibre Channel port in the data path.

Columns in Data Path Segment tab:

The Data Path Explorer Segment tab shows each segment in a data path.

Each segment has two endpoints: the beginning device and the ending device connected by a communication link. Each endpoint resides in a parent entity, which can be a computer, fabric, or subsystem.

The Data Path Segment tab columns display the following information:

Located
This column is associated with the Locate function in the Topology Viewer. If search text is present in a row, the Locate column displays a magnifier icon.

Segment Health
The health status of the segment is an aggregation of From Endpoint Health, To Endpoint Health, From Entity Health, To Entity Health, and Connection Health.

Label The number of the data path and its segments. The data path number corresponds to the numbering on the Data Path tab. Segment numbering advances from computer tile through fabric tile to subsystem tile.

From Entity Health
The health status of the parent entity beginning the segment.

From Entity Performance
The performance status of the parent entity that begins the segment. This column is displayed only when the performance overlay is enabled.

From Entity
The name of the computer, fabric, or subsystem that is the parent of the beginning device (endpoint) in the segment.

From Endpoint Health
The health status of the endpoint device that begins the segment.

From Endpoint Performance
The performance status of the endpoint device that begins the segment. This column is displayed only when the performance overlay is enabled.

From Endpoint
The name of the endpoint device that begins the segment.

To Entity Health
The health status of the parent entity that ends a segment.

To Entity Performance
The performance status of the parent entity that ends the segment. This column is displayed only when the performance overlay is enabled.

To Entity
The name of the computer, fabric, or subsystem that is the parent of the ending device (endpoint) in the segment.

To Endpoint Health
The health status of the endpoint device that ends the segment.

To Endpoint Performance
The performance status of the endpoint device that ends the segment. This column is displayed only when the performance overlay is enabled.

To Endpoint
The name of the endpoint device that ends the segment.

Connection Health
The health status of the connection between the endpoint device that begins the segment (From Endpoint) and the endpoint device that ends the segment (To Endpoint). The Connection Health is displayed as N/A if the health status data for the connection is not available.

Example use cases for Data Path Explorer:

The following examples show some of the use scenarios for Data Path Explorer:

- An application on a host is running slowly, which causes you to question the health of the I/O path for that application. Are there component level performance problems that might cause the slow application response?
- Disk utilization on the storage subsystem is high, triggering performance alerts. You decide to see which hosts or applications, residing on separate host disks, are causing the problem.
- You suspect that the I/O paths for two applications (on separate host disks) are in conflict with each other. They could be sharing a common component (such as a switch). After viewing the I/O paths for these two applications, you make the required zoning or connectivity change to alleviate the problem.
- I/O path information can be used do load balancing and storage planning. For example, a host adapter port on a host has I/O from three host disks passing through it, but a second port has I/O from ten host disks (and has more I/O). This information can be used to provision a new host disk using the first host adapter port.

Note: Application to host disk mapping is not shown in the Topology Viewer. This mapping must be done outside IBM Tivoli Storage Productivity Center. However, Tivoli Storage Productivity Center can provide mapping from database applications to host disks if configured to do so.

Planning for Data Path Explorer view:

Use the Data Path Explorer view to monitor the health and performance statistics of storage components.

The Data Path Explorer is a type of view in the topology viewer. It combines topology views into a visualization that you can use to follow the connectivity from one end point in the topology view to another. Storage administrators can use this view to debug connectivity and performance problems in the storage environment. Data Path Explorer displays the different data paths (SAN access paths or I/O paths) for each host disk so that you can check the health and performance statistics of the resources along the I/O path.

To use Data Path Explorer for debugging performance problems, complete the following steps to ensure that information in the view is current.

1. In the Topology Viewer, right-click and select **Global Settings** from the menu.
2. Select the **Performance** check box to enable the performance overlay.
3. Click **OK**.
4. Ensure that I/O is running on the host system.
5. Create a threshold alert. For information about threshold alerts, see "Creating a performance threshold and viewing violations" on page 100.
6. Create and run performance monitors on the switches and storage subsystems that you want to view.
7. Run a discovery for the CIM agents and Storage Resource agents.
8. Run a probe to collect data about the storage subsystems, fabrics, and Storage Resource agents.
9. Run a performance monitor to collect performance data for the storage subsystems and switches that you want to view. For information about how to create and run performance monitors, see "Creating performance monitors" on page 300.

Different agents have different roles when they discover information and affect what can be shown in the Topology Viewer. The Data Path Explorer view is also subject to those limitations. Different amounts of information can be shown by the Data Path Explorer view depending on the agents that are deployed throughout the fabric.

To display data paths from a host to a subsystem in this view, it is necessary to install and configure the following agents:

- CIM agent to monitor the storage subsystem
- Storage Resource agent to monitor the host
- Storage Resource agent to monitor the fabric
- CIM agent to collect performance information about the fabric

To show the host disk assignment for a storage subsystem, the Tivoli Storage Productivity Center Data server requires access to the CIM agents (SMI-S agent) that are managing it. For storage subsystems, the data path is defined by the host masking or mapping information that is maintained by the CIM agent.

Analysis on the data path of a SAN Volume Controller requires that one or more CIM agents are managing it and the back-end system.

To show accurate information, you must have the most current information from CIM agents and Storage Resource agents. To ensure that you have the most current information, run probes on all the resources in the data path before you view the Data Path Explorer.

To show physical fabric connectivity, you must have a Storage Resource agent, out-of-band Fabric agent, or CIM agent for monitoring the fabric.

To show performance overlay in the Data Path Explorer, you must first create and run performance monitors on the switches and storage subsystems that you want to view.

To show zone overlay, you must have a Storage Resource agent that is connected to the fabric along the data paths or a CIM agent for the fabric. To view the zone overlay, complete the following steps:

1. In the Topology Viewer, right-click and select **Global Settings**.
2. Select the **Show zone tab**.
3. Click **OK**.

If you want to start the Data Path Explorer view for a host, you must deploy a Storage Resource agent to it.

Complete the following steps to set up the Data Path Explorer for viewing:

1. Install and configure the CIM agents for the storage subsystems and fabric. Contact your vendor for information about installing and configuring the CIM agent for a storage subsystem or fabric.
2. Install and configure a Storage Resource agent or out-of-band Fabric agent to monitor the fabric.
3. Install the Storage Resource agent on the hosts.

Displaying membership in zones and zone sets

This topic describes how the topology viewer allows you to examine zones and zone sets.

Zone information is shown in the **Zone** tab in the table view. The **Zone** tab is disabled by default. Selecting a zone or zone set in the **Zone** tab highlights zone and zone set members in the graphical view. Ports and entities that are members of the selected zone or zone set are indicated in the graphical view by highlighting. You can select highlighted entities to show details (for example, connections).

The entities listed in the **Zone** tab are only those entities that are in the current view. However, zone or zone-set selections apply to all currently open views in the topology viewer. Therefore, switching to another view might indicate other entities in the selected zone.

Note: You cannot use the **Zone** tab to change membership status of entities.

To activate the **Zone** tab, follow this procedure:

1. Right-click in the graphical view and select **Global Settings**.
2. In the Settings window, select **Show zone tab** and click **OK**.

Displaying detailed information about entities

This topic describes how to view detailed information about entities using the topology viewer.

There are two ways to display information about entities:

1. Right-click the entity and select **Launch Detail Panel** to display detailed information about an entity (for example, the operating system of a computer, the used and unused capacity of a storage subsystem)
2. Right-click the entity or entities and select **Reports**. From the dialog displayed, you can choose which report to create. The reports you can choose depend on which entities you selected. Asset reports are available only if you selected a single entity. Select a report and click **Create Report**

Displaying storage resource connections

This topic describes how physical and logical connections among entities in the environment can be displayed using the topology viewer.

Connections are shown in both the graphical and table views. In the graphical view, connections are represented by lines between entities or groups. In the table view, a **Connections** tab lists both ends of a connection.

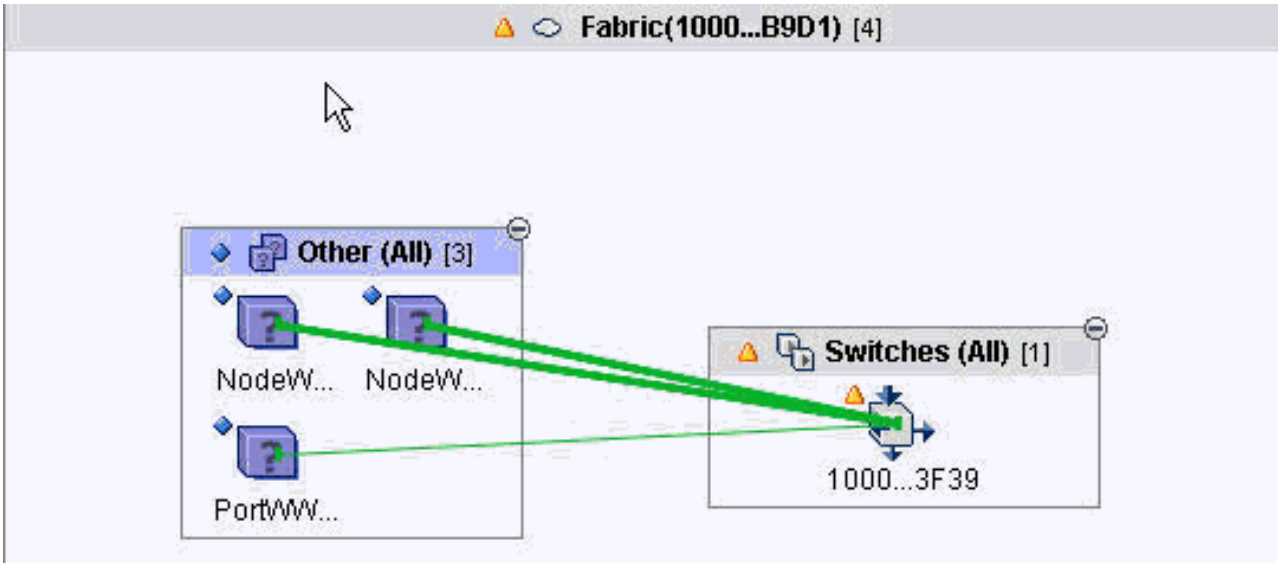

Figure 27. Connections

In the graphical view, the color of a connecting line between entities indicates the status of a connection. The color of the line between groups, individual entities, or between an individual entity and a group can be green (normal), yellow (warning), red (critical), black (missing), or grey (status unavailable).

Note: When you have a host with two HBAs or FC ports that are both in the same switch and you switch the ports between them, the connection will be shown with two yellow (warning) lines. This is because you have not reset the history retention setting. Until the settings are reset, the connecting lines will remain yellow.

Connections between two individual entities are represented by a thin line. Connections between groups or between a group and an individual entity are represented by a thick line. A connection to a collapsed group can represent either a single connection or several connections.

To show connections, follow this procedure:

1. Right-click a blank space in background of the graphical view and select **Global Settings**.

2. In the Settings window, select **Show Connections** and click **OK**.

Filtering entities based on health status

This topic describes how to quickly narrow the health status of normal or non-normal entities in your environment.

In views with many entities and connections, the display can become very cluttered and can make it difficult to determine where problems are occurring. You can filter entities to display only those entities whose health status is non-normal (critical, warning, missing, or unknown).

When you apply filtering, you can also choose to display entities with a health status of normal but that are connected to entities whose health status is either critical or warning. To filter entities by health status:

1. Right-click in the background area of the graphical view and select **Global Settings**.
2. In the Settings window, select **Hide Normal Entities** and, optionally, **Unless connected to non-normal**, and click **OK**.

Labeling entities

This topic describes how you can specify or change label entities.

In both the graphical and table views, entities are identified by labels (for example, the name of a computer). You can specify or change these labels. To specify or change the existing label for an entity, follow this procedure:

1. Right-click the entity and select **Launch Detail Panel**.
2. Type a label for the entity in the **Label** field.
3. Click the save icon in the menu bar to save the label, and then close the window.

Locating storage resources in the topology

This topic describes how you can search for a particular entity using the entity name or a keyword (for example, a host name or operating system).

The open view (for example, L0:Other) will be searched for the string you enter. The results of a locate action are shown by highlighting.

Figure 28. Highlighted entity

To locate an entity, type a search string in the **Locate** field in the table view and click **Find**. The current view (for example, L0:Other) will be searched for the string you enter. The locate function does not search the entire topology, only the currently open view.

When the entity is found, it will be displayed in both the graphical and table views. If the located entity is inside a collapsed group in the topology, you will not be able to see that entity until you expand that group. However, the entity will be

highlighted in the table view. If the entity is in a collapsed group in the table view, that group will automatically be expanded and the highlighted table entry will receive focus. You can then click on that table entry to see the located entity in the topology view.

If you specify a Locate criterion that applies to more than one entity in a view, all the applicable entities are displayed.

Managing Storage Subsystems

This topic describes how to display the Storage Subsystems page, where you can perform storage-subsystem management tasks.

Right-click a storage subsystem and select **Manage Storage Subsystems**.

Pinning entities

This topic describes how pinning entities is an easy way to identify selected entities for future reference or quick access.

Pinned entities display a small flag next to the entity. This flag is propagated to all views that contain that entity. This is useful for marking entities for various reasons, such as a reminder to look at this entity later and for upward propagation. For example, the user may pin a computer in an L0 or L1 Computers view and then quickly wish to see which fabric it belongs to by opening an L0 Fabrics view and looking at which fabric group contains the pinned entity - indicated by the small flag. Pinned objects remain visible as you change views in the topology viewer and when you log off and log back in. Pin lists are stored on a per user basis. This means that different users can have separate pin lists. Pinned objects are indicated by a pin icon in the **Pinned** column of a table or by a flag icon next to the entity in the graphical view.

Figure 29. Pinned entity

To pin an entity, right-click the entity in the graphical view and select **Pin**. To unpin an entity, right-click the entity in the graphical view and select **Unpin**. The same function is available in the **Action** list in the table view.

Pinned entities are visible in groups at all levels and in all views regardless whether the group is expanded or collapsed. You can only pin individual entities; you cannot pin groups (for example, a group of computers).

Printing information from the topology viewer

You can print information displayed in the topology viewer on a single page for reference.

Printing capability is available for the graphical view but is not available in the tabular view. To print the information displayed in the topology viewer, right-click the graphical view and select **Print** or **Print Preview**. The content is scaled to fit onto one page if necessary.

Refreshing views

This topic describes how you can refresh the topology viewer to view the most up-to-date details about your environment.

To provide you with the most up-to-date information, the topology viewer automatically refreshes all open views, including graphics and tables, every five minutes from the IBM Tivoli Storage Productivity Center database. You can also refresh views manually or change the default refresh rate. Opening a new view automatically shows the most current information. To manually refresh views or change the refresh rate, right-click in the graphical view or click the arrow in the **Action** field in the table view and choose one of the following options:

- **Refresh View**: Updates the view that is currently displayed.
- **Refresh All Views**: Updates all views that are currently open. You can also use this option to remove the alert indicators from the devices in the topology view.
- **Refresh Settings**: Opens the **Refresh Rate Settings** window, in which you can select a new automatic refresh rate (specified in minutes).

Removing entities from the database

Computers in the missing state were once in the topology viewer, but were not discovered by IBM Tivoli Storage Productivity Center during the latest refresh cycle. You might want to remove them from the view and the database in order to make it easier to work with other resources.

To remove a computer from the topology viewer (both graphical and tabular views) and the database, you first need to uninstall the agents on that system.

Note: This action is available only for entities that are in the missing state.

1. See *Uninstalling the agents* in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for instructions on uninstalling the agents.
2. In the topology viewer *graphical* view, right-click the computer entity and select **Remove From Database**..
3. In the topology viewer *tabular* view, right-click the computer entity and select **Remove From Database**.

Removing deleted entities

When a discovered entity (fabric, subysystem, switch, or volume) is deleted using another application, such as an element manager or other management application, the next discovery or probe of the topology viewer displays a missing entity. The status of missing entities is aggregated and propagated up through the device hierarchy.

It can be difficult to determine if there are other lower-level entities that have "missing" status, but have not been intentionally removed by a management application. The Removed Resource Retention setting in the topology viewer can be set to a low value, such as 0 or 1 day to remove the missing entity from the display. This removes the missing entities from the display the next time the Removed Resource Retention process runs.

You can also remove deleted entities from the topology viewer using the following procedure:

1. Highlight the intentionally deleted entities in the topology view and select **Remove Resource** in the **Action** list in the tabular view.

Note: Not all entity classes can be removed from the topology using the **Action** list. For those entities, use the Removed Resource Retention setting.

2. A prompt asks you if you are sure you want to remove the resource from the database. Click **Yes**.
3. To refresh the view, click **Refresh View** in the **Action** list. The missing entities will no longer be displayed in the topology viewer.

Selecting entities in graphical and table views

This topic describes how to select entities in the graphical and table views of the topology viewer.

To select an entity in the graphical view, click the entity image. In the table view, click the name of the entity. To make multiple selections, press CTRL and click.

Showing alerts

In the topology viewer, you can show the alert logs for resources such as computers and storage systems.

You can show alert logs for computers, fabrics, storage systems, and switches.

Right-click an entity and select **Show <entity type> Alerts**.

For example:

Right-click a computer and select **Show Computer Alerts**.

Showing the monitoring status of entities

View how Tivoli Storage Productivity Center is monitoring a storage entity that is displayed in the topology viewer.

Display all the data collection schedules, runs, and jobs that are related to a storage entity in the topology viewer.

To view monitoring status, right-click an entity and select **Show Monitoring Status**. For example, right-click a computer and select **Show Monitoring Status**.

When you select this action, the Job Management panel is displayed. Use this panel to view information about how a storage entity is monitored by Tivoli Storage Productivity Center. For storage resource groups, when you click this option the schedules for all the storage entities in that group are displayed.

Related reference:

"Managing schedules, jobs, and runs" on page 605
Use the Job Management panel in the stand-alone GUI to view and manage the schedules, runs, and jobs that are related to the resources that are monitored by Tivoli Storage Productivity Center.

Managing storage systems

You can use the stand-alone GUI to group and view information about storage systems.

Grouping and viewing storage systems

Group and view information about storage systems.

Grouping storage systems:

A storage system group is a user-defined set of storage systems. You might want to group storage systems to make it easier to run monitoring or reporting jobs.

Note: A storage system can belong to only one storage system monitoring group, but can belong to multiple storage system reporting groups.

Creating a storage system group:

This topic describes how to create a storage system group.

1. In the navigation tree pane, expand **Disk Manager** > **Monitoring** > **Groups**.
2. Right-click **Storage Subsystem** and then click **Create Storage Subsystem Group**. In the content pane, the Create Storage Subsystem Group window opens.
3. In the **Description** field, type a short description of the group.
4. In the **Available** field, click the storage systems that you want to add to the group, and then click . The storage systems are added to the **Current Selections** field.
5. Click **File** > **Save**. A window opens requesting a name for the storage system group.
6. Type the name of the storage system group and click **OK**.

Viewing storage system information:

You can view information about the storage systems that are in your environment. You can view general information about all storage systems or you can view detailed information about a specific storage system. This information includes the available and used space for the system.

Viewing general information about a storage subsystem

In the navigation tree, expand **Disk Manager** and click **Storage Subsystems**. On the Storage Subsystems Page, the **Storage Subsystems** table contains the following information about each storage subsystem:

Subsystem

A system ID that is in the following format:

Type-Model-ID-Vendor

Type The type of the storage system, for example, TotalStorage Enterprise Storage Server or SAN Volume Controller.

Model The model number, for example, 2105 or 2145.

ID The serial number or cluster name. Cluster names are used only for storage systems such as SAN Volume Controller and Storwize V7000.

Vendor The manufacturer of the storage system, for example, IBM.

Label A user-defined name for the storage system or the ID of the storage system.

Type The type of storage system.

Model The model of the storage system.

Status The status of the system:

- Normal

- Critical
- Warning
- Unreachable

Available Space (GB)
The total storage capacity that is not formatted in volumes or virtual disks and that is assigned to hosts.

For the following storage systems, this value represents managed disk space that is not allocated to virtual disk space within the managed disk groups and managed disk space that is not defined to a managed disk group:

- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

Tip: For space-efficient volumes, this value is the amount of storage space that is requested for these volumes, not the actual allocated amount. This value can result in discrepancies in the overall storage space reported for a storage system that uses space-efficient volumes. This situation also applies to other space calculations, such as the calculations for the storage system's Consumable (Available) Volume Space and FlashCopy Target Volume Space.

Consumed Space (GB)
The total storage capacity that is formatted in volumes or virtual disks and assigned to hosts.

For the following storage systems, this value represents all space that is allocated from MDisk groups to VDisks (space that is consumed by volumes):

- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

Configured Capacity Limit (GB)
The total amount of real available and formatted storage space in the storage pools that are associated with the storage system.

Remaining Configured Capacity (GB)
The amount of real storage space available to allocate to volumes in the storage pools that are associated with the storage system.

Tiers Indicates the tiers that are used by the storage system. Each tier is delimited by a semi-colon (;). Possible values are: generic_hdd and generic_ssd.

Information about tiers is available only for the following storage systems:

- SAN Volume Controller 6.1 or later
- Storwize V3500
- Storwize V3700

- Storwize V7000
- Storwize V7000 Unified

The value N/A is displayed in tier-related columns for other storage systems.

Tier Capacities (GB)
The total storage capacity of each tier that is used by the storage system. The capacity of each tier is delimited by a semi-colon (;).

Tier Free Capacities (GB)
The free storage capacity (available space) of each tier that is used by the storage system. The free capacity of each tier is delimited by a semi-colon (;).

Viewing detailed storage subsystem information
You can view detailed information about a specific storage system. This information includes the available and used space for the system. To view detailed information about a specific storage system, click the storage system, and then click the magnifying glass icon.

A tab that contains information about the storage system is displayed. The contents of the tab are determined by the type of storage system that you selected.

Storage systems other than SAN Volume Controller, Storwize family, Storwize V7000 Unified, or IBM SONAS
The following list shows all possible information fields that are on the tab. However, the fields that are displayed are determined by the type of storage system that you selected.

Storage Subsystem
The storage system identifier.

Label The user-defined name of the storage system.

Status The status of the storage system:
- Normal
- Critical
- Warning
- Unreachable

Vendor
The manufacturer of the storage system.

Type The type of storage system.

Model The model of the storage system.

Available Space (GB)
The total storage capacity that is not formatted in volumes or virtual disks and assigned to hosts.

Tip: For space-efficient volumes, this value is the amount of storage space that is requested for these volumes, not the actual allocated amount. This value can result in discrepancies in the overall storage space reported for a storage system that uses space-efficient volumes. This situation also applies to other space calculations, such as the calculations for the storage system's Consumable (Available) Volume Space and FlashCopy Target Volume Space.

Consumed Space (GB)
The total storage capacity that is formatted in volumes or virtual disks and assigned to hosts.

Configured Capacity Limit (GB)
The total amount of real available and formatted storage space in the storage pools that are associated with the storage system.

Remaining Configured Capacity (GB)
The amount of real storage space available to allocate to volumes in the storage pools that are associated with the storage system.

Serial Number
The serial number.

Revision
The revision number.

User-defined property 1 (UDP1)
A user-defined property.

User-defined property 2 (UDP2)
A user-defined property.

User-defined property 3 (UDP3)
A user-defined property.

IP Address
The IP address of the element manager.

Element Manager
Identifies the element manager. This value can be a URL, the fully qualified name of the executable file that starts the element manager, or the short name for the system.

Replication Information
Provides replication information for the storage system, if applicable.

SAN Volume Controller, Storwize family, or Storwize V7000 Unified
The following list shows the information fields that are on the tab:

Storage Subsystem
The storage system identifier.

Label The user-defined name of the storage system.

Status The status of the storage system:

- Normal
- Critical
- Warning
- Unreachable

Vendor
The manufacturer of the storage system.

Type The type of storage system.

Model The model of the storage system.

Available Space (GB)
The total storage capacity that is not formatted in volumes or virtual disks and assigned to hosts.

For these types of storage systems, this value represents managed disk space that is not allocated to virtual disk space within the managed disk groups and managed disk space that is not defined to a managed disk group.

Tip: For space-efficient volumes, this value is the amount of storage space that is requested for these volumes, not the actual allocated amount. This value can result in discrepancies in the overall storage space reported for a storage system that uses space-efficient volumes. This situation also applies to other space calculations, such as the calculations for the storage system's Consumable (Available) Volume Space and FlashCopy Target Volume Space.

Consumed Space (GB)
The total storage capacity that is formatted in volumes or virtual disks and assigned to hosts.

For these types of storage systems, this value represents all space that is allocated from MDisk groups to VDisks (space that is consumed by volumes).

Tiers Indicates the tier levels that are used by the storage system. Each tier is displayed as a row in a table.

Tier Indicates a tier level that is used by the storage system, such as generic_hdd or generic_ssd.

Tier Capacity (GB)
Total storage capacity of a tier.

Tier Free Capacity (GB)
Free storage capacity (available space) of a tier.

Serial Number
The serial number.

Revision
The revision number.

User-defined property 1 (UDP1)
A user-defined property.

User-defined property 2 (UDP2)
A user-defined property.

User-defined property 3 (UDP3)
A user-defined property.

IP Address
The IP address of the element manager.

Element Manager
Identifies the element manager. This value can be a URL or the fully qualified name of the executable file that starts the element manager.

Replication Information
Provides replication information for the storage system, if applicable.

Storwize V7000 File Module
The following list shows the information fields that are on the tab:

Available FileSystem Space
The amount of available file system space for the file module.

Consumed FileSystem Space
The amount of used file system space for the file module.

Version
The version of the file module.

IBM SONAS
The following list shows the information fields that are on the tab:

Storage Subsystem
The storage system identifier.

Label The user-defined name of the storage system.

Status The status of the storage system:

- Normal
- Critical
- Warning
- Unreachable

Vendor
The manufacturer of the storage system.

Type The type of storage system: IBM SONAS.

Model The model of the storage system.

Available FileSystem Space
The total amount of unique file system storage space that is not used or available to the computers where a Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe
- File systems that are mounted remotely and are not identified to a network-attached storage server
- File systems that are not mounted

Consumed FileSystem Space
The total amount of unique file system storage space that is used or consumed by the computers where a Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe
- File systems on discovered hypervisors that are assigned to virtual machines where an agent is installed
- File systems on virtual machines where an agent is installed, but the hypervisors for the virtual machines are not discovered and the total used space exceeds the capacity of the total file system capacity
- File systems that are mounted remotely and are not identified to a network-attached storage server
- File systems that are not mounted

Version
The version number.

User-defined property 1 (UDP1)
A user-defined property.

User-defined property 2 (UDP2)
A user-defined property.

User-defined property 3 (UDP3)
A user-defined property.

IP Address
The IP address of the element manager.

Element Manager
Identifies the element manager. This value can be a URL, the fully qualified name of the executable file that starts the element manager, or the short name for the IBM SONAS system.

Replication Information
Provides replication information for the storage system, if applicable.

Managing schedules, jobs, and runs

Use the Job Management panel in the stand-alone GUI to view and manage the schedules, runs, and jobs that are related to the resources that are monitored by Tivoli Storage Productivity Center.

Tip: In Tivoli Storage Productivity Center V5.2 and later, the following data collection jobs are managed in the web-based GUI: probes and performance monitors. Scans, pings, and TPC Server probes are still managed in the stand-alone GUI. For more information about the data collection jobs that you can manage in each GUI, see "Collecting data" on page 295.

Specifically, you can use the Job Management panel to complete the following tasks:

- View information about schedules, runs, and jobs in one central location. In previous versions of the product, this information was available only through nodes under each schedule type in the navigation tree.
- Identify the schedules that include a particular storage resource.
- Identify the schedules, runs, and jobs that are encountering problems when they run.
- Ensure that a storage resource is fully monitored by Tivoli Storage Productivity Center.
- Edit, rename, delete, and run schedules. You must be logged in as the user who created a schedule to complete some of these actions.

Use either of the following methods to access the Job Management panel:

- Click the job management icon in the toolbar.
- Open the definition of a schedule and click **View Job History** to view the runs and jobs that are related to the schedule. For example, in the navigation tree, click **IBM Tivoli Storage Productivity Center** > **Monitoring** > **Scans** > **TPCUser.Default Scan** to view the scan definition panel and click **View Job History**.

The following terms are used on this panel:

Schedule
A schedule is a definition of a monitoring task that collects data about storage resources. Monitoring tasks include scans, pings, quotas, and discoveries. For example, define a ping schedule to collect data about the availability of resources; define a scan schedule to collect file-level information about computers.

Run A run is a single invocation of a schedule.

Job A job is a unit of work within a run for a schedule. One run might have multiple jobs. For example, a scan schedule includes one job per computer for which it collects data.

The Job Management panel is divided into the following sections:

Filters section
Use this section to complete the following tasks:

- Determine the type of resources that are shown on the panel.
- View recommendations for resources that are not fully monitored by Tivoli Storage Productivity Center.

Filtering the contents of this panel is optional. Select **Show All** from the **Entity Type** menu to show all the schedules that are defined in Tivoli Storage Productivity Center.

Schedules section
Use this section to manage the schedules that are related to the resources that you selected in the **Filters** section. You can complete the following actions for each schedule:

- View and edit the definition of a schedule.
- Run a schedule immediately.
- Rename a schedule.
- Delete a schedule.
- Display schedules that have a certain status, such as Success, Running, or Failed.
- View basic information about a schedule, including the ID of the user who created it, description, job type (for example, ping or scan), the last time the schedule was run, and the next time the schedule is run.

Jobs for Selected Schedule section
Use this section to manage each run of a schedule and the jobs that are associated with a run. You can complete the following tasks for jobs and runs:

- View the status of a run and the date and time when it was run.
- View information about the jobs that are associated with each run, including the name of the computer on which a job was run, the type of the monitored resource, and the status of the job.
- View the log files for one or more jobs at a time.
- Display jobs that started or completed within a specified length of time. For example, display jobs that completed in the last two weeks.
- Display jobs that have a certain status, such as Success, Running, or Failed.

Click and drag the horizontal lines on the panel to resize each section. Click the names of columns in the **Schedules** and **Jobs for Selected Schedule** sections to sort the contents of lists based on the clicked column.

Filters

Refresh All
Click this button to update all the lists that are displayed on the panel.

Entity Type
Select the type of resource whose related schedule information you want to view. By default all schedules are displayed when you first access this panel. When you select a specific resource type, the following actions take place:

- The schedules, runs, and jobs that are related to the resource type are displayed in the appropriate sections on the panel.
- The names of the resources that belong to that resource type are displayed in the **Entity Name** list box.
- The text that is displayed for the **Show Recommendations** button indicates if the resources are fully monitored.

You can select from the following entity types in the **Entity Type** menu:

Show All
(default) View all the schedules that are defined in Tivoli Storage Productivity Center.

Storage Subsystem
View the schedules for storage subsystems that are monitored by Tivoli Storage Productivity Center.

Fabric View the schedules for fabrics that are monitored by Tivoli Storage Productivity Center.

Computer
View the schedules for computers that are monitored by Tivoli Storage Productivity Center.

Switch
View the schedules for switches that are monitored by Tivoli Storage Productivity Center.

Monitoring Group
View the schedules that contain monitoring groups.

Storage Resource Group
View the schedules for all the resources that are included in a storage resource group.

Entity Name Filter
Enter text to filter the names of resources that are displayed in the **Entity Name** list box. You can enter a combination of letters, numbers, dashes, underscores, and spaces. The filter text is not case-sensitive. You can also enter the following wildcard characters:

- Use an asterisk (*) to represent zero or more unknown characters of any value. For example, enter **DS8*** to display resources of the selected type that begin with **DS** in their names, such as DS6000-A12-IBM, DS8000-B34-IBM, and DS5000-B35-IBM.
- Use a question mark to represent a single unknown character of any value. For example, enter **D?000-B3?-IBM** to include resources such as DS8000-B34-IBM, DS8000-B36-IBM, and DS5000-B35-IBM in the **Entity Name** list box.

Text that is displayed for the Show Recommendations button
Warning text is displayed for the **Show Recommendations** button if the resources selected in the **Entity Name** list box are not fully monitored.

A resource is considered not fully monitored if it is not included in all data collection schedules that are available in the licensed edition of Tivoli Storage Productivity Center that is installed.

If a resource is considered not fully monitored, click **Show Recommendations** to view a list of recommended actions you can take to address the warning.

Show Recommendations
Click this button to view recommended actions for a selected resource that is not fully monitored. From the Recommendations panel, you can perform the following tasks:

- View the schedules that are recommended for a resource.
- Click the **Take Action** button to create the recommended schedule for a resource.

Schedules

View Detail
View the definition of a schedule that is selected in the schedule list. The schedule definition panel opens when you click this button. Use this panel to view and edit the definition of a schedule.

Run Now
Run the selected schedule immediately. This option is only available if the job is not currently running.

Refresh
Update the status of the schedule that is selected in the schedule list. The process to refresh the status of a schedule might take a few moments depending on the current server load and network speed.

action drop down menu
Use the options in this menu to complete actions against the selected schedule. The menu includes **View Detail**, **Run Now**, and the following actions:

Rename
Rename a schedule. You must be the owner of a schedule or the Tivoli Storage Productivity Center superuser to rename a schedule. This action is not available for the default schedules that are owned by TPCUser.

Delete Delete a schedule. When you delete a schedule, all the runs and jobs that are associated with the schedule are also deleted. You must be the owner of a schedule or the Tivoli Storage Productivity Center superuser to delete a schedule. This action is not available for the default schedules that are owned by TPCUser.

Only show schedules with status
Filter the schedules that are shown in the list according to their statuses. For example, use this menu to indicate that only schedules with a status of **Running** are displayed in the schedules list. The default setting displays all schedules.

You can select the one of the following statuses to filter schedules:

- Show All

- Show All Unsuccessful
- Success
- Running
- Warnings
- Failed
- The job is waiting
- The job is running

Tivoli Storage Productivity Center automatically updates the schedules list if you change the value that is selected in this menu.

schedules list

Use this list to view the schedules for the resources that are selected in the **Filters** section of the panel. The schedules that are displayed must also meet the filter selected in the **Only show schedules with status** menu. Click the name of a column to sort the contents of this list based on that column.

Some jobs that run in Tivoli Storage Productivity Center do not have a schedule that is associated with them. For example, jobs that are run through the command-line interface (CLI) do not have a schedule that is defined for them. To ensure that these job types are displayed on the Job Management panel, a new schedule called **CLI and Event Driven Jobs** was created. Select this schedule to view the runs and jobs that do not have user-created schedules that are associated with them. The job log files for the CLI and Event Driven Jobs schedule are automatically retained for 90 days and then deleted.

The following information is displayed for each schedule:

Creator
The ID of the Tivoli Storage Productivity Center user who created the schedule.

Name The name of a schedule.

Description
The description of a schedule.

Job Type
The type of the job. The following job types are shown:
- ping
- scan
- aggregator
- removed resource retention
- CLI and event-driven
- quota

Last Run Finished
The data and time when the most recent run of a schedule completed. This column is blank if the schedule was never run.

Last Run Status
The status of the most recent run of a schedule.

Next Run
The date and time when the next run of the schedule is planned. The value **On Demand** is displayed for schedules that are defined

with the **Run Now** option. The value **Never** is displayed for schedules that are defined with the **Run Once** option and have already run.

row number
Appears after the schedules list and indicates the number of rows in the schedules list. One row is displayed for each schedule. Use the horizontal lines on the panel to resize the section and increase or decrease the number of rows that are displayed.

Jobs for Selected Schedule

View Log File(s)
View the log files of jobs that are selected in the runs list. Use the Shift + click and Ctrl + click keyboard shortcuts to select multiple jobs. You can view up to 10 logs at the same time. The Log File Viewer panel opens when you click this button or when you double-click a row for a job. Use this panel to view all the processing performed by the Tivoli Storage Productivity Center server and agents for jobs.

Note: The number of job log files that are retained for a schedule is determined by the values that are defined on the Log File and Cached Batch Report Retention panel under **Administrative Services** > **Configuration**. However, for CLI and Event Driven Jobs schedules, job log files are automatically retained for 90 days.

Show Jobs, in the last
Use these fields to determine the runs that are displayed for the selected schedule. Use the following conditions in the filter:

- Select **Started** or **Completed** from the **Show Jobs** menu to display runs based on whether they started or completed in the time frame that is specified in the **last** field.
- Specify a time frame for started or completed runs by entering a number in the **last** field and selecting a unit of measurement from the menu.

For example, select **Started**, enter **4**, and select **weeks** to display the runs for a schedule that were started within the last 4 weeks. Tivoli Storage Productivity Center automatically updates the list of runs if you change the value that is selected in these fields.

Only show jobs with status
Filter the runs that are shown in the list according to their statuses. For example, use this menu to indicate that runs with a status of **Completed with warning(s)** are displayed. The default setting displays all runs for the selected schedule.

You can select the one of the following statuses to filter runs:

- Show All
- Show All Unsuccessful
- The job is waiting
- Running
- The job is running
- Completed successfully
- Running with warning(s)
- Completed with warning(s)
- Completed with error(s)

- System error occurred
- The agent did not respond (to a request from the server)
- The agent is being upgraded
- The agent is not licensed
- Agent not found
- Agent could not be reached
- The job delay limit was exceeded.
- The agent disconnected
- Host was shutdown or recycled
- Cluster resource group unreachable: no IP
- Scripts are disabled for this agent
- Script was not found or has an invalid extension
- The agent disk space is low

Tivoli Storage Productivity Center automatically updates the runs list if you change the value selected in this menu.

runs list

Use this list to view the runs and jobs of the schedule that is selected in the **Schedules** section of the panel. This list displays runs based on the criteria that are specified in the **Show jobs** and **in the last** fields.

Click the **+** symbol next to an entry in the **Run** field to view the jobs that are associated with a run. Click the name of a column to sort the contents of this list on that column.

The following information is displayed for each run or job:

Run Status
The status of a run.

Run The run number that is assigned to a run and the date and time when the run was started.

Computer
(jobs only) The name of the computer where a job is run.

Device
(jobs only) The name of the resource to which a job applies. This column is blank if the job is not related to a specific resource.

Status (jobs only) The status of a job.

Start Time
(jobs only) The date and time when a job started.

Finish Time
(jobs only) The data and time when a job completed.

Log File Name
(jobs only) The name of the log file for a job. Click the **View Log File** to view the contents of a job log.

row number

Shows the number of rows in the runs list. One row is displayed for each run. You can expand each run row to view the jobs that are related to that run. Use the horizontal lines on the panel to resize the section and increase or decrease the number of rows that are displayed.

Job Management - Recommendations

View and apply the recommendations for resources that are not fully monitored by Tivoli Storage Productivity Center.

A resource is not fully monitored if it is not included in all the data collection schedules that are available in the licensed edition of Tivoli Storage Productivity Center that is installed.

Recommendation Severity

Indicates whether a recommendation is a warning or only informational.

- A warning message indicates an action that you can take to more fully use Tivoli Storage Productivity Center to monitor a resource. For example, when you use a probe schedule in the Tivoli Storage Productivity Center license to collect data about a storage system, a warning message is displayed. This message recommends that you also include the storage system in a performance monitor schedule. You can then use the data that is collected by performance monitor to identify the bottlenecks in storage usage and performance for the storage subsystem.
- An informational message is a recommendation to upgrade Tivoli Storage Productivity Center to an edition that includes more monitoring functions. For example, when you use probe schedules in Tivoli Storage Productivity Center Select Edition to collect data about a fabric, an informational message is displayed. This message indicates that if you upgrade to the IBM SmartCloud Virtual Storage Center Storage Analytics Engine license, performance data can be collected about the fabric.

For warning messages, select the corresponding row and click **Take Action** to take the recommended action.

Entity Displays the name of a resource to which the recommendation applies.

Recommendation

Indicates the recommended actions for a resource. One row is displayed per recommendation. Click **Take Action** to resolve recommendations that have a warning severity.

If multiple recommendations are displayed for a resource, sometimes completing one of the recommended actions might resolve the other recommendations. If you click **Take Action** for the first recommendation, the second recommendation might also be resolved. Both recommendations are then automatically removed from the **Recommendation** list.

Take Action

Take the recommended action for a warning. When you click the **Take Action** button, the corresponding panel for the action is displayed.

Cancel

Exit the window without taking any actions.

Launch and launch in context

The launch and launch in context features are used to start the Tivoli Storage Productivity Center stand-alone GUI and to start other applications from Tivoli Storage Productivity Center. The terms inbound and outbound are used to differentiate starting Tivoli Storage Productivity Center (inbound) and starting another application from Tivoli Storage Productivity Center (outbound).

Inbound launch in context refers to starting Tivoli Storage Productivity Center at a specific point in the stand-alone GUI. You can use inbound launch in context in either of the following ways:

- Configure an application such as Tivoli Enterprise Portal to start Tivoli Storage Productivity Center at functionally related locations in the stand-alone GUI.
- Start Tivoli Storage Productivity Center from Java Web Start, the command line, or the **Start** menu or Start page in Windows. Use parameters to specify the starting location in the stand-alone GUI.

Outbound launch refers to starting an application at the starting point for that application. Use outbound launch to open the management GUIs of monitored storage systems and switches from the Tivoli Storage Productivity Center web-based GUI.

Inbound launch in context

Start Tivoli Storage Productivity Center at a specified point in the stand-alone GUI.

Prerequisite: You must have View and Modify permissions to use the launch in context feature.

You can use the launch in context feature to perform the following tasks in Tivoli Storage Productivity Center:

- Create a capacity report
- Create a wasted space report

You can run the launch in context feature in the following ways:

- **Use Java Web Start** from a Web browser. See "Launching with Java Web Start."
- **Use the command line on Windows, AIX, Linux, or UNIX** to enter the command and arguments on the command line. See "Launching from the command line" on page 615.
- **Use the Windows Start menu or Start page**. See "Launching from Windows" on page 616.
- **Use the Tivoli Enterprise Portal GUI** to define, edit, and launch commands. See "Launch from Tivoli Enterprise Portal" on page 617.

The launch in context feature has a number of parameters. See "Launch in context parameters" on page 1729, "Launch-in-context task parameters" on page 1731, and "Available windows" on page 1731.

Launching with Java Web Start:

Using Java Web Start, you can use the launch-in-context feature from a Uniform Resource Locator (URL).

The launch-in-context feature requires Java Web Start 6 on the computer that is running the web browser. If the correct release of Java Web Start is not present, you receive an error message when you enter the URL. The message states that you must download the Java Runtime Environment, which includes Java Web Start. Follow the link that is provided in the message to install the required software.

If you are using Internet Protocol Version 6 (IPv6) and are using Java Web Start to start IBM Tivoli Storage Productivity Center on a Windows system, use a system host name such as `tpc332.storage.newyork.xyz.com:9549` in the web browser

address field rather than an IPv6 address such as
`[2001:DB8:0000:1234:0000:0000:5678:ABCD]:9549`.

To use the launch-in-context feature to complete a task or display a window on a remote server from your browser, complete the following steps.

1. Open your browser and type the launch-in-context command.

   URLs cannot include spaces. In the command, a parameter and its value are separated by an equal sign (=). Parameter=value pairs are separated by an ampersand (&).

   The format of a command to view a window on a remote server follows.

   ```
   http://device_server:device_port/ITSRM/jsp/jnlp/
   tpcgui.jsp?consoleClass=com.ibm.usmi.console.lic.tpc.LICConsole&-
   server=server:port&-function=window_name
   ```

   Tip: The Data server port value is 9549. The port value for the Device server is 9551 for https and 9550 for http. If the wrong the Device server port value is used, the return page cannot be interpreted by the browser. If the wrong Data server port value is used, the following error message is returned.

   ```
   Unable to connect to the server
   ```

2. The first time a window or task starts, you receive a security warning that the downloaded JAR files were digitally signed by a certificate from a trusted company. To prevent this message from being displayed in the future, select **Always trust content from this publisher** and click **Run**. If the Tivoli Storage Productivity Center graphical user interface (GUI) was not downloaded on the local computer by using Java Web Start, or the local computer does not have the latest version of the Tivoli Storage Productivity Center GUI, the GUI is downloaded.
3. If you included the user ID and password in the command, the sign-in dialog is bypassed, otherwise log on to the Tivoli Storage Productivity Center GUI. The Tivoli Storage Productivity Center GUI goes to the tree node corresponding to the window name that you specified. You are connected with the user name and password that you entered. Depending on the value you specified for the **-function** parameter, the following actions might occur:
   - The specified window opens
   - For data.filesystem_capacity_by_computer, the Capacity Report opens.
   - For TPC.reports_data_wasted_space, the Wasted Space Report opens.

Restriction

There is a known limitation when you start the **File System Capacity** > **By Computer** report with launch-in-context, you must specify the ID for the computer upon which you want to report. For example:

```
http://tpcserver14:9550/ITSRM/jsp/jnlp/
tpcgui.jsp?consoleClass=com.ibm.usmi.console.lic.tpc.LICConsole&-
server=tpcserver14.storage.xyz.com:9549&-
function=data.filesystem_capacity_by_computer&-
computer_id=tpcserver14
```

The value that you enter for the **computer_id** parameter might not match the ID for that computer that is stored in the Tivoli Storage Productivity Center database repository. You receive a warning message if the ID for the computer that you enter in the URL does not match the ID that is stored in the data repository. Ensure that this report is displayed properly when you run from the launch-in-context facility. Enter an ID for the computer that

matches the ID that is displayed for it through the Topology Viewer or appropriate report in the product user interface.

Use the information in the following links to enter parameters. If any of the parameters are spelled incorrectly, an error message indicates the incorrect parameter.

"Launch in context parameters" on page 1729
You can use the launch in context parameters, definitions, and examples.

"Launch-in-context task parameters" on page 1731
Use the launch-in-context feature task parameters to complete storage tasks.

"Available windows" on page 1731
You can use windows with the inbound launch-in-context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the **-function** window name argument.

Launching from the command line:

Start windows and tasks from the Windows AIX, Linux, or UNIX command line.

The bat file `tpc.bat` is used for Windows operating systems. The `.sh` file `TPCD.sh` is used for AIX, Linux, or UNIX systems. The parameters follow the file name. Parameters are separated by spaces. Parameters are optional and can display in any order. See "Launch in context parameters" on page 1729, "Launch-in-context task parameters" on page 1731, and "Available windows" on page 1731 for the parameters.

Frequently used commands can be saved as unique command files and rerun as needed.

1. In the Command Prompt window, type the command for the action that you want to take, and press **Enter**.

   The action occurs on the remote Data server. The graphical user interface (GUI) window for that action opens on the local computer.
2. Enter parameters if you previously did not enter any parameters. If you do not provide a user name, password, or server name, the IBM Tivoli Storage Productivity Center GUI opens the connection dialog box. If you do not provide these values, the GUI presents an error message box and requests the values. If you specify incorrect parameters, you receive an error message.

The following example starts a task from the Windows command line.

```
tpc.bat -user admin -passwd password1 -server
tivolii122.storage.usca.ibm.com:9549 -function
data.filesystem_capacity_by_computer -computer_id <computerID>
```

The following example uses the launch-in-context feature to view a window on the remote system:

```
tpc.bat -user admin -passwd password1 -server
tivolii122.storage.usca.ibm.com:9549 -function data_report.AvailableFile
System Space
```

Restriction

There is a known limitation when you go to the **File System Capacity** > **By Computer** report with launch-in-context. You must specify the ID for the computer upon which you want to report. For example:

```
http://tpcserver14:9550/ITSRM/jsp/jnlp/
tpcgui.jsp?consoleClass=com.ibm.usmi.console.lic.tpc.LICConsole&-
server=tpcserver14.storage.xyz.com:9549&-
function=data.filesystem_capacity_by_computer&-
computer_id=tpcserver14
```

The value that you enter for the **computer_id** parameter might not match the ID for that computer that is stored in the Tivoli Storage Productivity Center database repository. You receive a warning message if the ID for the computer that you enter in the URL does not match the ID that is stored in the data repository. Ensure that this report is displayed properly when you run from the launch-in-context facility. Enter an ID for the computer that matches the ID that is displayed for it through the Topology Viewer or appropriate report in the product user interface.

Launching from Windows:

Learn how to open windows and start tasks on remote servers with Windows operating systems.

The bat file, `tpc.bat`, is used for Windows operating systems. The parameters follow the bat file name. Parameters are optional and can display in any order. To use the launch-in-context feature in Windows, complete the following steps.

1. Complete the steps that apply to the Windows operating system that you use.

Option	Description
Windows Server 2012	1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**. 2. On the Start page, right-click, and then click the **All apps** taskbar button. 3. Right-click **TPC Stand-alone GUI**, and then click the **Open file location** taskbar button. 4. On the IBM Tivoli Storage Productivity Center window, right click **TPC Stand-alone GUI**, and then select **Properties**.
Windows 7, Windows 2008, Windows Vista, Windows XP	1. Click **Start** > **All Programs** > **IBM Tivoli Storage Productivity Center**, right-click **TPC Stand-alone GUI**, and then select **Properties**.

2. Enter the command line arguments in the **Target** field and click **OK**.

   If the Data server is remote, the action occurs on the Data server and the graphical user interface (GUI) window the action opens is on the local computer.
3. Enter the parameters if you did not previously enter the necessary parameters. If you do not provide the user name, password, or server name, the Tivoli Storage Productivity Center GUI opens the connection dialog box. If you do not provide these values, the GUI presents an error message box and requests the values.

The following example starts a task from the Windows **Start** menu.

```
"C:\Program Files\IBM\TPC\gui\tpc.bat" -user admin -passwd password1
-server tpcserver1.tpc.example.com:9549 -function
data.filesystem_capacity_by_computer -computer_id computerID
```

The following example uses the Windows **Start** menu to view a window on the remote server:

```
"C:\Program Files\IBM\TPC\gui\tpc.bat" -user admin -passwd password1
-server tpcserver1.tpc.example.com:9549 -function disk.storage_subsystem
```

"Launch in context parameters" on page 1729
You can use the launch in context parameters, definitions, and examples.

"Launch-in-context task parameters" on page 1731
Use the launch-in-context feature task parameters to complete storage tasks.

"Available windows" on page 1731
You can use windows with the inbound launch-in-context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the **-function** window name argument.

Launch from Tivoli Enterprise Portal:

This topic shows you how to start an existing launch definition from the Tivoli Enterprise Portal graphical user interface (GUI).

You must configure Tivoli Enterprise Portal before you can use the launch in context feature. For information on prerequisites and details on how to configure the launch in context feature with Tivoli Enterprise Portal, see "Configuring Tivoli Enterprise Portal." After Tivoli Enterprise Portal has been configured, you can start the application from the configuration window or use **Launch** from the Tivoli Enterprise Portal navigation tree.

If you use the Tivoli Enterprise Portal navigation tree, **Launch** will open the Tivoli Enterprise Portal graphical user interface with your predefined launch in context commands. To use one of these commands, complete the following steps:

1. In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems *your_server_name* > Windows OS**. Right-click **Disk** and select **Launch**. Disk is one of several values you can select. Values include Disk, Network, Enterprise Services, Memory, Printer, Process, Processor, and System.
2. The Create or Edit Launch Definitions window opens showing a list of defined launch definitions in the **Existing Launch Definitions** area.
3. Select the launch definition you want to use from the list and click **Launch**. The window or task that is defined by that launch definition opens.

For information on editing the launch in context command with the Tivoli Enterprise Portal GUI, see:

Related tasks:

"Editing launch definitions" on page 619
You can edit an existing launch definition with the Tivoli Enterprise Portal graphical user interface (GUI).

Configuring Tivoli Enterprise Portal:

You can configure Tivoli Enterprise Portal so that you can use the launch in context feature.

To configure Tivoli Enterprise Portal complete the following steps:

1. The initial steps depend on whether you are going to open a window or perform a task.
   a. If you are opening a window, in the Tivoli Enterprise Portal navigation window on the local computer, select where you want the launch definition to originate from. This opens the corresponding view on the target computer. Right-click one of the following locations.
      - The navigator item
      - A row in a table or situation event console view
      - A slice of a pie chart
      - A bar of a bar chart

      If the launch definition originates from a table, chart, or situation event console view, open the workspace containing the view. If you think you might edit the argument later, create the launch definition from a table or chart view.
   b. If you are going to perform a task, create a launch definition for one of the following tasks:
      - Create a volume or VDisk
      - Create a capacity report
      - Create a wasted space report
2. In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems >** ***your_server_name*** **> Windows OS**. Right-click **Disk** and select **Launch**. Disk is one of several values you can select. Values include Disk, Network, Enterprise Services, Memory, Printer, Process, Processor, and System.
3. The Create or Edit Launch Definitions window opens. Click **Create New**.
4. In the **Name** text field, type a descriptive name for this launch definition. This name is used in the definition list.
5. In the **Target** text box click **Browse** to locate the `tpc.bat` file. The Select files for directories for Launch Definitions window opens. A typical location for Windows systems is `C:\Program Files\IBM\TPC\gui\tpc.bat`. Select the correct directory and click **OK**.

   When you start a launch definition from the command line, if any part of the path contains spaces, enclose the entire entry in quotation marks.
6. In the **Arguments** text field, enter any parameters that will be added when the application is started. Arguments are entered as *`-parameter value`*. In the Create or Edit Launch Definitions window, click **Browse** next to **Arguments** to locate the parameters that you want to use. The Select Browser Options window opens. Click **Substitutable Items** and then click **OK**. Select parameters from the **Selected Context** and **Extended Context** lists.

   The **help** argument lists the available **function** values. You must type the help parameter by itself into the **Arguments** text field or into the form `tpc.bat -help functions` on the command line for Windows or `TPCD.sh -help functions` for UNIX. The **-help functions** parameter provides a window that lists the nodes in the navigation tree that are available to start. If you type **-help** without the **functions** parameter, it results in an error message.

   When all parameters have been selected, click **OK** to return to the Create or Edit Launch Definitions window.
7. If the parameters can be changed at launch time (for example, you want to change the system ID), check the **Edit arguments before execution** box.

8. Specify a value for the **Start in** location. Click **Browse** to locate the IBM Tivoli Storage Productivity Center installation and select the folder **gui** where `tpc.bat` is located.
9. Click **Evaluate** to see how the command resolves when the application starts. Correct any of the parameters that are incorrect and reevaluate the command.
10. When you are satisfied with the launch definition, click **Apply** to save your changes. This keeps the window open so that you can start the current launch definition or create a new one. Click **OK** to save the launch definition and close the window.
11. To start the current definition, select **Launch**. The Tivoli Storage Productivity Center graphical user interface (GUI) with the function parameter that you have provided is opened.

Related reference:

"Launch in context parameters" on page 1729
You can use the launch in context parameters, definitions, and examples.

"Launch-in-context task parameters" on page 1731
Use the launch-in-context feature task parameters to complete storage tasks.

"Available windows" on page 1731
You can use windows with the inbound launch-in-context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the **-function** window name argument.

Editing launch definitions:

You can edit an existing launch definition with the Tivoli Enterprise Portal graphical user interface (GUI).

If the **Edit arguments before execution** box was checked when a launch definition was configured, you can edit the launch definition by changing the attribute values. To edit a launch definition, complete the following steps:

1. In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems > *your_server_name* > Windows OS**. Right-click **Disk** and select **Launch**. Disk is one of several nodes you can select. Nodes include Disk, Network, Enterprise Services, Memory, Printer, Process, Processor, and System.
2. The Create or Edit Launch Definitions window opens showing a list of launch definitions in the **Existing Launch Definitions** area.
3. Select the launch definition that you want to edit. The values are displayed.
4. Make the necessary changes in the text fields.
5. Click **Evaluate** to see how the arguments will resolve when the application starts. Correct any incorrect arguments and re-evaluate the command.
6. Click **Apply** to save the changes and keep the window open. Click **OK** to save the launch definition and close the window.

   If you do not want to overwrite the existing launch definition, save the new launch definition with a unique name.

Outbound launch

Resources such as storage systems and switches can have their own management GUIs or element managers. In theTivoli Storage Productivity Center web-based GUI, you can open the start page for these management GUIs.

The action for opening a management GUI is available for storage systems and switches under the following conditions:

- The resource is monitored by Tivoli Storage Productivity Center.
- The resource has a management GUI that is web-based or can be accessed with Java Web Start.

Related tasks:

"Opening the management GUI for a resource" on page 440
Resources such as storage systems and switches can have their own management GUIs or element managers. In Tivoli Storage Productivity Center, you can open the start page in these GUIs from resource list pages, such as the Storage Systems page and Switches page.

Working with invoices

With IBM Tivoli Storage Productivity Center you can develop invoice definitions and create invoices associated with data or resource usage.

Using Tivoli Storage Productivity Center you can:

- Charge users for the amount of *storage they are consuming* across a network and within relational database management system (RDBS) instances
- Charge users for the amount of *space they own* within the database and tablespace of RDBMS instances.

With Tivoli Storage Productivity Center you can collect usage information on a departmental, group, or user level. You can allocate costs by storage usage by user, disk capacity by computer, tablespace, or file system or physical device. You can create cost centers by creating user, computer, or tablespace groups, allowing organizations to chargeback individuals or business units for their storage usage. By understanding the costs associated with existing and future storage usage, you can improve the use of that storage and reduce the costs associated with its maintenances and upgrades.

Tivoli Storage Productivity Center uses the storage usage information gathered by probes and scans to generate invoices that chargeback for storage usage. Chargeback is the process of breaking down costs associated with data or resource usage and then distributing those costs back to the organizational consumers.

To collect the information needed for chargeback, you need to run probes and scans:

Probes
: Define probes when you want to charge by Disk Capacity. Probes collect statistics on the storage assets in your enterprise, such as computers, disk controllers, hard disks, file systems, storage subsystems, instances, logs, RDBMS objects, and so on. The results of probes are stored in the database repository and are used to supply the data necessary for Tivoli Storage Productivity Center's reporting function.

Scans
: Define scans when you want to charge by OS User groups. Scans collect statistics on the usage and trending of your actual storage consumption. The results of scans are stored in the database repository and are used to supply the data necessary for Tivoli Storage Productivity Center's Capacity, Usage, Usage Violations, and Backup Reporting facilities.

In addition to providing basic storage usage invoicing, Tivoli Storage Productivity Center integrates with the chargeback systems already implemented in your environment. It provides you with a higher level, application-specific CIMS output format. CIMS is IBM Tivoli Usage and Accounting Manager (ITUAM) which

provides IT financial knowledge by measuring, analyzing, reporting, and billing based on usage and costs of shared computing resources. With Tivoli Storage Productivity Center, you can export chargeback data directly into ITUAM, giving you the ability to integrate your storage usage data with other enterprise chargeback information and processes.

Planning for chargeback

Before generating data that will be used for chargeback, it is a good idea to plan ahead and gather basic information about your system and organization.

Understand your environment

Understanding the users, data, and resources within your organization will help you determine the most reasonable and enforceable chargeback policies. In the planning stages of developing a chargeback policy, it will help you to understand such things as:

- The types of data stored within your environment
- Patterns of data usage
- Different types of users and user groups
- The IT resources that exist within your organization

Determine the chargeback policy you want to implement

It is important to learn and understand the advantages and disadvantages of the charge back policy you want to implement. For example, with IBM Tivoli Storage Productivity Center, you can:

- Divide the cost of storage usage by all users
- Divide the cost of storage usage by individual users or groups of users
- Charge by resource
- Charge by individual usage of resources

Define monitoring groups

Organize the IT resources within your environment into functional groups. Defining groups enables you to classify your data and resources into billable entities, which in turn provides the information on who is using what storage. This enables you to bill specific groups according to their usage of storage space. See "Creating a group" on page 337.

Computer groups

Computer groups collect operating system and RDBMS storage statistics on computers that are related to each other or belong to a common group such as a department, a location, a territory, and so forth. Examples of computer groups are Sales, Accounting, Development, IT.

A computer can belong to one computer group only.

User groups

User groups are users that are related to each other or belong to a common group such as a department, a location, a territory, and so forth. Examples of user groups are Administrators, Developers, Marketing staff, HR staff.

A user can belong to one user group only.

OS User groups

OS user groups are users that are grouped together within the operating system.

Database-tablespace groups

Database-tablespace groups monitor entire databases or tablespaces

across multiple instances or databases. Examples of tablespace groups are System Tablespaces, System Databases, databases-tablespaces related to a specific application or applications, and so forth.

A database-tablespace can belong to one database-tablespace group only.

Determine what information you want to gather

Determining what information you want to gather about the storage resources within your enterprise is critical to helping you implement a chargeback strategy. Probes and scans are the monitoring jobs you can define to collect information about operating system and RDBMS storage resources. The information gathered by these jobs is reflected in the data included within a Tivoli Storage Productivity Center invoice.

As you become more familiar with how a specific charge back model works for your organization, you can continue to review and implement to more advanced charge back models.

Setting up invoices

This topic takes you through the task of setting up an invoice.

To set up an invoice, complete the following steps:

1. In the Navigation Tree pane, expand **Data Manager for Chargeback** > **Parameter Definition**. Data Manager for Chargeback - Configuration window opens.
2. In the **Output Format** area, specify the format for the output. You can click either of the following options:

   Invoice
   Creates a formatted text file suitable for opening in or importing into word processing applications.

   CIMS Creates a text file with comma-separated fields suitable for importing into a ITUAM application.
3. In the **Type of Charges** area, specify the information that you want the invoice to contain. You can select one or more of the following check boxes:

   OS Storage Usage By User
   The resulting invoice includes the following information:
   - Total space consumed by user groups and the cost of that usage
   - Total space consumed by each user in the environment and the cost charged to each user

   OS Disk Capacity By Computer
   The resulting invoice includes the following information:
   - Total space consumed by computer groups and the cost of that usage
   - Total space consumed on each computer in the environment and the cost of that usage

   Storage Usage By Database User
   The resulting invoice includes the following information:
   - Total space consumed by DB user groups and the cost of that usage
   - Total space consumed by each DB user in the environment and the cost charged to each user

Total Size By Database-Tablespace

The resulting invoice includes the following information:

- Total space consumed by database-tablespace groups and the cost of that usage
- Total space consumed on each database or tablespace in the environment and the cost of that usage

4. (Invoice output format only) For each type of charge that you selected in step 3 on page 622, select the unit of measurement (kilobyte, megabyte, or gigabyte) and type a numerical value. For example, if you want to charge $1.25 per megabyte of storage usage, select **Megabyte** and type 1.25 in the **Cost per** field.

   (CIMS output format only) If you select CIMS as the output type, the **Cost Per** field will be deselected. This is because you will set the rates and costing within the ITUAM application.
5. If you selected **OS Storage Usage by User** or **Storage Usage by Database User**, select the profile that is used in the scan job. By default, this is either TPCUser.Summary By Owner for **OS Storage by User** or TPCUser.Db User Space for **Storage Usage by Database User**.
6. Click to save the invoice definition.

Invoice definitions are not saved in actual files, but the last invoice definition saved is the one that will be used to create the invoice.

Generating an invoice

Generate an invoice and save it as a file.

To generate an invoice, complete the following steps:

1. In the Navigation Tree pane, expand the **Data Manager for Chargeback** node.
2. Click **Online Execution**. The Online Execution page opens.
3. Click the **Browse** button. The Open window opens, listing previously generated invoice files in the default directory. The default directory, or the directory where you have chosen to save the invoice files, is not part of the IBM Tivoli Storage Productivity Center application but on the computer where Tivoli Storage Productivity Center is installed.
   a. Click on an existing file or type a new file name, including extension, in the **File name** text field. A text file (.txt) is the default file type for both output types.

      Note: Unless you save the invoice as a separate file name, the default file (StorageChargeback.txt) will be overwritten each time you run an invoice. It is recommended that you save your invoices as a separate file name each time you run them.
   b. Click the **Open** button. You will be returned to the Online Execution page.

   Alternatively, in the Online Execution page, without clicking the **Browse** button, type the full path name of the invoice file to be generated in the **Output File** text field.
4. Click the **Generate Output File** button to generate the invoice.

The invoice is saved to the file name and path you specified.

Viewing and printing invoices

This topic discusses viewing or printing a generated invoice in both Invoice and CIMS output formats.

If you have generated the invoice in the Invoice output format, you can view and print that invoice with any text editor or word processing program. If you chose CIMS as the output format, use your ITUAM application to work with the invoice data.

To view and print an invoice generated in the Invoice output format, complete the following steps:

1. Start the text editor or word processing program.
2. Follow the steps in the word processing program to open the invoice. Open the invoice from the location you specified in the **Output File** field in the Online Execution page.
3. View the file and make any changes. If you make any changes to the file, be sure and save it under a different name.
4. Print the file from the word processing program.

To view and print an invoice generated in the CIMS output format, integrate IBM Tivoli Storage Productivity Center accounting data into the Tivoli Usage and Accounting Manager, by following the instructions in the Tivoli Usage and Accounting Manager Information Center. Alternatively, you can view the original invoice data in a text editor or word processing program. Remember that the data will be in the form of comma-separated fields.

Sample Invoices

This section gives examples of sample invoices.

OS Storage Usage By User:

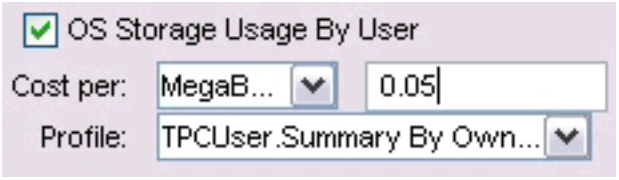

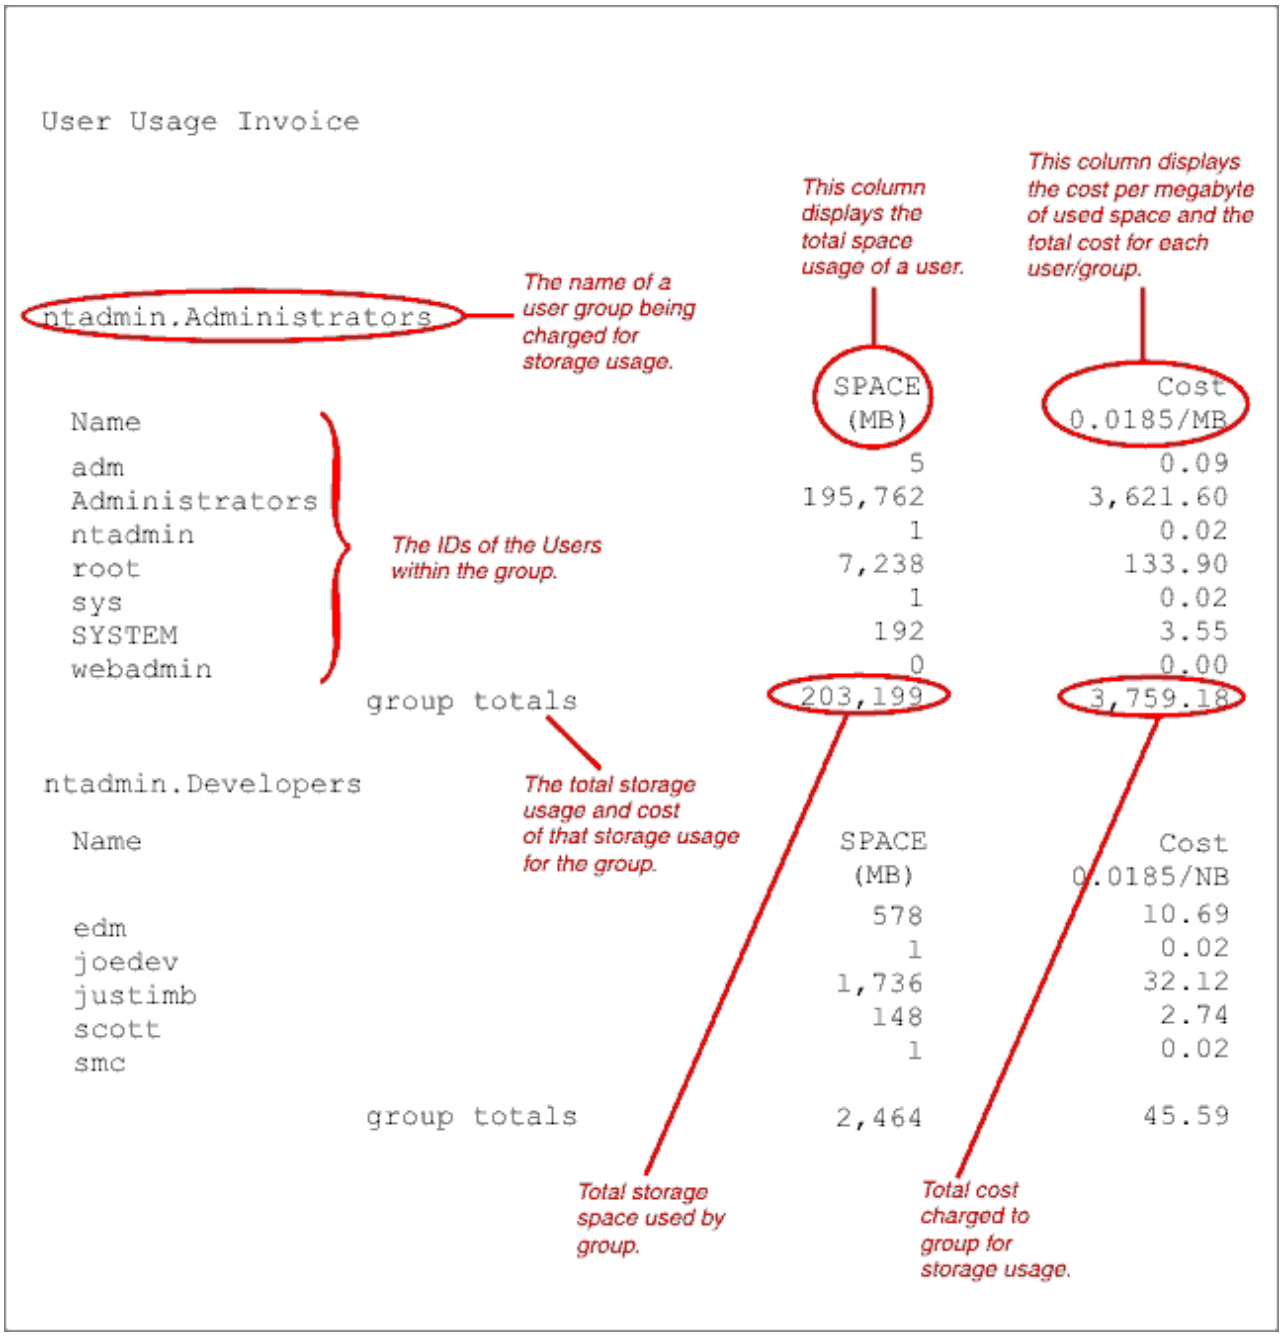

Figure 30. Sample Invoice — OS Storage Usage By User

OS Disk Capacity By Computer:

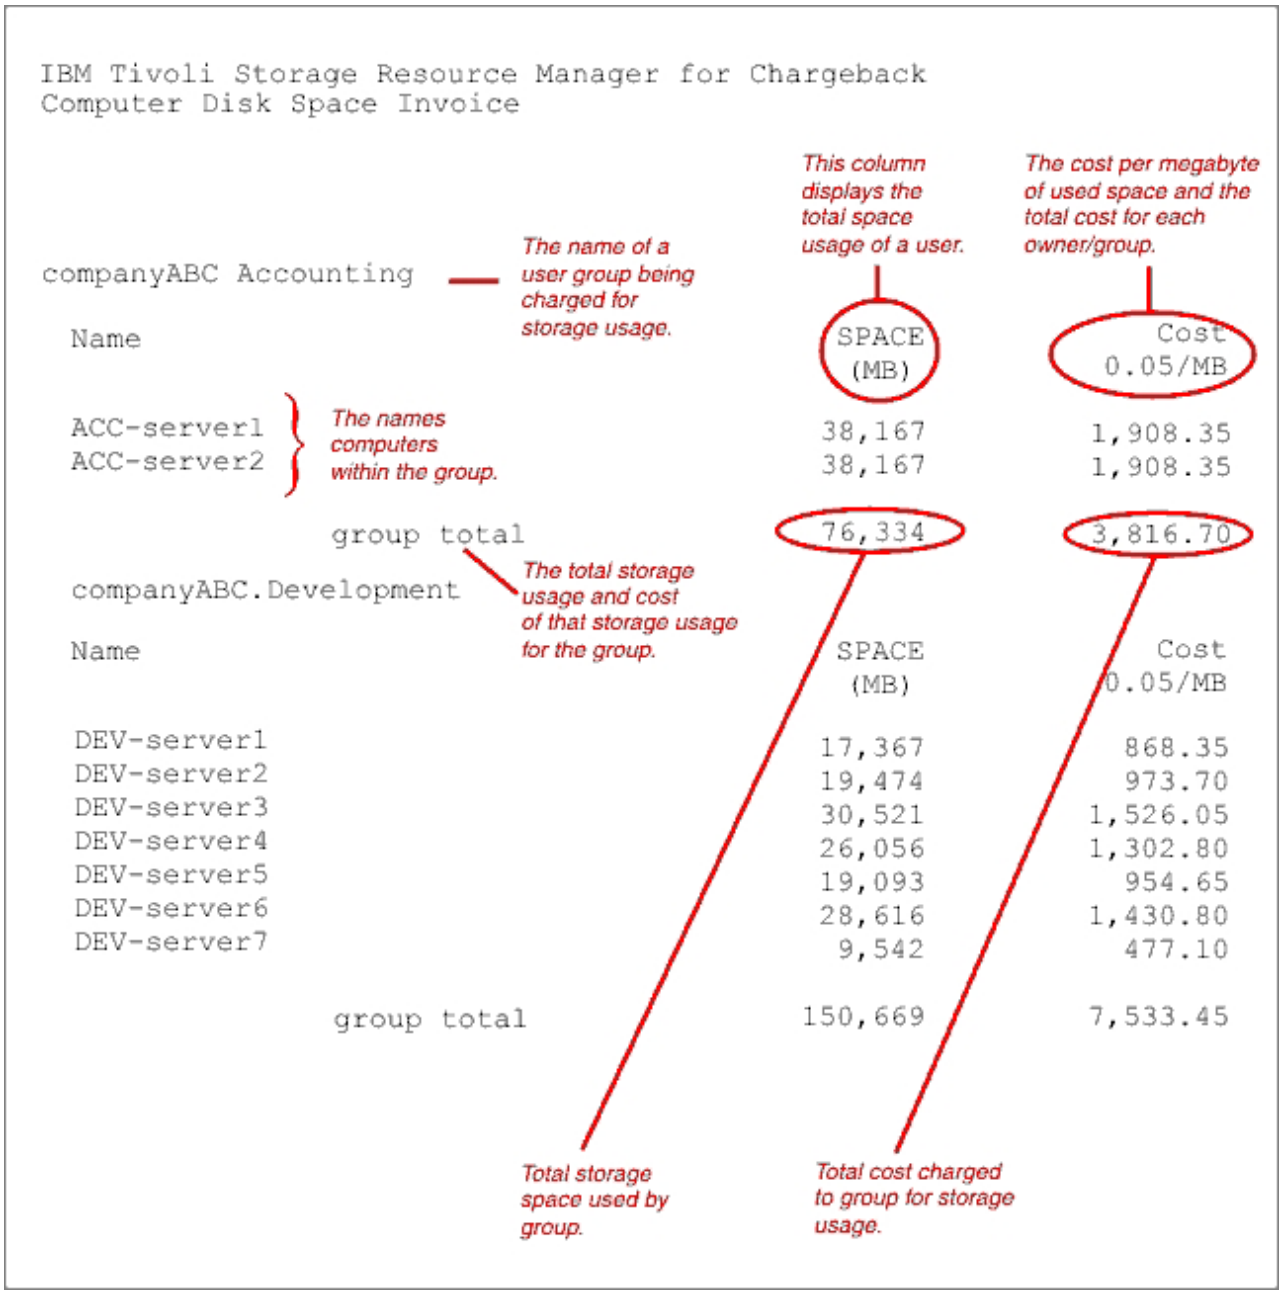

Figure 31. Sample Invoice — OS Disk Capacity by Computer

Storage Usage By Database User:

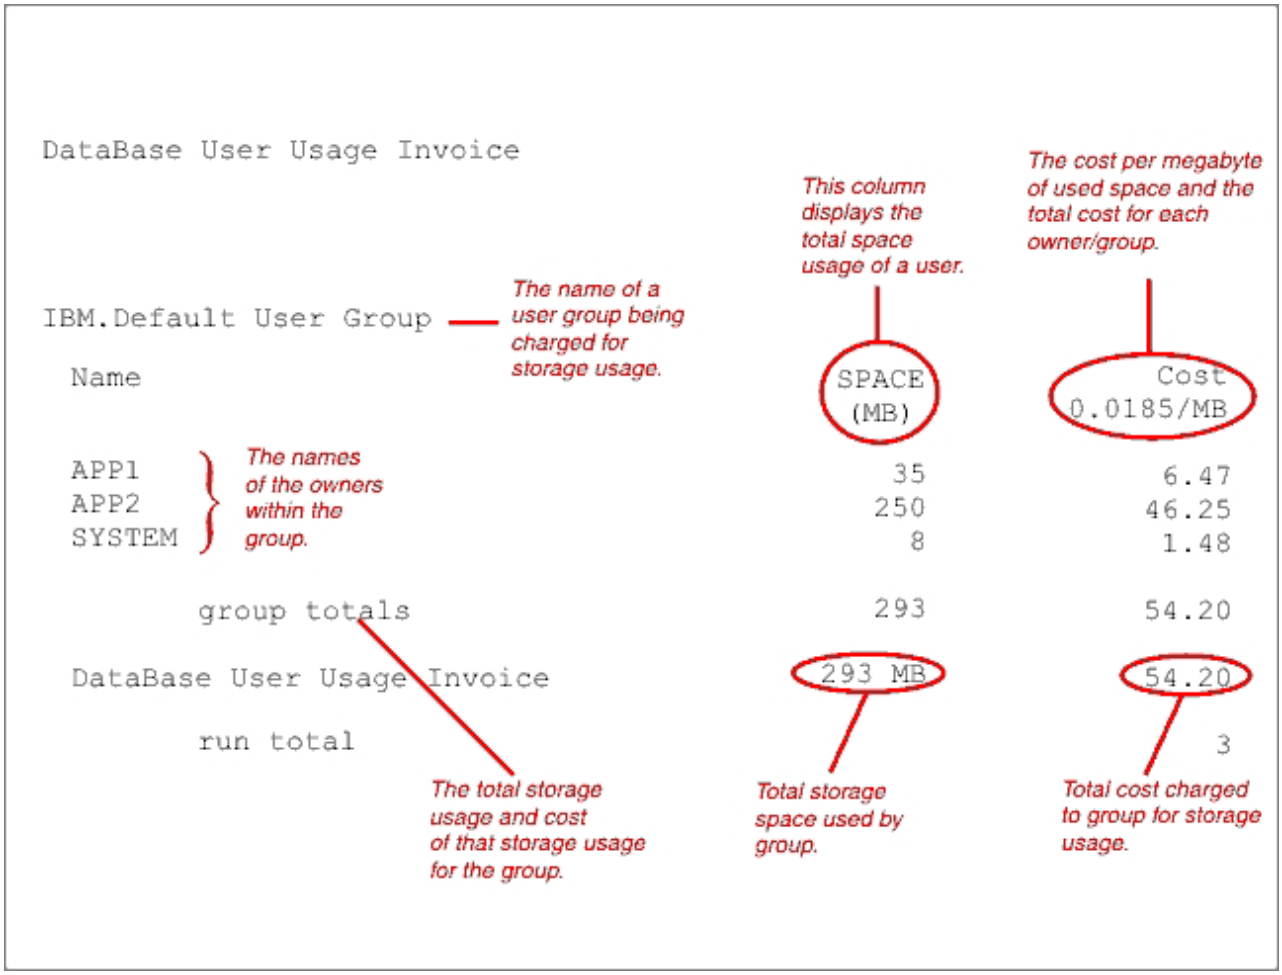

Figure 32. Sample Invoice — Storage Usage by Database User

Total Size By Database -Tablespace :

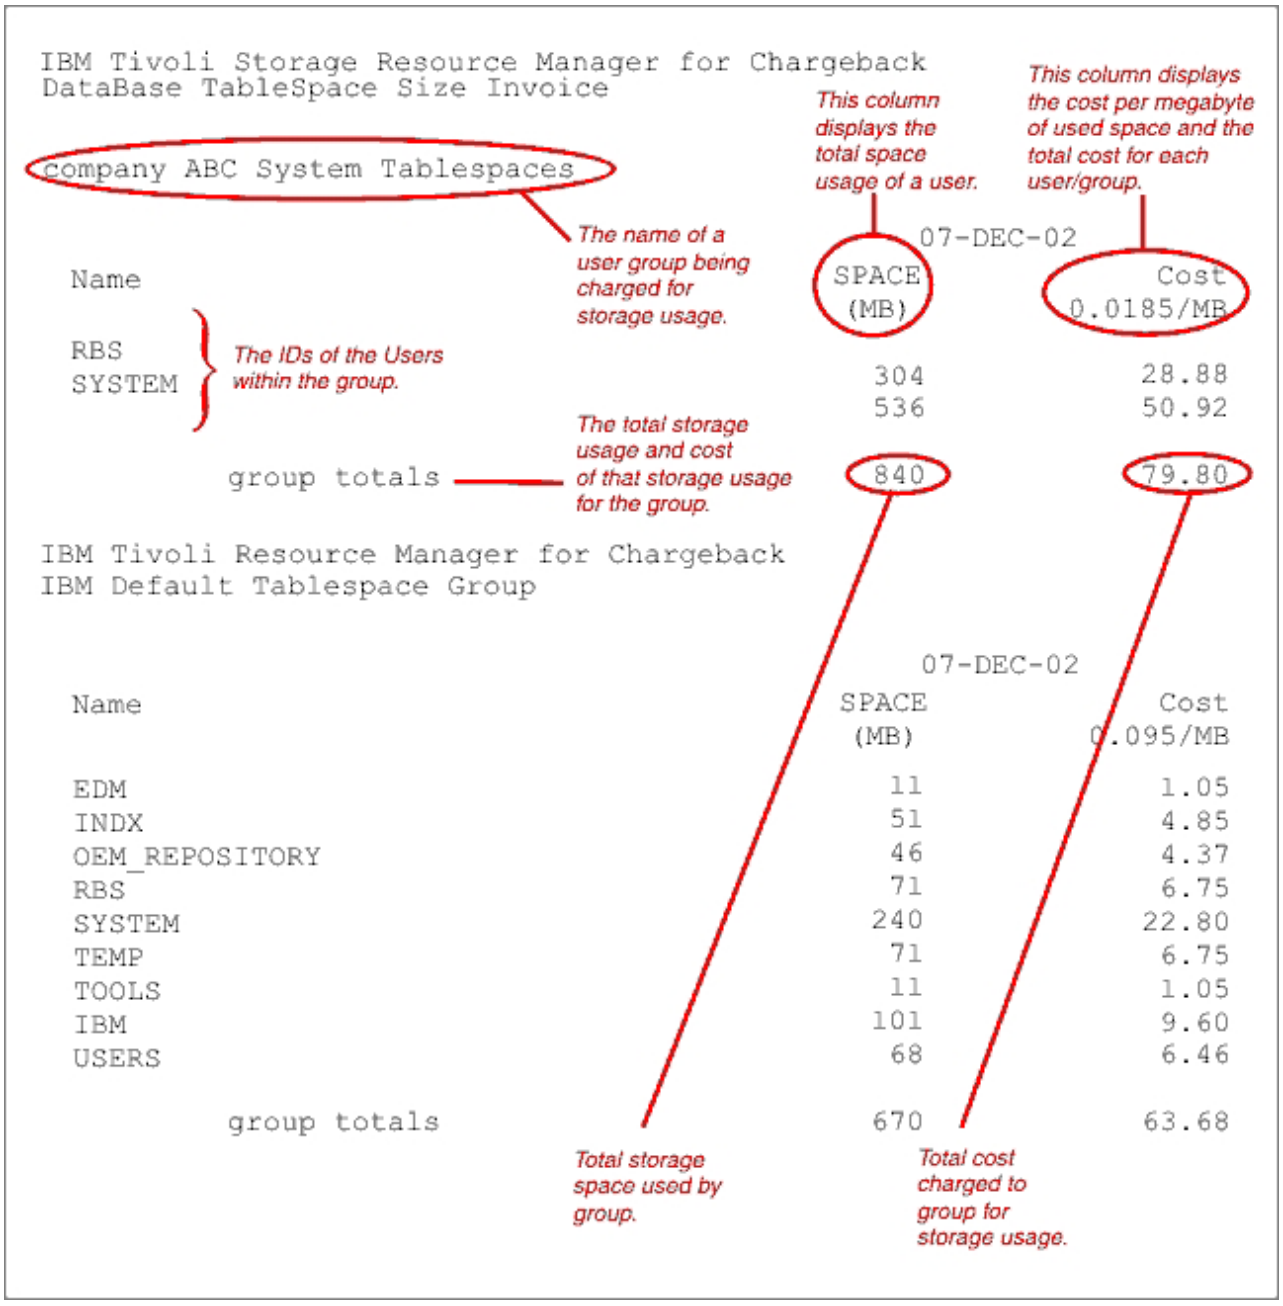

Figure 33. Sample Invoice — Total Size by Database-Tablespace

Chapter 6. Managing replication

Manage replication components, including setting up the management servers, adding storage connections and paths, creating sessions and copy sets, setting up data replication, performing planned maintenance, and recovering from a disaster.

Administering

Administer IBM Tivoli Storage Productivity Center for Replication to authorize users, start and use the graphical user interface (GUI), start and stop services, and many other administrative tasks.

Starting and stopping Tivoli Storage Productivity Center for Replication

Use these procedures to start and stop Tivoli Storage Productivity Center for Replication, including the embedded WebSphere Application Server and DB2.

Starting the Tivoli Storage Productivity Center for Replication server

The commands that are required to start the Tivoli Storage Productivity Center for Replication server depend on the operating system.

The Tivoli Storage Productivity Center for Replication server is started by default after the application is installed.

Starting Tivoli Storage Productivity Center for Replication on Windows:
To start Tivoli Storage Productivity Center for Replication, issue the following command from the command prompt:

```
TPC_install_directory\scripts\startTPCReplication.bat
```

Where the variable *TPC_install_directory* is the installation directory. The default directory is `C:\Program Files\IBM\TPC`.

Starting Tivoli Storage Productivity Center for Replication on AIX or Linux:
To start Tivoli Storage Productivity Center for Replication, issue the following command from the command line:

```
/TPC_install_directory/scripts/startTPCReplication.sh
```

Where *TPC_install_directory* is the installation directory. The default directory is `/opt/IBM/TPC`.

Starting Tivoli Storage Productivity Center for Replication on z/OS:
To start Tivoli Storage Productivity Center for Replication for System z, complete the following steps:

1. Ensure that the Tivoli Storage Productivity Center for Replication for System z HFS is mounted on the UNIX System Services. The root production directory is *path_prefix*`/opt/Tivoli/RM`.
2. Issue the command from the IBM z/OS System Display and Search Facility (SDSF) panel to start Tivoli Storage Productivity Center for Replication for System z, as shown in the following example. The address spaces that are shown might be different in your environment.

© Copyright IBM Corp. 2001, 2013

`/START IWNSRV`

The initiator is IWNSRV and the job name is IWNSRV. For information about the IWNSRV job, see the IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide.

Stopping Tivoli Storage Productivity Center for Replication server

The commands that are required to stop the Tivoli Storage Productivity Center for Replication server depend on the operating system.

Stopping Tivoli Storage Productivity Center for Replication on Windows:
To stop the Tivoli Storage Productivity Center for Replication server, issue the following command from the command prompt:

TPC_install_directory`\scripts\stopTPCReplication.bat`

Where the variable *TPC_install_directory* is the installation directory. The default directory is `C:\Program Files\IBM\TPC`.

Stopping Tivoli Storage Productivity Center for Replication on AIX or Linux:
To stop Tivoli Storage Productivity Center for Replication, issue the following command from the command line:

`/`*TPC_install_directory*`/scripts/stopTPCReplication.sh`

Where *TPC_install_directory* is the installation directory. The default directory is `/opt/IBM/TPC`.

Stopping Tivoli Storage Productivity Center for Replication on z/OS:
To stop Tivoli Storage Productivity Center for Replication for System z, complete the following steps:

1. Issue a command from the IBM z/OS System Display and Search Facility (SDSF) panel to stop Tivoli Storage Productivity Center for Replication for System z, as shown in the following example command. The address space that is shown might be different in your environment.

   /STOP IWNSRV

   Where **/STOP IWNSRV** stops the server address space.
2. Issue the **/D A,L** command to show the processes that are running and verify that the stop command completed. If the IWNSRV address space is displayed, issue the **/CANCEL** command for the address space.

Verifying that the Tivoli Storage Productivity Center for Replication server is running

You can verify that the Tivoli Storage Productivity Center for Replication server is running in several ways.

Complete one or more of these tasks to determine whether the Tivoli Storage Productivity Center for Replication server is running:

- Start the Tivoli Storage Productivity Center for Replication GUI or command-line interface shell. If either of these methods are successful, then the Tivoli Storage Productivity Center for Replication server is running.
- View the `csmTrace.log` file in the *TPC_install_directory*`\wlp\usr\servers\replicationServer\logs\CSM\` directory. If the `csmTrace.log` file is updated regularly and is increasing in size, then Tivoli Storage Productivity Center for Replication is running.

Identifying the version of IBM Tivoli Storage Productivity Center for Replication

The version of code on the IBM Tivoli Storage Productivity Center for Replication server determines the available features and enhancements added to the product.

By clicking **About**, located in the Main menu, you can view the version and release of the IBM Tivoli Storage Productivity Center for Replication. The **-ver** command that is described in the *IBM Tivoli Storage Productivity Center Command-Line Interface Reference* also displays the current version of the product.

Backing up and restoring IBM Tivoli Storage Productivity Center for Replication configuration data

You can back up the entire IBM Tivoli Storage Productivity Center for Replication database. You can also back up copy sets in a specific session. You can then use the backup files to restore a previous configuration or recover from a disaster.

Back up and recovery

You can backup and recover copy set data for a specific session and the complete Tivoli Storage Productivity Center for Replication database.

Copy sets

You can export data about all copy sets in a specific session to maintain a backup copy that you can use to recover if you lose the session or to upgrade to a different management server.

When exporting copy sets, Tivoli Storage Productivity Center for Replication creates a backup of the session and saves the data in a comma separated value (CSV) file, which you can view or edit in a spreadsheet program such as Microsoft Excel. The exported CSV file includes the session name, session type, date that the data was exported, and the copy sets for each role pair. There is one line per copy set, and the volumes in the copy set are separated by a comma (for example: `ESS:2105.FCA57:VOL:17C7,ESS:2105.12043:VOL:17C7`).

The following example illustrates the content of the CSV file for a FlashCopy session. Note that the first valid row must contain the appropriate role names for the session. The order of the copy sets does not matter, and you can include extra roles. A copy set is created from each row that follows the role names. All rows must have data in each column to be a valid row. Note that the number sign (#) indicates that the line is a comment. Lines that are comments are ignored.

```
#Session1,
#FlashCopy,
#Oct 2, 2009 10:03:18 AM

H1,T1
DS8000:2107.FRLL1:VOL:1004,DS8000:2107.FRLL1:VOL:1104
DS8000:2107.FRLL1:VOL:1011,DS8000:2107.FRLL1:VOL:1101
DS8000:2107.FRLL1:VOL:1005,DS8000:2107.FRLL1:VOL:1105
```

Important: You must manually save this file on the local system when you export copy sets from the Tivoli Storage Productivity Center for Replication Web interface.

IBM Tivoli Storage Productivity Center for Replication database

Tivoli Storage Productivity Center for Replication database contains all product data, including data about storage systems, sessions, copy sets, paths, user administration and management servers. You can back up this data and use the backup file to recover from a disaster or restore a previous configuration.

Important: You must have Administrator privileges to back up and recover the database.

The current data is stored in a new file each time you create a backup. The backup file is named *yyyyMMdd_HHmmssSSS*.zip, where *yyyy* is the year, *MM* is the month, *dd* is the day, *HH* is the hour, *mm* is the minute, *ss* is the seconds, *SSS* is the milliseconds when the backup command was run. It is your responsibility to delete backup versions that are no longer needed.

By default, the backup file is stored in the following location:

Tivoli Storage Productivity Center for Replication
TPC_install_directory\wlp\usr\servers\replicationServer\database\backup\

Tivoli Storage Productivity Center for Replication for System z
path_prefix/opt/Tivoli/RM/wlp/usr/servers/replicationServer/database/backup/

You can change the default location by editing the **db.backup.location** property in the rmserver.properties file, which is in the following location:

Tivoli Storage Productivity Center for Replication
TPC_install_directory\wlp\usr\servers\replicationServer\properties\

Tivoli Storage Productivity Center for Replication for System z
path_prefix/opt/Tivoli/RM/wlp/usr/servers/replicationServer/properties/

The backup file contains the Tivoli Storage Productivity Center for Replication database data at the time the backup was performed. Any changes that were made after the backup are not reflected when the backup files are used to restore an Tivoli Storage Productivity Center for Replication database. It is recommended that you create a new backup file:

- After changing the Tivoli Storage Productivity Center for Replication database data, such as adding or deleting a storage system, changing properties, and changing user privileges
- After an Tivoli Storage Productivity Center for Replication session changes direction. For example, if a Metro Mirror session was copying data from H1 to H2 when the backup was taken, and later, the session was started in the H2 to H1 direction. The session must be in the Prepared state before you create the backup.
- After a site switch has been declared and the Enable Copy To Site command has been issued. After you create a backup, consider deleting the previous backup to prevent Tivoli Storage Productivity Center for Replication from starting the copy in the wrong direction.

When you create a backup, ensure that all Tivoli Storage Productivity Center for Replication sessions are either in the Defined, Prepared or Target Available state.

Restoring the Tivoli Storage Productivity Center for Replication database from a backup copy puts Tivoli Storage Productivity Center for Replication back to the point in time when the backup was made. Relationships that exist on the storage systems that were created by Tivoli Storage Productivity Center for Replication after the backup was made will no longer be managed by Tivoli Storage Productivity Center for Replication until you add the copy set to the session and Tivoli Storage Productivity Center for Replication assimilates the relationship into the session. Copy sets that were deleted after the backup will be restored and a subsequent Start command to the session will create new relationships; therefore, you must remove the deprecated copy sets before issuing the Start command.

After restoring a Global Mirror session, you must stop the Global Mirror master and subordinates before restarting the Global Mirror session. Refer to your DS6000 and DS8000 storage system documentation for more information.

Exporting copy set data

You can export data about all copy sets in a specific session, to maintain a backup copy that can be used to recover if you lose your session or upgrade to a different server.

Perform these steps to export the copy sets in a specific session:

1. In the navigation tree, select **Sessions**. The Sessions panel is displayed
2. Select the session for which you want to export copy sets.
3. Select **Export Copy Sets** from the **Actions** list, and click **Go**. The Export Copy Set wizard displays the status of the export and a link to the exported file.
4. Click that link and save the file to the local system.

   Important: You must save the file to your local system. After you close the panel, the data will be lost.
5. Click **Finish**.

Importing copy set data

You can import copy set data that was previously exported to a comma separated value (CSV) file.

Perform the following steps to import copy sets into an existing session:

1. In the navigation tree, select **Sessions**. The Session panel is displayed.
2. Select the session for which you want to import copy sets.
3. Select **Add Copy Sets** from the **Actions** list, and click **Go**. The Add Copy Sets wizard is displayed.
4. Select **Use a CSV file to import copy sets**.
5. Type the location and name of the CSV file to import, or use **Browse** to select the file. Then, click **Next**.
6. Verify that the matching results were successful, and then click **Next**.
7. Select the copy sets you want to add, and then click **Next**.
8. Confirm the number of copy sets that you want to create, and click **Next**. A progress bar displays.
9. Click **Next**.
10. Verify the matches, and click **Finish**.

Verifying the Java heap size

The Java heap size determines the number of volume pairs that Tivoli Storage Productivity Center for Replication can manage.

Verifying the Java heap size for Tivoli Storage Productivity Center for Replication

The default Java heap size is 1024 MB, which supports fewer than 75,000 pairs. If you increase the Java heap size by 256 MB, support increases by 25,000 pairs. You can continue to increase the Java heap size until you reach the maximum heap size.

For more information about the maximum heap size, see the information about sizing the Java heap size in the IBM SDK for Java Tools information center at http://publib.boulder.ibm.com/infocenter/javasdk/tools/index.jsp.

For more information about how to set up the Java heap size, see the *WebSphere Application Server V8.5 Liberty profile* technote at http://www.ibm.com/support/docview.wss?uid=swg21596474.

Verifying the Java heap size for Tivoli Storage Productivity Center for Replication for System z

The default Java heap size is 768 MB, which supports fewer than 50,000 pairs. If you increase the Java heap size by 256 MB, support increases by 25,000 pairs. You can continue to increase the Java heap size until you reach the maximum heap size.

For more information about the maximum heap size, see the information about sizing the Java heap size in the IBM SDK for Java Tools information center at http://publib.boulder.ibm.com/infocenter/javasdk/tools/index.jsp.

For more information about how to set up the Java heap size, see the *WebSphere Application Server V8.5 Liberty profile* technote at http://www.ibm.com/support/docview.wss?uid=swg21596474.

Managing management servers

This section provides information about how to set up active and standby management servers, restore a lost connection between the management servers, or perform a takeover on the standby management server.

Management servers

The *management server* is a system that has IBM Tivoli Storage Productivity Center for Replication installed. The management server provides a central point of control for managing data replication.

You can create a high-availability environment by setting up a *standby management server*. A standby management server is a second instance of Tivoli Storage Productivity Center for Replication that runs on a different physical system, but is continuously synchronized with the primary (or active) Tivoli Storage Productivity Center for Replication server. The active management server issues commands and processes events, while the standby management server records the changes to the active server. As a result, the standby management server contains identical data to the active management server and can take over and run the environment without

any loss of data. If the active management server fails, you can issue the Takeover command to make the standby management server take over.

Connecting the active management server to the standby management server

Ensure that the active management server is connected to the standby management server. This connection creates the *management server relationship* that begins the synchronization process. Each management server can be in only one management server relationship.

Problems might occur during the synchronization process. For example, the management server relationship might become disconnected for a number of reasons, including a connectivity problem or a problem with the alternate server.

If a problem occurs during the synchronization process, the alternate server database is restored to its original state before the synchronization process began. Issue the Reconnect command to restore synchronization.

Completing a takeover on the standby management server

If you must complete a takeover and use the standby server, ensure that you shut down the active management server first. You must ensure that you do not have two active management servers. If there are two active management servers and a condition occurs on the storage systems, both management servers respond to the same conditions, which might lead to unexpected behavior.

If you complete an action on the active management server when the servers are disconnected, the servers will be out of synch.

Viewing the status of the management servers

You can view the status of the active and standby management severs from the Management Servers panel in the Tivoli Storage Productivity Center for Replication graphical user interface (GUI). If you are logged on to the active management server, the icons on this panel show the status of the standby management server. If you are logged on to the standby management server, the icons on this panel show the status of the active management server.

When the status is Synchronized, the standby management server contains the same data that the active management server contains. Any update to the active management server database is replicated to the standby server database.

Managing volumes on storage systems

When you add direct connections, Hardware Management Console (HMC) connections, or z/OS connections on the active management server, Tivoli Storage Productivity Center for Replication automatically enables the management of attached extended count key data (ECKD™) volumes, non-attached count key data (CKD) volumes, and all fixed-block volumes on the storage system. To disable management of volumes on the storage system, use the volume protection function.

Setting exceptions for ephemeral port numbers

For management servers that use an Internet Protocol version 4 (IPv4) connection, the port that is shown on the Management Servers page of the Tivoli Storage Productivity Center for Replication GUI is the *listener* port. This port was set when Tivoli Storage Productivity Center for Replication was installed.

However, when the management server initiates a connection to standby server, the management server uses an *ephemeral* port.

An ephemeral port is a temporary port number that is assigned for the duration the connection. When the connection terminates, the ephemeral port is available for reuse. When the management server reconnects to the standby server, a new ephemeral port number is assigned.

If you are using a firewall, you must create exceptions that enable the management server to connect to the standby server by using outbound ephemeral ports.

Information specific to management servers in z/OS environments

If the standby management server is not in the active server z/OS sysplex, the standby server is not able to communicate with the storage systems by using a z/OS connection. Therefore, another connection must be made by using a TCP/IP connection.

If DB2 is configured for data sharing mode across the z/OS sysplex, one of the Tivoli Storage Productivity Center for Replication servers must be configured to use the zero-administration embedded repository. If the embedded repository is not used, the two servers overwrite the same data in the Tivoli Storage Productivity Center for Replication database.

SNMP alerts

This topic describes the SNMP alerts that are sent by IBM Tivoli Storage Productivity Center for Replication and the associated object IDs (OIDs).

SNMP alerts are sent during the following general events:
- Session state change
- Configuration change
- Suspending-event notification
- Communication failure
- Management Server state change

Session state change SNMP trap descriptions

This topic lists the SNMP traps that are sent during a session state change. A different trap is sent for each state change. These alerts are sent by only the active management server.

A session state change SNMP trap is sent each time the session changes to one of the following states:
- Defined
- Preparing
- Prepared
- Suspended
- Recovering
- Flashing

- Target Available
- Suspending
- (Metro Global Mirror only) SuspendedH2H3
- (Metro Global Mirror only) SuspendedH1H3

An SNMP trap is also sent when a recovery point objective (RPO) threshold is exceeded for a role pair that is in the session.

Table 81. Session state change traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.1	The state of session *X* has transitioned to Defined.
1.3.6.1.4.1.2.6.208.0.2	The state of session *X* has transitioned to Preparing.
1.3.6.1.4.1.2.6.208.0.3	The state of session *X* has transitioned to Prepared.
1.3.6.1.4.1.2.6.208.0.4	The state of session *X* has transitioned to Suspended.
1.3.6.1.4.1.2.6.208.0.5	The state of session *X* has transitioned to Recovering.
1.3.6.1.4.1.2.6.208.0.6	The state of session *X* has transitioned to Target Available.
1.3.6.1.4.1.2.6.208.0.19	The state of session *X* has transitioned to Suspending.
1.3.6.1.4.1.2.6.208.0.20	The state of session *X* has transitioned to SuspendedH2H3.
1.3.6.1.4.1.2.6.208.0.21	The state of session *X* has transitioned to SuspendedH1H3.
1.3.6.1.4.1.2.6.208.0.22	The state of session *X* has transitioned to Flashing.
1.3.6.1.4.1.2.6.208.0.23	The state of session *X* has transitioned to Terminating.
1.3.6.1.4.1.2.6.208.0.26	The recovery point objective for the role pair of *X* in session *Y* has passed the warning threshold of *Z* seconds.
1.3.6.1.4.1.2.6.208.0.27	The recovery point objective for the role pair of *X* in session *Y* has passed the severe threshold of *Z* seconds.

Related reference:

"Session states" on page 751
You can view the health and status of a session in the Tivoli Storage Productivity Center for Replication GUI.

Configuration change SNMP trap descriptions

This topic lists the SNMP traps that are sent when the configuration changes. These alerts are sent by only the active management server.

Configuration change SNMP traps are sent after the following configurations changes are made:

- One or more copy sets have been added or deleted from a session

  An alert is sent for each set of copy sets added to or removed from a session. Note that an alert for copy set changes is sent only once within 15 minutes of a configuration change, so you might not see alerts from successive changes that occur within that 15-minute period. For example, if you make a copy set configuration change that causes an alert to be sent at 10:41:01, and if you were to make additional copy set changes at 10:42:04 and 10:50:09, no alerts would be sent for these two changes because they occurred within the 15-minute minimum interval from the first alert.
- PPRC path definitions have been changed

  An alert is sent for each path configuration change made.

Table 82. Configuration change traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.7	One or more copy sets have been added or deleted from this session. **Note:** An event is sent for each session at least every 15 minutes.
1.3.6.1.4.1.2.6.208.0.8	Peer-to-Peer Remote Copy (PPRC) path definitions have been changed. An event is sent for each path configuration change.

Suspending-event notification SNMP trap descriptions

These SNMP traps that are sent during a suspending-event notification. The traps are sent by the active and standby management server.

Suspending-event notification SNMP traps indicate that a session has transitioned to a Severe status due to an unexpected error.

Table 83. Suspending-event notification traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.9	The session is in a Severe state due to an unexpected error.

Communication-failure SNMP trap descriptions

This topic lists the SNMP traps that are sent during a communication-failure. These alerts are sent by both the active and standby management servers.

Communication-failure SNMP traps are sent after the following events occur:

- A server times out attempting to communicate with a storage system.
- A server encounters errors attempting to communicate with a storage system.
- An active server terminates communication with a standby server as a result of communication errors.
- A standby encounters communication errors with an active server.

After an SNMP trap for a given failure is sent, it is not resent unless communication has been reestablished and failed again.

Table 84. Communication-failure traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.10	Server *X* has timed out attempting to communicate with storage system *Y*.
1.3.6.1.4.1.2.6.208.0.11	Server *X* has encountered errors attempting to communicate with storage system *Y*.
1.3.6.1.4.1.2.6.208.0.12	Active server *X* has terminated communication with standby server *Y* as a result of communication errors.
1.3.6.1.4.1.2.6.208.0.13	Standby server *X* has encountered communication errors with active server *Y*.

Management Servers state-change SNMP trap descriptions

This topic lists the SNMP traps that are sent when the state of the management server changes. These alerts are sent by both the active and standby management servers.

A management server state change SNMP trap is sent each time the management server changes to one of the following states:

- Unknown
- Synchronization Pending
- Synchronized
- Disconnected Consistent
- Disconnected

Table 85. Management Servers state-change traps

Object ID (OID)	Description
1.3.6.1.4.1.2.6.208.0.14	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Unknown (previously Offline).
1.3.6.1.4.1.2.6.208.0.15	The IBM Tivoli Storage Productivity Center for ReplicationServer Management Server connection X->Y has changed state to Synchronized.
1.3.6.1.4.1.2.6.208.0.16	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Disconnected Consistent (previously Consistent Offline).
1.3.6.1.4.1.2.6.208.0.17	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Synchronization Pending.
1.3.6.1.4.1.2.6.208.0.18	The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Disconnected.

Setting up a standby management server

You can set up a standby management server in two ways: setting up the management server you are logged in to as the standby management server, or designating another server as the standby management server.

Note: When you define a standby management server, the IBM Tivoli Storage Productivity Center for Replication code must be at the same level on both the standby and active management servers.

Setting the local management server as the standby server

This topic describes how to set the management server on which you are currently logged in as the standby management server.

Attention: When you set the local management server, the server you are logged in at will be wiped of all session information, and replaced with the session information belonging to the server you specified.

Important: The port number for the standby management server must be the same on both the management server and the standby management server in a high-availability relationship. If you change the standby management server port number on either the management server or the standby management server, you must also change the port number on the other server.

Perform these steps to set the local management server as the standby server:

1. In the navigation tree, select **Management Servers**.
2. From the **Action** menu, select **Set This Server as Standby**, and click **Go**. This opens the Set This Server As Standby panel.
3. Enter the domain name or IP address of the desired active management server.

4. Click **OK** to connect to the active server. You have now designated the server you are logged in at as the standby server.

Setting a remote management server as the standby server

This topic describes how to set up a management server on which you are not logged in as the standby server.

Attention: When you set a remote management server as the standby server, the remote management server is wiped of all session information and replaced with the session information belonging to the management server you specified.

Important: The port number for the standby management server must be the same on both the management server and the standby management server in a high-availability relationship. If you change the standby management server port number on either the management server or the standby management server, you must also change the port number on the other server.

Perform these steps to set a remote management server as the standby server:

1. In the navigation tree, select **Management Servers**.
2. Select **Define Standby** from the drop-down menu, and click **Go**. The Define Standby panel appears.
3. Type the domain name or IP address of the server that you are defining as the new standby management server. Log in to the standby management server by entering the username and password.
4. Click **OK** to connect to the standby management server.

Applying maintenance to an active management server

If an active management server develops a problem during an active session, you can reinstall Tivoli Storage Productivity Center for Replication on the server without affecting the session.

Follow these steps to reinstall Tivoli Storage Productivity Center for Replication on the active management server.

1. If the heartbeat is enabled, disable it:
   a. In the navigation tree, select **Advanced Tools**.
   b. Click **Disable Heartbeat**.
2. On the standby server, Server 2, issue a take-over. This makes Server 2 the active server. It is possible that the original active server, Server 1, is still listed on the Server 2 Management Servers page. If so, select **Remove Standby**.
3. Disable the heartbeat on Server 2, in case there are any problems.
4. Uninstall Tivoli Storage Productivity Center for Replication on Server 1.
5. Reinstall Tivoli Storage Productivity Center for Replication on Server 1.

   Note: If no changes have been made to the configuration while Server 1 was being reinstalled, steps 6 and 7 are not necessary.
6. When Tivoli Storage Productivity Center for Replication is running on Server 1, log into Server 2 and set Server 1 as the standby server for Server 2.

   This step copies the configuration from Server 2 to Server 1. This process takes a few minutes.
7. When the management servers status is synchronized, issue a take-over on Server 1.

   This makes Server 1 an active server, able to control sessions.

Note: It is possible that Server 2 is still listed on the Server 1 Management Servers page. If this is the case, select **Remove Standby**.

8. Disable the heartbeat on Server 1 to make sure this active server does not have any problems.

   Note: If you do not need to reinstall Tivoli Storage Productivity Center for Replication on Server 2, skip steps 9 and 10.

9. Uninstall Tivoli Storage Productivity Center for Replication from Server 2.
10. Reinstall Tivoli Storage Productivity Center for Replication on Server 2.
11. On Server 2, go to the Management Servers page and select the **Set This Server As Standby** option, entering the information for Server 1. When this step is complete, Server 1 is the active server, and Server 2 is the standby server.
12. When you are confident that the active server is running without any problems, enable the heartbeat again, if desired.

Reconnecting the active and standby management servers

If the active and standby management servers become disconnected, reestablish that connection.

Perform these steps to cause the standby management server to become the active management server:

1. In the navigation tree, select **Management Servers**. The Management Servers panel is displayed.
2. Select **Reconnect** from the **Actions** list, and click **Go**.

Performing a takeover on the standby management server

If the active management server fails, you can force the standby management server to take over monitoring and managing replication responsibilities.

Important: If the current active management server is still active, you must not attempt to control the replication environment simultaneously from both management servers. Instead, either reconfigure the current active management server to be a standby management server, or shut it down.

Perform these steps to cause the standby management server to become the active management server:

1. If the active management server is functioning, take it offline so you do not have two active management servers managing the same sessions.
2. Log in to the IBM Tivoli Storage Productivity Center for Replication Web interface running on the standby management server.
3. In the navigation tree, select **Management Servers**. The Management Servers panel is displayed.
4. Select **Takeover** from the **Actions** list, and click **Go**.
5. To reestablish high-availability, perform one of these steps:
   - Choose another server to be the standby management server. See instructions for setting up a standby management server.
   - Bring the failed management server back online, and then make that server the standby management server. See “Setting up a standby management server” on page 639.

- Bring the failed management server back online, and then make that server the active management server to return to the original configuration. Repeat the steps in this section and then add the original standby server as the standby server.

Important: Do *not* use the **Reconnect** command if you perform a takeover. You would use the **Reconnect** command when the active server loses its connection with the standby server; it reconnects the two servers. Do *not* use the **Reconnect** command after a takeover to reconnect to the original active server.

Configuring SNMP

The SNMP community name has a default value of `public`.

To change the community name, modify or add the **`csm.server.snmp_community_string`** property in the `rmserver.properties` file. This file is in the *`TPC_install_directory`*`\wlp\usr\servers\replicationServer\properties\` directory.

Adding SNMP managers

Use the **mksnmp** command to add an SNMP manager to the list of servers to which Tivoli Storage Productivity Center for Replication sends SNMP alerts.

Tivoli Storage Productivity Center for Replication uses management information base (MIB) files to provide a textual description of each SNMP alert that is sent by Tivoli Storage Productivity Center for Replication. You must configure the SNMP manager to use both the `SYSAPPL-MIB.mib` and `ibm-TPC-Replication.mib` files. These files are on the installation DVD in the `replication\CSM-Client\etc` directory. Follow the directions that are provided by your SNMP manager application to configure it to use the MIB files.

Tivoli Storage Productivity Center for Replication sends all SNMP alerts to each registered SNMP manager. SNMP alerts are not specific to any particular session, and all alerts for any session are sent. You cannot choose to send a subset of SNMP alerts.

Changing the standby management server port number

The standby management server port is used is used for communication between the active and standby management server. This port is initially defined during the installation. You can manually change this port after installation.

Important: The port number for the standby management server must be the same on both the management server and the standby management server in a high-availability relationship. If you change the standby management server port number on either the management server or the standby management server, you must also change the port number on the other server.

1. Open the `rmserver.properties` file in the *`TPC_install_directory`*`\wlp\usr\servers\replicationServer\properties\` directory.
2. Modify the port number for the following property:
   `communications.haPort=`*`port_number`*
3. Restart Tivoli Storage Productivity Center for Replication. You must restart Tivoli Storage Productivity Center for Replication to activate property changes. Properties are not synchronized between the Tivoli Storage Productivity Center for Replication management servers and must be done on each Tivoli Storage Productivity Center for Replication management server.

Changing the client port number

The client port is used to log in to the graphical user interface and command-line interface from a remote system. This port is initially defined during the installation. You can manually change this port after installation.

Important: The client port number must be the same on both the management server and the standby management server in a high-availability relationship. If you change the client port number on either the management server or the standby management server, you must also change the port number on the other server.

1. Open the `rmserver.properties` file in the *`TPC_install_directory`*`\wlp\usr\servers\replicationServer\properties\` directory.
2. Modify the port number for the following property:
   `communications.port=`*`port_number`*
3. Open the `repcli.properties` file in the *`TPC_install_directory`*`\cli\` directory.
4. Modify the port number for the following property:
   `port=`*`port_number`*
5. Restart Tivoli Storage Productivity Center for Replication. You must restart Tivoli Storage Productivity Center for Replication to activate property changes. Properties are not synchronized between the Tivoli Storage Productivity Center for Replication management servers and must be maintained on each Tivoli Storage Productivity Center for Replication management server.

Managing storage systems

To replicate data among storage systems using IBM Tivoli Storage Productivity Center for Replication, you must add connections to the storage systems. After a storage system is added, you can associate a location, modify connection properties, set volume protection, and refresh the storage configuration for that storage system.

Storage systems

A *storage system* is a hardware device that contains data storage. Tivoli Storage Productivity Center for Replication can control data replication within and between various storage systems.

To replicate data among storage systems using Tivoli Storage Productivity Center for Replication, you must manually add a connection to each storage system in the Tivoli Storage Productivity Center for Replication configuration. You can then omit storage systems for which Tivoli Storage Productivity Center for Replication does not manage replication and omit storage systems that are being managed by another Tivoli Storage Productivity Center for Replication management server.

For redundancy, you can connect a single storage system using a combination of direct, Hardware Management Console (HMC), and IBM z/OS connections.

You can use the following storage systems:

- IBM TotalStorage Enterprise Storage Server Model 800
- IBM System Storage DS6000
- IBM System Storage DS8000
- IBM System Storage SAN Volume Controller
- IBM Storwize V3500
- IBM Storwize V3700

- IBM Storwize V7000
- IBM Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)
- IBM XIV Storage System

SAN Volume Controller can virtualize various storage systems. Although Tivoli Storage Productivity Center for Replication does not support all storage systems, you can manage other storage systems through a single SAN Volume Controller cluster interface. Tivoli Storage Productivity Center for Replication connects directly to the SAN Volume Controller clusters.

You can define a location for each storage system and for each site in a session. When you are adding copy sets to the session, only the storage systems whose location matches the location of the site are allowed for selection. This ensures that a session relationship is not established in the wrong direction.

Easy Tier heat map transfer

The IBM System Storage DS8000 Easy Tier heat map transfer function transfers heat map information from a source storage system to one or more target storage systems. Each target storage system then generates volume migration plans based on the heat map data and the physical configuration of the storage system.

This process ensures that the performance characteristics of the target storage systems are consistently updated to reflect the performance characteristics of the source storage system.

The Easy Tier heat map transfer function is available for System Storage DS8000 Release 7.1 and later. For more information about the Easy Tier heat map transfer function, go to the System Storage DS8000 information center at http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp and search for *heat map transfer*.

You can enable heat map transfers for System Storage DS8000 storage systems from the IBM Tivoli Storage Productivity Center for Replication graphical user interface (GUI). The storage systems must meet the following requirements:

- The source and target storage systems must be connected to Tivoli Storage Productivity Center for Replication by using a Hardware Management Console (HMC) connection.
- The Easy Tier heat map transfer function must be enabled on the source and target storage systems. If the function is disabled on either the source or target storage system, an error is displayed when the next transfer is attempted.

To enable, disable, and configure heat map transfers, see “Configuring the use of the Easy Tier heat map transfer function” on page 656.

Important:

When you enable or disable the use of the Easy Tier heat map transfer function in Tivoli Storage Productivity Center for Replication, the function is not enabled or disabled on the storage systems that are connected to Tivoli Storage Productivity Center for Replication. The configuration options that you set for Easy Tier heat map transfer in Tivoli Storage Productivity Center for Replication are used only by Tivoli Storage Productivity Center for Replication.

Do not manage heat map transfers for the same System Storage DS8000 storage systems from multiple Tivoli Storage Productivity Center for Replication servers. Doing so causes transfers to occur more frequently than the default transfer schedule, which is 12 hours after the previous transfer.

Storage connections

You must create a connection from the IBM Tivoli Storage Productivity Center for Replication management server to each storage system.

The following types of connection are available. The connection type that you can use depends on the storage system type.

Direct Connection
This connection type is available for all storage systems other than IBM System Storage DS8700 or later storage systems or System Storage DS8000 storage systems that are on an Internet Protocol version 6 (IPv6) network.

Hardware Management Console (HMC) Connection
This connection type is available for System Storage DS8700 or later storage systems and System Storage DS8000 storage systems that are on an IPv4 or IPv6 network.

IBM z/OS Connection
This connection type is available for System Storage DS8000, System Storage DS6000, and IBM TotalStorage Enterprise Storage Server Model 800 storage systems. If the storage system is connected to a z/OS system on which Tivoli Storage Productivity Center for Replication is installed, a native connection is automatically added. If the storage system is connected to another z/OS system, you can use an IP connection to connect to the system.

A single storage system can be connected by using multiple connections for redundancy. For example, you can connect to a System Storage DS8000 storage system by using an HMC connection and a z/OS connection.

When you add a storage connection to the Tivoli Storage Productivity Center for Replication configuration, the storage system and the connection are added to the active management server and the standby server configuration.

The storage systems are not required to be connected to the standby management server. However, if a storage system does not have a connection on the standby management server, you cannot manage copy services on the storage system from the standby server.

Important: If the Metro Mirror heartbeat is enabled, do not connect to a TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 storage system by using both an HMC connection and a direct connection. If you have both types of connections and the direct connection is lost, the session changes to the suspended state even though the HMC connection is still valid. If both connections are lost and the session is in the suspended state, restart the session when connectivity is regained to synchronize the session with the hardware.

Direct connection

By using a TCP/IP connection, the Tivoli Storage Productivity Center for Replication management server can connect directly to certain storage systems. The

TCP/IP connection is required to discover a system's configuration (such as LSSs, volumes, volume size, and format), issue queries, and receive asynchronous events.

The Tivoli Storage Productivity Center for Replication management server can connect directly with the following systems:

- TotalStorage Enterprise Storage Server
- DS6000
- DS8000
- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified
- The XIV system

DS8000 storage systems on an IPv4 network can be connected directly to the management server. A direct connection requires an Ethernet card in the cluster. DS8000 systems on an IPv6 network cannot use a direct connection. They can be connected only through an HMC or z/OS connection.

When you add a direct connection to a DS or ESS cluster, specify the following information for cluster 0 and 1:

- IP addresses or domain names
- Ports
- User names
- Passwords

System Storage SAN Volume Controller or Storwize V7000 can virtualize various storage systems. Although Tivoli Storage Productivity Center for Replication does not support all storage systems, you can manage these storage systems through a single System Storage SAN Volume Controller or Storwize V7000 cluster interface. Tivoli Storage Productivity Center for Replication connects directly to the System Storage SAN Volume Controller or Storwize V7000 clusters. When you add a direct connection to a System Storage SAN Volume Controller or Storwize V7000 cluster to the Tivoli Storage Productivity Center for Replication configuration, specify the cluster IP address of the System Storage SAN Volume Controller or Storwize V7000 cluster, which in turn points to multiple System Storage SAN Volume Controller or Storwize V7000 storage systems. Ensure that the user name and password are correct for the cluster. If incorrect values are used, significant communication problems can occur, such as never advancing to the Prepared state.

Important: The storage system user name must have privileges to maintain SSH keys. For information about troubleshooting Secure Shell connections to the storage system, see the Ethernet Connection Restrictions on SAN Volume Controller website at www-01.ibm.com/support/docview.wss?uid=ssg1S1002896.

Hardware Management Console connection

The IBM Tivoli Storage Productivity Center for Replication management server can connect to DS8000 storage systems through a Hardware Management Console (HMC). An HMC can have multiple DS8000 storage systems connected to it. When you add an HMC to the IBM Tivoli Storage Productivity Center for Replication configuration, all DS8000 storage systems that are behind the HMC are also added. You cannot add or remove individual storage systems that are behind an HMC.

You can also add a dual-HMC configuration, in which you have two HMCs for redundancy. You must configure both HMCs identically, including the user ID and password.

If planned maintenance is necessary on the HMC, it is recommended that you disable the Metro Mirror heartbeat on the management server while the maintenance is performed.

If the HMC needs to go down frequently or reboots frequently, it is recommended that you disable the Metro Mirror heartbeat. If the Metro Mirror heartbeat is required, the direct connection is recommended instead of an HMC connection.

Important: If a DS8000 storage system uses an HMC connection, the Metro Mirror heartbeat could trigger a freeze on the storage system and impact applications for the duration of the long busy timeout timer if the HMC is shut down for any reason, including upgrading microcode. The long busy timeout timer is the time after which the storage system will allow I/O to begin again after a freeze occurs if no run command has been issued by Tivoli Storage Productivity Center for Replication. The default value is two minutes for CKD volumes or one minute for fixed block volumes.

Notes:

- The user ID that you use to connect to the HMC must have admin, op_storage, or op_copy_services privileges on the DS8000 storage system. In addition, the user ID must have Copy Services access to any volumes that are in resource groups on the DS8000 storage system.
- Although you can use either an IP address or a host name to connect to an HMC, the use of a host name ensures that connection to the HMC is not lost if the IP address changes.
- You cannot modify the primary IP address or host name for the HMC in Tivoli Storage Productivity Center for Replication after you add the connection. If you want to change the IP address or host name, you must add a new connection to the HMC and then remove the old connection.
- For minimum microcode requirements to connect to a DS8000 through a management console, see the Supported Storage Products List website at www-01.ibm.com/support/docview.wss?uid=swg21386446.

z/OS connection

This connection type is available for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server Model 800 storage systems. If the storage system is connected to an IBM z/OS system on which Tivoli Storage Productivity Center for Replication is installed, a native connection is automatically added. If the storage system is connected to another z/OS system, you must you use an IP connection to connect to the z/OS system.

To add a z/OS connection, the system must be listed on the Host Systems page and in the Connected state. You can add a connection to a z/OS system from Tivoli Storage Productivity Center for Replication that is running on any operating system. For example, you can add a connection to a z/OS system in Tivoli Storage Productivity Center for Replication that is running on Windows, Linux, or AIX. Or, you can add a connection in Tivoli Storage Productivity Center for Replication that is running on z/OS to another z/OS system.

The z/OS connection is used to issue replication commands and queries for attached count key data (CKD) volumes over an existing Fibre Channel network

and to receive asynchronous events. When a storage system is added to Tivoli Storage Productivity Center for Replication through a z/OS connection, all CKD volumes that are attached to the Tivoli Storage Productivity Center for Replication management system are added to the Tivoli Storage Productivity Center for Replication configuration. CKD volumes that are not attached to the Tivoli Storage Productivity Center for Replication z/OS management server are not added to the Tivoli Storage Productivity Center for Replication configuration through the z/OS connection.

Ensure that all volumes in the logical storage subsystem (LSS) that you want to manage through a z/OS connection are attached to a z/OS system. Either the entire LSS must be attached to z/OS or none of the volumes in the LSS should be attached to z/OS for Tivoli Storage Productivity Center for Replication to properly manage queries to the hardware. For example, if there are two CKD volumes in an LSS, and one volume is attached to Tivoli Storage Productivity Center for Replication by using a z/OS connection and the other is attached through a direct connection, Tivoli Storage Productivity Center for Replication issues commands to both volumes over the Fibre Channel network. However, commands issued to the direct connection volume fail, and Tivoli Storage Productivity Center for Replication shows that the copy set that contains that volume has an error.

Tip: It is recommended that you create both a Hardware Management Console (HMC) and direct connection for CKD volumes to allow for greater storage accessibility.

If a storage system was previously added to the Tivoli Storage Productivity Center for Replication configuration through a z/OS connection and later the storage system is added through a TCP/IP connection, all non-attached CKD volumes and fixed block volumes are added to the Tivoli Storage Productivity Center for Replication configuration.

The Metro Mirror heartbeat

The Metro Mirror heartbeat is not supported through a z/OS connection. To use the Metro Mirror heartbeat, the storage systems must be added by using an HMC or direct connection. If the Metro Mirror heartbeat is enabled and a storage system is added through a direct connection and a z/OS connection, a suspend results if the direct connection becomes disconnected because there is no heartbeat through the z/OS connection.

Removing connections to a z/OS system

If Tivoli Storage Productivity Center for Replication has multiple connections to a storage system, the order in which you remove the connections produces different results:

- If you remove direct and HMC connections first, the fixed-block and non-attached CKD volumes that are attached through these connection types are removed from the Tivoli Storage Productivity Center for Replication configuration.
- The remaining CKD volumes that are attached through the z/OS host connection remain in the Tivoli Storage Productivity Center for Replication configuration until the z/OS host connection is removed.
- If you remove the z/OS host connection first and there is an HMC or direct connection to volumes, those volumes are not removed from the Tivoli Storage Productivity Center for Replication configuration.

Protected volumes

You can mark volumes as protected if you do not want those volumes used for replication.

When a volume is marked as protected, you cannot include that volume in a copy set. This protection applies only to IBM Tivoli Storage Productivity Center for Replication.

You might want to protect a volume in the following instances:

- The volume contains data that you never want to be copied to another volume. For example, the volume is secure, but if the data is copied to an unsecured volume, the data could be read. For this reason, the volume should not be the source for a relationship.
- The volume contains data that you do not want to be overwritten. For this reason, the volume should not be the target of a relationship.

Only administrators can change the volume protection settings.

Site awareness

You can associate a location with each storage system and each site in a session. This *site awareness* ensures that only the volumes whose location matches the location of the site are allowed for selection when you add copy sets to the session. This prevents a session relationship from being established in the wrong direction.

Note: To filter the locations for site awareness, you must first assign a site location to each storage system.

IBM Tivoli Storage Productivity Center for Replication does not perform automatic discovery of locations. Locations are user-defined and specified manually.

You can change the location associated with a storage system that has been added to the IBM Tivoli Storage Productivity Center for Replication configuration. You can choose an existing location or add a new one. Locations are deleted when there is no longer a storage system with an association to that location.

When adding a copy set to a session, a list of candidate storage systems is presented, organized by location. Storage systems that do not have a location are displayed and available for use when you create a copy set.

You can also change the location for any site in a session. Changing the location of a session does not affect the location of the storage systems that are in the session.

Changing the location of a storage system might have consequences. When a session has a volume role with a location that is linked to the location of the storage system, changing the location of the storage system could change the session's volume role location. For example, if there is one storage system with the location of A_Location and a session with the location of A_Location for its H1 role, changing the location of the storage system to a different location, such as B_Location, also changes the session's H1 location to Site 1. However, if there is a second storage system that has the location of A_Location, the session's role location is not changed.

Important: Location matching is enabled only when adding copy sets. If you change the location of a storage system or volume role, IBM Tivoli Storage Productivity Center for Replication does not audit existing copy sets to confirm or deny location mismatches.

Adding a storage connection

You must add connections to storage systems before you can use Tivoli Storage Productivity Center for Replication to manage data replication.

To add a storage system connection, you must belong to the Administrator role.

Adding a connection to a System Storage DS8000 storage system by using an HMC connection

You must use a Hardware Management Console (HMC) connection to connect to a System Storage DS8000 storage system if the system is a DS8700 or later or is on an Internet Protocol version 6 (IPv6) network.

You can also optionally use an HMC connection if the storage system is on an IPv4 network.

Complete the following steps to add a storage system connection:

1. In the navigation tree, select **Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **DS8000** icon.
4. On the Connection page of the Add Storage System wizard, click **HMC** and complete the following fields.

   Primary HMC
   Specify the following information about the primary HMC that is used to connect to the storage system.

   IP Address/Domain Name
   Enter the IP address or host name for the primary HMC. Although either an IP address or a host name are acceptable, the use of a host name ensures that connection to the HMC is not lost if the IP address changes.

   You cannot modify the primary IP address or host name for the HMC in Tivoli Storage Productivity Center for Replication after you add the connection. If you want to change the IP address or host name, you must add a new connection to the HMC and then remove the old connection.

   Username
   Enter the user name for the HMC.

   Password
   Enter the password for the HMC.

   Secondary HMC (Optional)
   Specify information about the secondary HMC, if applicable. A secondary HMC is used for redundancy.

   Both HMCs must be configured identically and must have the same user name and password.

   IP Address/Domain Name
   Enter the IP address or host name for the secondary HMC.

5. Click **Next** and follow the instructions in the wizard.

Adding a connection to a System Storage DS8000 storage system by using a direct connection

You can connect to a System Storage DS8000 storage system by using a direct connection if the system is earlier than a DS8700 and is on an Internet Protocol version 4 (IPv4) network.

Direct connection requires that each of the storage system clusters contains an Ethernet card.

Complete the following steps to add a storage system connection:

1. In the navigation tree, select **Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **DS8000** icon.
4. On the Connection page of the Add Storage System wizard, click **Direct Connect** and complete the following fields.

   Cluster 0
   Specify the following information about connection to cluster 0:

   IP Address/Domain Name
   Enter the IP address or host name for the cluster.

   Port Enter the port number for the cluster. The default value is 2433.

   Username
   Enter the user name for the cluster.

   Password
   Enter the password for the cluster.

   Cluster 1
   Specify the following information about connection to cluster 1. The port, user name, and password for cluster 0 are automatically entered in the corresponding fields for cluster 1. You can edit this information, if required for your environment.

   IP Address/Domain Name
   Enter the IP address or host name for the cluster.

   Port Enter the port number for the cluster. The default value is 2433.

   Username
   Enter the user name for the cluster.

   Password
   Enter the password for the cluster.
5. Click **Next** and follow the instructions in the wizard.

Adding a connection to a System Storage DS6000 or TotalStorage Enterprise Storage Server storage system

You can connect to a System Storage DS6000 or TotalStorage Enterprise Storage Server storage system by using a direct connection to the storage system.

Complete the following steps to add a storage system connection:

1. In the navigation tree, select **Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **DS6000** or **ES800** icon.

4. On the Connection page of the Add Storage System wizard, enter the following information for the storage system clusters:

   Cluster 0
   Specify the following information about connection to cluster 0:

   IP Address/Domain Name
   Enter the IP address or host name for the cluster.

   Port Enter the port number for the cluster. The default value is 2433.

   Username
   Enter the user name for the cluster.

   Password
   Enter the password for the cluster.

   Cluster 1
   Specify the following information about connection to cluster 1. The port, user name, and password for cluster 0 are automatically entered in the corresponding fields for cluster 1. You can edit this information, if required for your environment.

   IP Address/Domain Name
   Enter the IP address or host name for the cluster.

   Port Enter the port number for the cluster. The default value is 2433.

   Username
   Enter the user name for the cluster.

   Password
   Enter the password for the cluster.
5. Click **Next** and follow the instructions in the wizard.

Adding a connection to a SAN Volume Controller or Storwize storage system

You can connect to a SAN Volume Controller or Storwize storage system by using a direct connection to the storage system.

Complete the following steps to add a storage system connection:

1. In the navigation tree, select **Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **SAN Volume Controller**; **Storwize Family** for Storwize V3500, Storwize V3700, or Storwize V7000; or **Storwize V7000 Unified** icon.
4. On the Connection page of the Add Storage System wizard, enter the following information for the storage system cluster:

   Cluster IP / Domain Name
   Enter the IP address or host name of the storage system node.

   Username
   Enter the user name for the node.

   Password
   Enter the password for the node.
5. Click **Next** and follow the instructions in the wizard.

Adding a connection to an XIV system

You can connect to a IBM XIV Storage System by using a direct connection to the storage system.

Complete the following steps to add a storage system connection:

1. In the navigation tree, select **Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **XIV** icon.
4. On the Connection page of the Add Storage System wizard, enter the following information for a node on the storage system. When you specify the information for one node, the remaining nodes are discovered automatically.

   IP Address/Domain Name
   Enter the IP address or host name for a node in the storage system.

   Username
   Enter the user name for the node.

   Password
   Enter the password for the node.
5. Click **Next** and follow the instructions in the wizard.

Adding a connection to a z/OS system

You can connect to storage systems that are connected to an IBM z/OS system regardless of the operating system on which Tivoli Storage Productivity Center for Replication is installed.

If a z/OS system is defined on the Host Systems page, you can add a connection to a z/OS from any operating system. For example, you can add a connection to a z/OS system in Tivoli Storage Productivity Center for Replication that is running on Windows, Linux, or AIX.

Complete the following steps to add a z/OS system connection:

1. In the navigation tree, select **Storage Systems**.
2. On the Storage Systems page, click **Add Storage Connection**.
3. Click the **z/OS Direct Connection** icon.
4. On the Connection page of the Add Storage System wizard, select the storage systems that you want to add.
5. Click **Next** and follow the instructions in the wizard.

Removing a storage system

You can remove a storage system from the IBM Tivoli Storage Productivity Center for Replication configuration.

Prerequisites: You must have Administrator privileges to remove a storage system.

Removing a storage system removes all volumes on that storage system from management server control. All copy sets with a volume on the removed storage system are removed from their respective sessions, making the target volume unrecoverable. All connections to the removed storage system are removed, and any storage systems sharing these connections are also removed.

Perform these steps to remove a storage system:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Select the storage system that you want to remove.

Important: All connections to this storage system will be removed, all volumes on the storage system will be removed from management server control, and all copy sets that have a volume on this storage system will be removed from their respective sessions, leaving the target volume unrecoverable. Any storage systems sharing these connections will be removed as well.

3. Select **Remove Storage System** from the **Actions** list, and click **Go**.
4. Click **Yes** to remove the storage system.

Modifying the location of storage systems

You can associate a location with a storage system after a connection has been made to that storage system.

Prerequisites: You must have Administrator privileges to modify the location of a storage system.

Changing the location of a storage system might have consequences. When a session has a volume role with a location that is linked to the location of the storage system, changing the location of the storage system could change the session's volume role location. For example, if there is one storage system with the location of A_Location and a session with the location of A_Location for its H1 role, changing the location of the storage system to a different location, such as B_Location, also changes the session's H1 location to Site 1. However, if there is a second storage system that has the location of A_Location, the session's role location is not changed.

Perform these steps to modify the location of a storage system:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Change the location of the storage system by selecting a previously defined location from the drop-down list or type a new name in the table cell.

   To disable site awareness, set the location to `None`.

   Note: Locations are deleted from the drop-down list when there is no longer a storage system with an association to that location.

Modifying storage connection properties

You can modify the connection properties for a storage system, including IP addresses, user name and password.

Prerequisites:

- You must have Administrator privileges to modify storage connection properties.
- The storage system must be in the Disconnected state to change most storage connection parameters. You can add a secondary HMC to an existing HMC connection without the HMC being disconnected.

A storage system can lose connection to the management server, for example, if a port is blocked by a firewall or the user name or password is changed on the storage system. If the storage system loses connection, you might need to modify parameters (for example, user name or password) manually on the storage system, and then update the parameters in IBM Tivoli Storage Productivity Center for Replication.

Perform these steps to modify storage connection properties:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Click the **Connections** tab.
3. Perform one of these steps to view details for a specific storage connection:
   - Click the storage connection ID.
   - Select the storage connection, click **View/modify Connection Details** from the actions list, and then click **Go**.
4. Modify the appropriate settings to match the settings for the storage system.
5. Click **Apply** to continue making changes, and click **OK** when finished.

Refreshing the storage system configuration

You can refresh the storage system configuration to query the storage system for changes, such as which volumes are contained in an LSS. You might do this when you reconfigure a storage system and you want IBM Tivoli Storage Productivity Center for Replication to be aware of the changes.

Prerequisites: You must have Administrator privileges to modify storage connection settings.

Perform these steps to refresh the storage configuration:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Select the storage system for which you want to refresh the configuration.
3. Select **Refresh Configuration** from the **Actions** list, and click **Go**.

Setting volume protection

To ensure that data on a volume is not overwritten, you set its status to protected. Protected volumes are excluded from replication.

You must have Administrator privileges to change the protection setting of a volume.

1. In the navigation tree, select **Storage Systems**. The Storage Systems page is opened.
2. Click **Volume Protection**. The Volume Protection wizard is opened.
3. Select a storage system.
4. Optional: Depending on the type of storage system:
   a. Select **All IO Groups** or a specific I/O group.
   b. Select **All Logical Storage Subsystems** or a specific logical storage subsystem.
   c. Select **All Pools** or a specific pool.
5. Optional: In the **Volume** field, select a single volume.
6. Optional: In the **Volume Mask** field, enter a sequence of characters and wildcards that match user-defined or system-defined volume IDs. To protect a specific volume, enter the volume ID such as `ESS:2105.FCA57:VOL:1000`. To use a pattern to retrieve one or more volume IDs, you can enter a partial volume ID and use the wildcard character (*) to represent zero or more characters. For example, to retrieve all volume IDs that contain the characters FCA57, you enter `*FCA57*`.
7. Click **Next**.
8. Verify the search results, and click **Next**.

9. Click **Select All** to protect all the volumes. Alternatively, select a check box next to the volumes that you want to protect.
10. Click **Next**.
11. Click **Finish**.

Restoring data from a journal volume

This topic provides information on restoring data from a journal (J) volume. It describes how to restore data from journal volume used as part of a ESS/DS6000/DS8000 Global Mirror session, or as part of a Metro Global Mirror session, if data was corrupted on a host volume after you issued a **Recover** command. Following these steps will enable you to return to a consistent copy of the data on the host volume.

Perform the following steps to move the data from the journal volume back to the host volume:

Note: Follow these instructions only if you have already issued a **Recover** command to the site containing the journal volume. After the **Recover** command is issued, the journal volume will hold a copy of the consistent data at the time the command was issued.

1. Outside of IBM Tivoli Storage Productivity Center for Replication, using the DS8000 GUI /CLI, issue withdraw initiate background copy (issue a **rmflash -cp** command) on pairs containing the journal volume (for example, `H2J2`). This copies the remaining uncopied tracks from the host to the journal. Then, ensure all the Out of Sync (OOS) tracks reach zero.
2. Create a separate FlashCopy session either with IBM Tivoli Storage Productivity Center for Replication, or with the DS8000 GUI /CLI (issue a **mkflash** command with background copy), with the following conditions:
   - The journal volume (J*x*) is the source volume.
   - The host volume (H*x* or I*x* if using a session with Practice capabilities) is the target volume.
   - *x* is the site the **Recover** command was issued to.

Configuring the use of the Easy Tier heat map transfer function

You can enable Easy Tier heat map transfers for System Storage DS8000 storage systems that are connected to IBM Tivoli Storage Productivity Center for Replication from the Tivoli Storage Productivity Center for Replication graphical user interface (GUI).

The heat map transfer configuration options that are available in Tivoli Storage Productivity Center for Replication include adding storage systems for which you want to transfer heat map data and enabling or disabling the transfer of Easy Tier heat map data for those storage systems.

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is opened.
2. Click the **Easy Tier Heat Map Transfer** tab.
3. The following configuration options are available:

   Status Indicates whether the transfer of Easy Tier heat map data is enabled or disabled.

Enable Transfer
Click to enable the transfer of Easy Tier heat map data, and then click **Yes** in the confirmation window. Easy Tier heat map data is copied from the storage systems that are listed on this page to the target storage systems that are associated with those storage systems.

Disable Transfer
Click to disable the transfer of Easy Tier heat map data, and then click **Yes** in the confirmation window. The storage systems that are listed on this page remain displayed with a status of **Inactive**.

Add Storage Systems
Click to add System Storage DS8000 storage systems that you want to include in the Easy Tier heat map transfer configuration. If there are available systems, a list of storage systems is displayed. Select the storage systems from this list. Only those storage systems that meet the following requirements are displayed:

- The storage system must be version 7.1 or later.
- The storage system must be connected to IBM Tivoli Storage Productivity Center for Replication by using a Hardware Management Console (HMC) connection.
- The storage system is not already displayed on this page.

Select Action
Select a storage system that is listed on this page and then select one of the following actions. When you have made your selection, click **Go**:

Remove Storage System
Click to remove a storage system from this page. The Easy Tier heat map transfer process from the storage system to the target storage systems is stopped.

View Transfer Details
Click to view detailed information about the transfer from the source storage system to its target storage systems. This information includes whether the transfer succeeded or failed and the date and time that the transfer started and stopped. This action is not available if the source storage system does not have target storage systems.

Storage system details
The following information is shown for each storage system:

Storage System
Lists each source storage system and its target storage systems, if available.

Connection Status
One of the following states is displayed for each storage system connection:

Connecting
Indicates that connection to the storage system is in progress.

Connected
Indicates that connection to the storage system is established.

Disconnected
Indicates that connection to the storage system was lost. The connection can be lost for several reasons, including:

- The storage system is not available.
- The user name and password that is used to connect to the storage system was changed.
- There are firewall issues that prevent connection to the storage system.

Inactive
Indicates that the transfer of Easy Tier heat map data is disabled.

Related concepts:

"Easy Tier heat map transfer" on page 644
The IBM System Storage DS8000 Easy Tier heat map transfer function transfers heat map information from a source storage system to one or more target storage systems. Each target storage system then generates volume migration plans based on the heat map data and the physical configuration of the storage system.

Transfer Results for Storage System window

Use the Transfer Results for Storage System window to view details about the transfer of Easy Tier heat map data to target storage systems.

The following information is displayed in this window. This information includes the number of successful and failed transfers and the last start and stop time that a transfer occurred.

Target Storage System
Select the storage system for which you want to view the transfer details.

Attempts
Shows the number of attempts that were made to transfer the Easy Tier heat map data to the target storage system.

Successful Transfers
Shows the number transfers that were successful.

Failed Transfers
Shows the number of transfers that failed.

Last Transfer Start Time
Shows the date and time that the last transfer started. If a transfer is in progress, **In Progress** is shown.

Last Transfer Stop Time
Shows the date and time that the last transfer stopped. If a transfer is in progress, **Unknown** is shown.

Last Transfer Result
Provides a message that explains the results of the last transfer. If a transfer is in progress, a message specifying that the transfer is in progress is shown.

Next Transfer Time
Shows the date and time that the transfer is scheduled to occur again. If a transfer is in progress, **Unknown** is shown. Transfers are automatically scheduled for 12 hours after the previous transfer.

Managing host systems

A host system is an IBM z/OS or AIX system that connects to storage systems to enable certain replication features for those systems.

A connection to a z/OS host system is required if you want to enable z/OS features such as HyperSwap and hardened freeze in Tivoli Storage Productivity Center for Replication sessions.

If Tivoli Storage Productivity Center for Replication is installed on the host system, the host system connection is automatically added to Tivoli Storage Productivity Center for Replication. This connection is referred to as the native z/OS connection.

If Tivoli Storage Productivity Center for Replication is not installed on the host system, you must add the connection to the host system by using an IP connection. To use an IP connection, you must set certain Resource Access Control Facility (RACF®) settings on the z/OS host as described in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

A connection to an AIX host system is required if you want to use the Open HyperSwap feature. This feature enables the automatic swap of input/output (I/O) to the volumes on a secondary site when a failure occurs when I/O is written to the primary site. For the swap to occur, the storage systems that contain the volumes must be attached to an AIX host and the AIX host must be connected to Tivoli Storage Productivity Center for Replication. The Open HyperSwap feature is optional for the Metro Mirror with Failover/Failback session type.

Restriction: Open HyperSwap is not supported for AIX host servers that are in a clustered environment such as IBM PowerHA SystemMirror for AIX.

You can add both z/OS and AIX host connections in Tivoli Storage Productivity Center for Replication regardless of the operating system on which the application is installed. For example, you can add a connection to a z/OS host system in Tivoli Storage Productivity Center for Replication that is running on Windows, Linux, or AIX.

Adding a host system connection

You can add a connection to one or more host systems to the IBM Tivoli Storage Productivity Center for Replication configuration.

Prerequisites: You must have Administrator privileges to add a host system connection.

Follow these steps to add a host system connection:

1. In the navigation tree, select **Host Systems**.
2. Click **Add Host Connection**.
3. In the Add Host Connection window, enter the following information:

   Connection Type
   Select the host system type: z/OS or AIX.

   Host name / IP address
   Enter the host name or IP address for the host system.

   Port Enter the port for the host system.

For z/OS, the port number must be the same as the port number that is specified for the HyperSwap management address space IOSHMCTL SOCKPORT parameter on the z/OS system. For more information about the SOCKPORT parameter, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

For AIX, use the default port number 9930 unless the port was modified in Subsystem Device Driver Path Control Module (SDDPCM).

User name
Enter the user name. The **User name** field is displayed only for z/OS host systems. The user name and password must be the same name and password that are specified in the Resource Access Control Facility (RACF) settings on the host system. For more information about the RACF settings, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

Password
Enter the password. The **Password** field is displayed only for z/OS host systems.

4. Click **Add Host**.

The host system is displayed on the Host Systems page. If connection to the host system is successful, the status **Connected** is displayed. If the connection to the host system failed, the status **Disconnected** is displayed.

Modifying a host system connection

You can modify host system connections that are in the IBM Tivoli Storage Productivity Center for Replication configuration.

Prerequisites: You must have Administrator privileges to modify a host system connection and the connection must be in the **Disconnected** state.

Follow these steps to modify a host system connection:

1. In the navigation tree, select **Host Systems**.
2. On the Host Systems page, click the link for the host system in the **Host System** column. Complete the information in the following fields. The fields that are displayed depend on the host system and the type of connection to the host system.

   Note: All fields are presented for informational purposes; however, some fields are information-only fields.

   Host name or IP address
   For AIX host systems or IBM z/OS host systems that are connected by using an IP connection, enter the host name or IP address for the host system.

   If Tivoli Storage Productivity Center for Replication is installed on the z/OS host system, **ZOS_NATIVE_CONNECTION** is displayed.

   Port Enter the port for the host system.

   For z/OS, the port number must be the same as the port number that is specified for the HyperSwap management address space IOSHMCTL SOCKPORT parameter on the z/OS system. For more information about the SOCKPORT parameter, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

For AIX, use the default port number 9930 unless the port was modified in Subsystem Device Driver Path Control Module (SDDPCM).

Type Shows the type of host system connection.

System name
Shows the name of the z/OS host system. If there is no connection to the host system, **Unknown** is displayed in this field.

Sysplex name
Shows the name of the z/OS sysplex that the host system is in or specifies that the system is not in a sysplex. If there is no connection to the host system, **Unknown** is displayed.

User name
Enter the user name for the z/OS host system. The user name and password must be the same name and password that are specified in the Resource Access Control Facility (RACF) settings on the host system. For more information about the RACF settings, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

Password
Enter the user password for the z/OS host system. If you change the user name, you must change the password.

Sessions
Shows the session that is associated with the AIX host system. The session is displayed only if the Open HyperSwap feature is enabled for the session.

The session name that is displayed might not be a session that is defined on the current Tivoli Storage Productivity Center for Replication server. For example, another Tivoli Storage Productivity Center for Replication server was used to connect to the host and a session name was defined on that server, but was not defined on the current server.

3. Click **Update Host**.

The host system is displayed on the Host Systems page. If connection to the host system is successful, the status **Connected** is displayed. If the connection to the host system failed, the status **Disconnected** is displayed.

Removing a host system connection

You can remove host system connections from the Tivoli Storage Productivity Center for Replication configuration.

Prerequisites: You must have Administrator privileges to remove a host system.

Removing a host system connection disables the ability to use features that are supported by the host such as Open HyperSwap or HyperSwap. Sessions that use the host system to provide these features can no longer communicate with the host and the features are disabled for the entire session.

Follow these steps to remove a host system connection:

1. In the navigation tree, select **Host Systems**.
2. Select the host system connection that you want to remove.
3. Select **Remove Host Connection** from the **Select Action** list, and click **Go**.

4. Click **OK** to remove the host system connection.

The host system is removed from the Host Systems page.

Removing a session from an AIX host system connection

You can remove a session that is associated with an IBM AIX host system from the Tivoli Storage Productivity Center for Replication configuration. When you remove a session, the host no longer recognizes that the session is managing the volumes that are attached to that host. This function is to be used primarily for cleanup purposes.

Prerequisites: You must have Administrator privileges to remove a session association.

When a session has Open HyperSwap enabled, the session communicates with the host system and the host system stores an association to that session on the Tivoli Storage Productivity Center for Replication server. If the Tivoli Storage Productivity Center for Replication server that made the association becomes inaccessible, it might be necessary to clean up and remove the session association from a different Tivoli Storage Productivity Center for Replication server.

If a host system has an associated session, the session name is displayed in the **View/Modify Host Connection** window for the host system. To open this window, click the link for the host system in the **Host System** column on the Host Systems page.

Follow these steps to remove a session association from a host system connection:

1. In the navigation tree, select **Host Systems**.
2. Select the host system connection that contains the session that you want to remove.
3. Select **Remove Session Association** from the **Select Action** list, and click **Go**.
4. Click **OK** to remove the session association.

The session name is removed from the **Sessions** field in the **View/Modify Host Connection** window.

Managing logical paths

Logical paths define the relationship between a source logical subsystem (LSS) and a target LSS that is created over a physical path. To configure logical paths for TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000, use the ESS/DS Paths panel in Tivoli Storage Productivity Center.

To configure partnerships for the following storage systems, use the graphical user interface (GUI) or command-line interface (CLI) for the storage system:

- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified
- The XIV system

Viewing logical paths

You can view all logical paths that are defined on an IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, or IBM System Storage DS6000 storage system.

Perform one of these procedures to view logical paths:

- From the ESS/DS Paths panel of IBM Tivoli Storage Productivity Center for Replication:
  1. In the navigation tree, select **ESS/DS Paths**. The ESS/DS Paths panel is displayed.
  2. Click the storage system ID to display logical paths for that storage system.
- From the Storage Systems panel:
  1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
  2. Select an ESS, DS6000, or DS8000 storage system for which you want to view logical paths.
  3. Select **View Paths** from the **Select Action** list, and click **Go**. The ESS/DS Paths panel is displayed with a list of defined logical paths.

Adding logical paths

This topic describes how to add IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 logical paths.

Ensure that you have defined the appropriate storage systems on the Storage Systems panel.

Perform these steps to add logical paths:

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **ESS/DS Paths**.
2. Click **Manage Paths**. The Path Management wizard is displayed.
3. From the drop-down boxes in the Path Management wizard, select the source storage system, source logical storage system, target storage system, and target logical storage system. Then, click **Next**.
4. From the drop-down boxes in the Path Management wizard, select the source port and target port and click **Add**. You can add multiple paths between the logical storage subsystems, or just one at a time. When you have made your selections, click **Next**.
5. Confirm your selections and click **Next**.
6. Verify the remaining wizard panels and click **Next**.
7. Click **Finish**.

Adding logical paths by using a CSV file

You can create a comma separated (CSV) file to define logical paths. The CSV file specifies storage systems pairings and associated port pairings that are used for replication. Tivoli Storage Productivity Center for Replication uses the port pairings that are defined in the CSV file to establish logical paths

Perform these steps to add IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 logical paths by using a CSV file:

1. Create a CSV file named `portpairings.csv` in the *TPC_install_directory*`\wlp\usr\servers\replicationServer\properties\` directory. You can create the CSV file in a spreadsheet such as Microsoft Excel or in a text editor. The `properties` directory contains an example CSV file that is named `portpairings.csv.sample`, which contains the following example port pairs:

   ```
   #
   # Example CSV file
   #
   2107.04131:2107.01532,0x0331:0x0024,0x0330:0x0100,0x0331:0x000C
   2107.05131:2107.01532,0x0330:0x0029,0x0331:0x0001
   ```

   Each line represents a storage system to storage system pairing. The first value represents the storage systems, which are delimited by a colon. The remaining values are the port pairs, which are delimited by a colon. All values are separated by a comma. Commented lines must start with #,
2. To enable the changes in the file, complete a task that requires new paths to be established. For example, suspend a session to remove the logical paths and then issue the **Start H1->H2** command to enable the paths to use the port pairings in the CSV file.

Considerations when creating and using the CSV file:

- The CSV file does not affect Global Mirror control paths.
- Port mapping is bidirectional. A logical path is established from system A to system B and from system B to system A depending on the direction of the pairs on the hardware.
- If the CSV file contains multiple lines that specify the same storage system to storage system pairing, Tivoli Storage Productivity Center for Replication uses the last line. This rule applies regardless of the order of the storage system pairing. For example, if you have storage systems 2107.04131:2107.01532 defined on the first line of the CSV file and then have 2107.01532:2107.04131 defined on the second line, Tivoli Storage Productivity Center for Replication uses second line.
- If a line in the CSV contains information that is not formatted correctly, the line is ignored. This rule includes lines that specify storage systems but do not include ports or include ports that are not formatted correctly.
- If the CSV file contains valid and invalid port pairs, the valid port pairs might or might not be established. Invalid port pairs can cause the following errors to be displayed in the Tivoli Storage Productivity Center for Replication console and in the ESS/DS Paths panel:
  - Return Code F52: This error is displayed if a port is invalid.
  - Return Code 0400: This error is displayed if a port is invalid and out of the range for the device.

  Other storage system error codes might be displayed also, depending on the path topology, types of paths, and the incorrect port pairings that are specified in the CSV file.
- If the CSV file contains no valid port pairs, no logical paths are established and subsequent commands to the storage systems that require logical paths might fail. If there are existing logical paths for a storage system, those paths are used until they are removed.

Removing logical paths

This topic describes how to remove IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 logical paths.

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **ESS/DS Paths**.
2. Click the link for the storage system that contains the paths that you want to remove.
3. Select the paths that you want to remove.
4. From the drop-down box, select **Remove**.
5. Click **Go**.

Setting up data replication

This topic describes the how to set up data replication in your environment, including creating sessions and adding copy sets to those sessions.

A *session* is a container of multiple copy sets managed by a replication manager. A *copy set* is a set of volumes that contain copies of the same data. All the volumes in a copy set are the same format (count key data [CKD] or fixed block) and size. In a replication session, the number of volumes in a copy set and the role that each volume in the copy set plays are determined by the session type.

Sessions

A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

The type of data replication that is associated with the session determines the actions that you can perform against all copy sets in the session, the number of volumes that a copy set can contain, and the role that each volume plays.

Attention: Use only the Tivoli Storage Productivity Center for Replication graphical user interface (GUI) or command-line interface CLI to manage session relationships, such as volume pairs and copy sets. Do not modify session relationships through other interfaces such as the System Storage DS CLI. If you modify relationships through other interfaces, a loss of consistency can occur across the relationships that are managed by the session. The exceptions to this requirement are failover operations that are managed by external applications for certain session and storage system types as described in "Failover and failback operations" on page 671.

Copy sets

During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a *copy set*.

Each volume in a copy set must be of the same size and volume type. For example, IBM System Storage SAN Volume Controller volumes must be used with other SAN Volume Controller volumes. The number of volumes in the copy set and the role that each volume plays is determined by the session type that is associated with the session to which the copy set belongs.

The following table lists the estimated number of role pairs and volumes per copy set that are supported for each session type.

Table 86. Number of role pairs and volumes per copy set for each session type

Session	Role Pairs	Volumes
Basic HyperSwap	1	2
FlashCopy	1	2
Snapshot[1]	0	1
Metro Mirror Single Direction	1	2
Metro Mirror Failover/Failback	1	2
Metro Mirror Failover/Failback with Practice	3	3
Global Mirror Single Direction (IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, IBM System Storage DS6000)	3	3
Global Mirror Single Direction (SAN Volume Controller, Storwize)	1	2
Global Mirror Failover/Failback (TotalStorage Enterprise Storage Server, System Storage DS8000, System Storage DS6000)	3	3
Global Mirror Failover/Failback (SAN Volume Controller, Storwize)	1	1
Global Mirror Failover/Failback with Practice (TotalStorage Enterprise Storage Server, System Storage DS8000, System Storage DS6000)	5	4
Global Mirror Failover/Failback with Practice (SAN Volume Controller, Storwize)	3	3
Global Mirror Either Direction with Two-Site Practice	8	6
Global Mirror Failover/Failback with Change Volumes	1	4
Metro Global Mirror	6	4
Metro Global Mirror with Practice	8	5
1. An XIV Snapshot session requires that the user define only the H1 volumes. All target volumes are created on the same storage pool as the source volumes.		

Use the Add Copy Sets wizard to add copy sets to a session. You can select a storage system; a logical subsystem (LSS), I/O group, or pool; or a single volume for reach role and then create one or more copy sets for the session.

You can use one of the following volume pairing options to automatically create multiple copy sets in the same session.

Storage system matching (System Storage DS8000, System Storage DS6000, or TotalStorage Enterprise Storage Server Model 800 Metro Mirror sessions only)
Creates copy sets by matching volumes (based on the volume IDs) across all LSSs for the selected storage systems. For example, volume 01 on the source LSS is matched with volume 01 on the target LSS.

You can select the storage system and select **All Logical Subsystems** in the list of LSSs. You can also do auto-matching at the LSS level for Metro Mirror sessions.

LSS, I/O group, or pool matching
Creates copy sets by matching all volumes based on the selected LSS, I/O group, or pool for each role in copy set.

Select the storage system and LSS, I/O group, or pool, and then select **All Volumes** in the **Volume** list.

If you do not want to use the auto-generated volume pairing for a copy set, clear that copy set so that it is not added during the wizard. Then, add the remaining copy sets and reopen the Add Copy Set wizard and manually enter the volume pairings that you want.

Invalid copy set are not added to the session. Copy sets can be invalid if their volumes are not the same type or size.

You can remove copy sets that you do not want to add to the session, even if they are valid. This process enables you to filter and eliminate unwanted copy sets before they are added to the session.

You can export the copy sets to take a snapshot of your session at a particular point in time for backup purposes.

Note: You can copy an entire storage system only for Metro Mirror sessions.

Adding copy sets

When you create a copy set for a session, a warning is displayed if one or more of the selected volumes exist in another session. If the volume you selected is in another session, confirm whether you want to add the volume to this session. Ensure that you intend to have the volume in multiple sessions because having the volume in multiple sessions can create conflicts. A scenario in which you would want the same volumes in multiple sessions is one in which you create a session for normal replication and then create another session for a disaster recovery practice. You must use the same target volumes in both sessions.

You can use extent space-efficient volumes as copy set volumes depending on the capabilities of the storage system. Extent space-efficient volumes must be fixed block (FB). You cannot use count key data (CKD) volumes.

Removing copy sets

You remove a copy set or range of copy sets by selecting the following items:

- Source volume
- LSS, I/O group, or pool
- Storage system

When the list of copy sets that meet your criteria is displayed, you can select the copy sets that you want to remove.

The consequence of removing copy sets varies depending on the state of the session:

Defined
There is no relationship on the hardware. The copy set is removed from Tivoli Storage Productivity Center for Replication data store.

Preparing or Prepared
The copy set is currently copying data, so Tivoli Storage Productivity Center for Replication terminates the hardware relationship for the copy set. The rest of the copy sets continue to run uninterrupted.

Suspended or Target Available
Any existing relationships on the hardware are removed for the copy set.

Before removing all copy sets from that session, terminate the session. Removing the copy sets when the session is active can considerably increase the amount of time it takes for the copy set removal to complete. Copy sets are removed one at a time, and when the session is active, commands are issued to the hardware. However, if you terminate the session first, then commands are not issued to the hardware and the removal process completes faster.

Tip: When you a remove copy set from Tivoli Storage Productivity Center for Replication, you might want to keep hardware relationships on the storage systems. These relationships are useful when you want to migrate from one session type to another or when resolving problems. For more information about keeping the hardware relationships when removing copy sets, see Removing Copy Sets.

Depending on the storage system, the following actions occur when a copy set is removed:

ESS 800, DS6000, and DS8000:

- The complete copy set is removed from Tivoli Storage Productivity Center for Replication.
- Any peer-to-peer remote copy (PPRC) pair that is part of a Global Mirror consistency group is removed from the consistency group on the storage system.
- If the PPRC pair is part of a Global Mirror consistency group and is the last remaining source volume in a subordinate session, the subordinate session is removed from the storage system.
- If the PPRC pair is the last remaining participant in a Global Mirror session, the Global Mirror session is removed from the storage system.
- Any PPRC relationship remains on the storage system.
- A Metro Mirror (synchronous PPRC) pair that is in a HyperSwap configuration is removed from that configuration but the pair remains on the hardware.
- FlashCopy relationship remains on the storage system if the hardware has not completed any background copy.

SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, Storwize V7000 Unified, or the XIV system

- The complete copy set is removed from Tivoli Storage Productivity Center for Replication.
- FlashCopy, Metro Mirror, and Global Mirror relationships are pulled out of their consistency group. If they are the last remaining relationship in a consistency group, that consistency group is removed from the hardware.

When you specify the force removal option, all knowledge of the specified copy set is removed from Tivoli Storage Productivity Center for Replication, even if the relationship itself still exists. In this situation, you are not able to remove the relationship using Tivoli Storage Productivity Center for Replication, because no information about the relationship exists. If you force a removal of a copy set and the removal fails, you must manually remove the relationship from the hardware. If you do not, you cannot to create new relationships.

One benefit of forcing a removal of the copy sets is that Tivoli Storage Productivity Center for Replication does not manage the consistency of copy sets that it has no knowledge of. This means that additional commands to the session do not affect the removed copy sets, even though they are still in a relationship on the hardware.

If you do not specify the force removal option and an error occurs that prevents the hardware relationships from being removed, the copy set is not removed from Tivoli Storage Productivity Center for Replication. The copy set remains as part of the session, and you can still perform actions on it.

Related tasks:

"Creating sessions and adding copy sets" on page 728
A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

"Removing copy sets" on page 741
This topic describes how to remove copy sets.

Volume roles

Volume roles are given to every volume in the copy set. The role defines how the volume is used in the copy set and the site location of the volume. For example, the H1 role is made up of host-attached volumes that are located at the primary site.

The site determines the location of the volumes. The number of sites in a copy set is determined by the session type. Tivoli Storage Productivity Center for Replication supports up to three sites:

Site 1 The location of the primary storage that contain the source data. Upon initial configuration, this site contains the host volumes with updates that are copied to the target volumes.

Site 2 The location of the secondary storage that receives the copy updates from the primary storage.

Site 3 The location of the tertiary storage that receives data updates from the secondary storage.

The volume roles that are needed in a copy set are determined by the type of replication that is associated with the session. Tivoli Storage Productivity Center for Replication supports these volume roles:

Host volume
A host volume is a volume that is connected to a server that reads and writes I/O. A host volume can be the source of updated tracks when the server connected to the host volume is actively issuing read and write input/output (I/O). A host volume can also be the target of the replication. When the host volume is the target, writes are inhibited.

Host volumes are abbreviated as H*x*, where *x* identifies the site.

Journal volume
A journal volume stores data that has changed since the last consistent copy was created. This volume functions like a journal and holds the required data to reconstruct consistent data at the Global Mirror remote site. When a session must be recovered at the remote site, the journal

volume is used to restore data to the last consistency point. A FlashCopy replication session can be created between the host or intermediate volume and the corresponding journal volume after a recover request is initiated to create another consistent version of the data.

Journal volumes are abbreviated as J*x*, where *x* identifies the site.

Intermediate volume
An intermediate volume receives data from the primary host volume during a replication with practice session. During a practice, data on the intermediate volumes is flash copied to the practice host volumes.

Depending on the replication method being used, data on intermediate volumes might not be consistent.

Intermediate volumes are abbreviated as I*x*, where *x* identifies the site.

Target volume
A target volume receives data from a host or intermediate volume. Depending on the replication type, that data might or might not be consistent. A target volume can also function as a source volume. For example, a common use of the target volume is as a source volume to allow practicing for a disaster (such as data mining at the recovery site while still maintaining disaster recovery capability at the production site).

Change volume (only SAN Volume Controller or Storwize storage systems)
A volume that contains point-in-time images that are copied from the host or target volume.

Change volumes are abbreviated as C*x*, where *x* identifies the site.

Related concepts:

"Session types" on page 53
Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

"Copy sets" on page 665
During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a *copy set*.

Role pairs

A *role pair* is the association of two volume roles in a session that take part in a copy relationship. For example, in a Metro Mirror session, the role pair can be the association between host volumes at the primary site and host volumes at the secondary site (H1-H2).

The flow of data in the role pair is shown using an arrow. For example, H1>H2 denotes that H1 is the source and H2 is the target.

Participating role pairs are role pairs that are currently participating in the session's copy.

Non-participating role pairs are role pairs that are not actively participating in the session's copy.

Snapshot sessions do not use role pairs.

Practice volumes

You can use a *practice volume* to practice what you would do in the event of a disaster, without interrupting current data replication. Practice volumes are available in Metro Mirror, Global Mirror, and Metro Global Mirror sessions.

To use the practice volumes, the session must be in the prepared state. Issuing the Flash command against the session while in the Prepared state creates a usable practice copy of the data on the target site.

Note: You can test disaster-recovery actions without using practice volumes; however, without practice volumes, you cannot continue to copy data changes between volumes while testing disaster-recovery actions.

Consistency groups

For Global Mirror and Metro Global Mirror sessions, IBM Tivoli Storage Productivity Center for Replication manages the consistency of dependant writes by creating a consistent point-in-time copy across multiple volumes or storage systems. A *consistency group* is a set of target volumes in a session that have been updated to preserve write order and are therefore recoverable.

Data exposure is the period when data is written to the storage at the primary site until data is replicated to storage at the secondary site. Data exposure is influenced by factors such as:

- Requested consistency-group interval time
- Type of storage systems
- Physical distance between the storage systems
- Available bandwidth of the data link
- Input/output (I/O) load on the storage systems

To manage data exposure, you can change the consistency group interval time. The *consistency group time interval* specifies how often a Global Mirror and Metro Global Mirror session attempts to form a consistency group. When you reduce this value, it might be possible to reduce the data exposure of the session. A smaller interval value causes the session to attempt to create consistency groups more frequently, which might also increase the processing load and message-traffic load on the storage systems.

Failover and failback operations

If a session has *failover* and *failback* capabilities, you can perform a site switch in which you move the operation of Tivoli Storage Productivity Center for Replication from one site to another, and change the direction of the copy without completing a full copy.

Failover is the process of temporarily switching production to a backup facility following a scheduled maintenance period or a disaster at a production site. A failover operation can be followed by a failback operation, which is the process of returning production to its original location and resuming the copy between the two sites.

The failover and failback operations use remote mirror and copy functions for these tasks:

- To reduce the time that is required to synchronize volumes after switching sites during planned or unplanned outages.

- To eliminate the need to perform a full volume copy from the recovery site to the production site, which can reduce the time that is required to resume operations at the production site.

Failover operations that are managed by other applications

Applications such as the IBM Series i Toolkit, VMware Site Recovery Manager, and Veritas Cluster Server can also manage failover operations for the session types that are listed in the following table.

Session type	Storage systems
Metro Mirror Single Direction	• System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified
Metro Mirror Failover/Failback	• System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified • The XIV system
Metro Mirror Failover/Failback with Practice	• System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified
Global Mirror Single Direction	• System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified
Global Mirror Either Direction with Two-Site Practice	• System Storage DS8000 • System Storage DS6000
Global Mirror Failover/Failback	• System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified • The XIV system

Session type	Storage systems
Global Mirror Failover/Failback with Practice	• System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified
Global Mirror Failover/Failback with Change Volumes	• SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified

If an application completes a failover operation for a session, the Severe status is displayed for the session on the Session Details page in Tivoli Storage Productivity Center for Replication and an error message is generated for the role pairs for which the failover occurred. To view the error message:

1. On the Session Details page, click the link for each role pair for which the failover occurred.
2. On the Role Pair Details page, click the **Show** link for each role pair. This link is in the **Details** column.

To change the session status to Normal, use the application to confirm that the role pairs that are in the session are consistent and then restart the session.

Session types

Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

The session types that are available depend on the storage system type and the edition of Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication that you are using as shown in this topic.

Basic HyperSwap

Basic HyperSwap is a copy services solution for IBM z/OS version 1.9 and later. Basic HyperSwap provides high availability of data if a storage system failure occurs.

The Basic HyperSwap session type enables Tivoli Storage Productivity Center for Replication to manage Basic HyperSwap by using Metro Mirror or Metro Global Mirror replication.

Basic HyperSwap sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 87. Basic HyperSwap

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication Basic Edition for System z • Tivoli Storage Productivity Center for Replication for System z	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000

FlashCopy

FlashCopy replication creates a point-in-time copy of the data on a source volume to a target volume. Data that existed on the target volume is replaced by the copied data.

FlashCopy sessions are available for the following Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 88. FlashCopy

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication Basic Edition for System z • Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3500 • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Snapshot

Snapshot replication creates a point-in-time copy of a volume or set of volumes without having to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

Snapshot sessions are available for the following Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 89. Snapshot

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication Basic Edition for System z • Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	The XIV system

Metro Mirror Single Direction

Metro Mirror is synchronous replication that operates between two sites that are up to 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system.

Metro Mirror replication maintains identical data in both the source and target volumes. When a write operation is issued to a source volume, the changes are propagated to the target volume before the write operation finishes processing.

With the Metro Mirror Single Direction session type, Metro Mirror replication is available only from the primary site.

Metro Mirror Single Direction sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 90. Metro Mirror Single Direction

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Metro Mirror Failover/Failback

Metro Mirror Failover/Failback replication enables you to switch the direction of the data flow so that you can use your secondary site as your production site. You can then copy changes that are made at the secondary site back to the primary site.

Metro Mirror Failover/Failback sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 91. Metro Mirror Failover/Failback

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage) • The XIV system

Metro Mirror Failover/Failback with Practice

Metro Mirror Failover/Failback with Practice replication combines Metro Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Metro Mirror Failover/Failback with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 92. Metro Mirror Failover/Failback with Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Global Mirror Single Direction

Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system.

Global Mirror replication maintains identical data in both the source and target volumes. When a write operation is issued to a source volume, the changes are typically propagated to the target volume a few seconds after the data is written to the source volume.

With the Global Mirror Single Direction session type, Global Mirror replication is available only from the primary site.

Global Mirror Single Direction sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 93. Global Mirror Single Direction

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Global Mirror Either Direction with Two-Site Practice

Global Mirror Either Direction with Two-Site Practice replication enables you to run Global Mirror replication from either the primary or secondary site. This session type provides volumes on the primary and secondary site that you can use to practice for disaster recovery without losing your disaster recovery capability.

Global Mirror Either Direction with Two-Site Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 94. Global Mirror Either Direction with Two-Site Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000

Global Mirror Failover/Failback

Global Mirror Failover/Failback replication enables you to switch the direction of the data flow so that you can use your secondary site as your production site. You can then copy changes that are made at the secondary site back to the primary site.

Global Mirror Failover/Failback sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 95. Global Mirror Failover/Failback

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage) • The XIV system

Global Mirror Failover/Failback with Change Volumes

Global Mirror Failover/Failback with Change Volumes replication provides the same capabilities as the Global Mirror Failover/Failback session type. However, Global Mirror Failover/Failback with Change Volumes sessions also provide the option of enabling or disabling the use of change volumes. Change volumes are available in SAN Volume Controller and Storwize storage systems.

Global Mirror Failover/Failback with Change Volumes sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 96. Global Mirror Failover/Failback with Change Volumes

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• SAN Volume Controller Version 6.3 and later • Storwize V3700 Version 6.4.1 and later • Storwize V7000 Version 6.3 and later • Storwize V7000 Unified Version 1.3 and later (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Global Mirror Failover/Failback with Practice

Global Mirror Failover/Failback with Practice replication combines Global Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Global Mirror Failover/Failback with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 97. Global Mirror Failover/Failback with Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified (Tivoli Storage Productivity Center for Replication copy services are available only for Storwize V7000 Unified block storage)

Metro Global Mirror

Metro Global Mirror replication provides continuous, remote data replication that operates between three sites that are varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source. Metro Global Mirror replication enables you to switch the direction of the data flow so that you can use your secondary or tertiary site as your production site.

Metro Global Mirror sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 98. Metro Global Mirror

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 (only H1 site) • System Storage DS8000

Metro Global Mirror with Practice

Metro Global Mirror with Practice replication combines Metro Mirror, Global Mirror, and FlashCopy capabilities to provide a point-in-time copy of the data on the tertiary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

Metro Global Mirror with Practice sessions are available for the following Tivoli Storage Productivity Center for Replication editions and storage systems.

Table 99. Metro Global Mirror with Practice

Supported editions	Supported storage systems
• Tivoli Storage Productivity Center for Replication for System z • All editions of Tivoli Storage Productivity Center	• TotalStorage Enterprise Storage Server Model 800 (only H1 site) • System Storage DS8000

Related concepts:

"Basic HyperSwap"
Basic HyperSwap is a copy services solution for z/OS version 1.9 and later. It provides high availability of data if a disk storage system failure occurs. Basic HyperSwap does not replace a disaster recovery solution. If a session is suspended but the suspend operation was not caused by a HyperSwap trigger, the session is not frozen to ensure that it is consistent.

"FlashCopy" on page 681
FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established.

"Snapshot" on page 682
Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes. You do not have to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

"Metro Mirror" on page 683
Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is in one storage system and the target is in another storage system.

"Global Mirror" on page 689
Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system. Global Mirror replication maintains identical data in both the source and target volumes.

"Metro Global Mirror" on page 694
Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that are varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source.

Basic HyperSwap

Basic HyperSwap is a copy services solution for z/OS version 1.9 and later. It provides high availability of data if a disk storage system failure occurs. Basic HyperSwap does not replace a disaster recovery solution. If a session is suspended but the suspend operation was not caused by a HyperSwap trigger, the session is not frozen to ensure that it is consistent.

When HyperSwap is combined with Metro Mirror and Metro Global Mirror replication, you can prepare your system for disaster recovery and ensure high availability of data. If a session is suspended but the suspend operation was not caused by a HyperSwap trigger, the session is frozen to ensure that it is consistent.

Restriction: This replication method is available only for ESS, DS6000, and DS8000 storage systems.

Basic HyperSwap replication performs the following actions:

- Manages CKD volumes in Metro Mirror to manage synchronous peer-to-peer remote copy (PPRC) relationships.
- Permits only CKD volumes to be added to the HyperSwap session. The graphical user interface (GUI) shows only CKD volumes when you add a copy set. The command-line interface (CLI) does not add a copy set if a fixed block volume is specified.
- Monitors events that indicate a storage device failed.
- Determines whether the failing storage device is part of a Metro Mirror synchronous PPRC pair.

- Determines the action to be taken from policy.
- Ensures that data remains consistent.
- Swaps the I/O between the primary logical devices in the consistency group with the secondary logical devices in the consistency group. A swap can occur from the preferred logical devices to the alternate logical devices or from the alternate logical devices to the preferred logical devices.

Metro Mirror Failover/Failback with HyperSwap

Metro Mirror Failover/Failback uses HyperSwap to configure and manage synchronous PPRC pairs.

Metro Global Mirror with HyperSwap

Metro Global Mirror with HyperSwap is a z/OS replication feature that provides the three-site continuous replication that you can use if a disaster occurs.

Example

Jane is using multiple DS8000 storage systems. The host applications run on a z/OS operating system and the z/OS environment is connected to the DS8000 storage systems. She has a site in Manhattan and a secondary DS8000 in Hoboken. While it is not required that she has a disaster recovery solution in effect, she does need to have a high-availability solution to keep her applications running all the time. Jane is worried that if a volume fails on the DS8000 in Manhattan, her database application might fail. Even a small downtime can be costly to Jane.

Jane uses a Basic HyperSwap session to mirror the data on the DS8000 in Manhattan to the secondary DS8000 in Hoboken. If a volume at the Manhattan site fails, Basic HyperSwap automatically directs application I/O to the mirrored volumes in Hoboken.

FlashCopy

FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established.

The ESS, DS6000, and DS8000 provide multiple logical subsystems (LSSs) within a single physical subsystem, while the following platforms provide multiple I/O groups:

- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

All platforms can use local replication in which the source volume is located in one LSS or I/O group and the target volume is located in the same or another LSS or I/O group. Using the FlashCopy feature, you can reference and update the source volume and target volume independently.

The following figure illustrates how a FlashCopy session works.

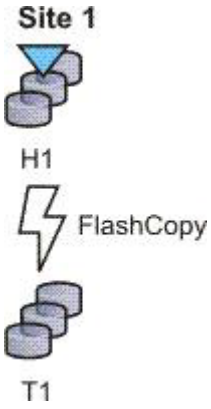

Example

Jane works for a bank. Jane uses a FlashCopy session to make a point-in-time copy of the customer data in existing international accounts. Every night, the bank's servers start batch processing. Jane uses the FlashCopy session to create checkpoint restarts for the batch processing if the batch processing fails. In the batch processing, the first step is to balance all international accounts and to make a FlashCopy point-in-time copy of the resulting data. The second step in the batch processing is to complete the international disbursements.

If the second step in the batch process fails, Jane can use the data from the FlashCopy session that was taken of the first step to repeat the second step. In this way, Jane does not have to begin the entire process again. Jane also uses a CLI script that performs a FlashCopy operation every night at 11:59 PM, and another script that quiesces the database. She backs up the data to tape on the target storage system, and then sends the tape to the bank's data facility for storage.

Snapshot

Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes. You do not have to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

The XIV system uses a snapshot session to create a large number of volume copies without affecting system performance. By using the snapshot function to create a point-in-time copy, and to manage the copy, you can save storage. With the XIV system snapshots, no storage capacity is used by the snapshot until the source volume or the snapshot is changed.

The following figure illustrates how a Snapshot session works.

Example

Jane's host applications are using an XIV system for their back-end storage. With the XIV system, Jane can create a large number of point-in-time copies of the data. The snapshot function ensures that if data becomes corrupted, she can restore the data to any point in time.

Jane sets up a Snapshot session by using Tivoli Storage Productivity Center for Replication and specifies the volumes on the XIV system that are used by the host applications. Jane does not have to provision target volumes for all the snapshots she intends to make. She can quickly configure a single Snapshot session.

When the session is configured, Jane uses a CLI script that runs a **Create Snapshot** command to the session every two hours. If a problem occurs, such as data becoming corrupted, Jane can find a snapshot of the data from a time before the problem occurred. She can restore the data to that point.

By creating a set of snapshots of the data, Jane can also schedule batch processing against that data every day. She can use the batch processing to analyze certain trends in the market without affecting the host applications.

Metro Mirror

Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is in one storage system and the target is in another storage system.

Attention: If you have Tivoli Storage Productivity Center for Replication for System z sessions that contain Metro Mirror relationships, ensure that the session does not contain system volumes (such as paging volumes) unless you select the **Manage H1-H2 with HyperSwap** or the **Enable Hardened Freeze** option for the session. By using these options, z/OS Input/Output Supervisor (IOS) manages freeze operations for the volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself.

Metro Mirror replication maintains identical data in both the source and target. When a write operation is issued to the source copy, the changes that are made to the source data are propagated to the target before the write operation finishes processing. If the storage system ends, no data is lost when you use Metro Mirror if data must be used from the recovery site.

A Metro Mirror session in Global Copy mode creates an asynchronous relationship to accommodate the high volume of data that is migrated. As a result, the data on the target system might no longer be consistent with the source system. The Metro Mirror session switches back to a synchronous relationship when Metro Mirror reissues a **Start** command. In addition, you can start a Metro Mirror session in Global Copy mode and toggle between Metro Mirror and Global Copy modes to accommodate periods in which you require host I/O response time over data consistency.

Tip: To determine whether there is any out-of-sync data that must to be copied before the session is consistent, check the percent that is complete in the session details page.

Metro Mirror Single Direction

The following figure illustrates how a Metro Mirror Single Direction session works.

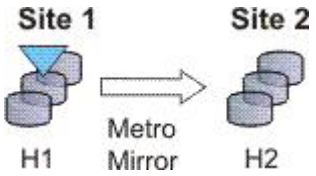

Metro Mirror Failover/Failback

Using Metro Mirror Failover/Failback, the data exists on the second site, which is less than 300 KM away, and you can switch the direction of the data flow. You can use this session type to run your business from the secondary site, and to copy changes that are made at the second site back to the primary site when you want to resume production at the primary site.

The following figure illustrates how a Metro Mirror with Failover/Failback session works.

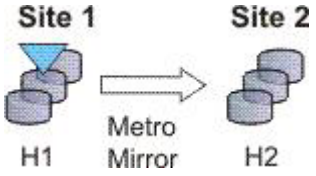

Metro Mirror Failover/Failback with Practice

A Metro Mirror Failover/Failback with Practice session combines Metro Mirror and FlashCopy features to provide a point-in-time copy of the data on the remote site. You can use this session type to practice what you might do if a disaster occurs, without losing your disaster recovery capability.

This solution consists of two host volumes and an intermediate volume.

The following figure illustrates how a Metro Mirror Failover/Failback with Practice session works.

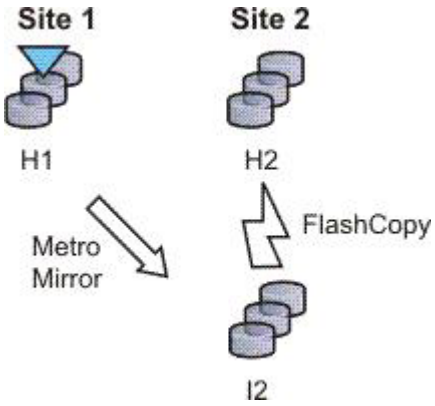

Metro Mirror Failover/Failback with HyperSwap

You can enable a Metro Mirror Failover/Failback session to have HyperSwap capabilities. To enable HyperSwap processing, see "Setting up the environment for HyperSwap" on page 697.

Metro Mirror Failover/Failback with HyperSwap combines the high availability of Basic HyperSwap with the redundancy of a two-site Metro Mirror Failover/Failback solution for managing count key data (CKD) volumes on z/OS. If the primary volumes encounter a permanent I/O error, the I/O is automatically swapped to the secondary site with minimal effect on the application.

A swap can be planned or unplanned. A planned swap occurs when you issue a HyperSwap command from the **Select Action** list in the graphical user interface (GUI) or when you issue a **cmdsess -action hyperswap** command.

The following figure illustrates how a Metro Mirror Failover/Failback session that is enabled for HyperSwap works.

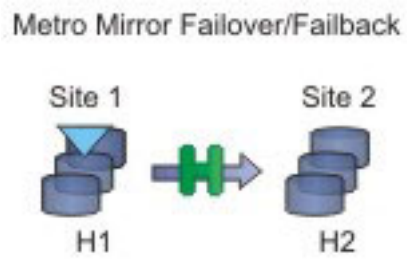

Metro Mirror Failover/Failback with Open HyperSwap

You can enable a Metro Mirror Failover/Failback session to have Open HyperSwap capabilities. To enable Open HyperSwap processing, see "Setting up the environment for Open HyperSwap" on page 698.

Metro Mirror Failover/Failback with Open HyperSwap combines the high availability of Basic HyperSwap on z/OS for fixed-block AIX volumes with the redundancy of a two-site Metro Mirror Failover/Failback solution. If the primary volumes encounter a permanent I/O error, the I/O is automatically swapped to the secondary site with minimal effect on the application.

A swap can be planned or unplanned. A planned swap occurs when you issue a HyperSwap command from the **Select Action** list in the GUI or when you issue a **cmdsess -action hyperswap** command.

The following figure illustrates how a Metro Mirror Failover/Failback session that is enabled for Open HyperSwap works.

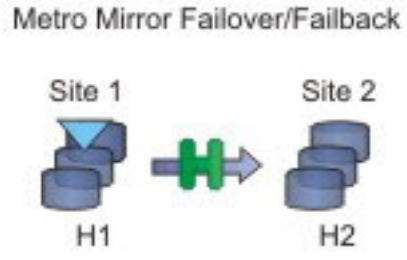

Examples

Read the following scenarios for information on using Metro Mirror for synchronous, remote data replication between two sites.

Metro Mirror Single Direction

At the beginning of a work week, Jane is notified that between 10:00 AM and 11:00 AM on the next Friday, power in her building is going to be shut off. Jane does not want to lose any transactions during the power outage, so she decides to transfer operations to the backup site during the outage. She wants a synchronous copy method with no data loss for the critical business functions, so she chooses Metro Mirror, which can be used between locations that are less than 300 KM apart.

In a synchronous copy method, when a write is issued to change the source, the change is propagated to the target before the write is posted. This method of replication maintains identical data in both the source and target. The advantage of this method is when a disaster occurs, there is no data loss at the recovery site because both writes must complete before signaling completion of a write to the source application. Because the data must be copied to both System Storage DS8000 devices before the write is completed, Jane can be sure that the data is safe.

The night before the planned outage, Jane quiesces the database and servers in San Francisco and starts the database and servers in Oakland. To accomplish this task, Jane issues the **Suspend** and **Recover** commands, and then issues the **Start** command on the secondary site. She shuts down the equipment in San Francisco to avoid any power spikes when she restarts the system after the power is turned on.

Metro Mirror in Global Copy mode

At the beginning of a work week, Jane is notified that between 10:00 AM and 11:00 AM on the next Friday, power in her building is going to be shut off. Jane does not want to lose any transactions during the power outage, so she decides to transfer operations to the backup site during the outage. She wants a synchronous copy method with no data loss for the critical business functions, so she chooses Metro Mirror, which can be used between locations that are less than 300 KM apart.

Jane wants to limit the effect on any applications while completing the initial Metro Mirror synchronization, so she begins the session in Global Copy mode. After she sees that approximately 70% of the data is copied, Jane decides to switch the session to Metro Mirror mode, assuring data consistency.

Metro Mirror with Practice

Jane wants to run a Metro Mirror with Practice from San Francisco to Oakland. She wants to verify the recovery procedure for the Oakland site, but she cannot stop running the Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane can practice the disaster recovery scenario in Oakland while the Metro Mirror session runs uninterrupted. By practicing running the applications at the Oakland site, Jane is better prepared to recover data if a disaster occurs at the San Francisco site.

While her session is running in Prepared state, Jane practices a recovery at the Oakland site by issuing the **Flash** command. This command momentarily pauses the session and starts a FlashCopy to the H2 volumes.

As soon as the FlashCopy is started, the session is restarted. These FlashCopy files create a consistent version of the data on the H2 volume that she can use for recovery testing, while the session continues to replicate data from San Francisco to Oakland. As a result, she can carry out the recovery testing without stopping the replication for any extended duration of time.

If, at some point, the Metro Mirror session is suspended because of a failure, Jane can use the practice session to restart the data replication process. She maintains a consistent copy of the data at the Oakland site, in case of a failure during the resynchronization process. When the session is suspended, she can issue a **Recover** command to create a consistent version of the data on the H2 volumes. After the **Recover** command completes, she can issue the **Start H1->H2** command to resynchronize the data from the San Francisco site to the Oakland site.

If a failure occurs before the restarted session is in Prepared state, she has a consistent version of the data on the H2 volumes. She only has to issue the **Recover** command to put the session into Target Available state and make the H2 volumes accessible from the servers. If the session was not in Prepared state when it was suspended, the subsequent **Recover** command does not issue the FlashCopy files to put the data on the H2 volumes. This means that the consistent data on the H2 volumes is not overwritten if the data to be copied to them is not consistent.

Metro Mirror Failover/Failback enabled for Open HyperSwap

Jane wants to run a Metro Mirror with Practice from San Francisco to Oakland. She wants to verify the recovery procedure for the Oakland site, but she cannot stop running the Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane can practice the disaster recovery scenario in Oakland while the Metro Mirror session runs uninterrupted. By practicing running the applications at the Oakland site, Jane is better prepared to recover data if a disaster occurs at the San Francisco site.

While the session is running in a Prepared state, Jane practices a recovery at the Oakland site by issuing the **Flash** command. This command momentarily pauses the session and starts a FlashCopy to the H2 volumes. As soon as the FlashCopy is started, the session is restarted. These FlashCopy files create a consistent version of the data on the H2 volume that she can use for recovery testing, while the session continues to replicate data from San Francisco to Oakland. As a result, she can carry out the recovery testing without stopping the replication for any extended duration of time.

If the Metro Mirror session is suspended because of a failure, Jane can use the practice session to restart the data replication process while she maintains a consistent copy of the data at the Oakland site, in case of a failure during the resynchronization process. When the session is suspended, she can issue a **Recover** command to create a consistent version of the data on the H2 volumes. After the **Recover** command completes, she can issue the **Start H1->H2** command to resynchronize the data from the San Francisco site to the Oakland site.

If a failure occurs before the restarted session is in Prepared state, she has a consistent version of the data on the H2 volumes. She only has to issue the **Recover** command to put the session into Target Available state and make the H2 volumes accessible from the servers. If the session was not in

Prepared state when it was suspended, the subsequent **Recover** command does not issue the FlashCopy files to put the data on the H2 volumes. This means that the consistent data on the H2 volumes is not overwritten if the data to be copied to the volumes is not consistent.

Selecting a HyperSwap session

A global insurance company decided to use Tivoli Storage Productivity Center for Replication to manage its disaster recovery environment. Jane wants minimal data exposure, both for planned outages such as routine maintenance, and for unplanned disasters. They have CKD volumes on System Storage DS8000 devices, and use z/OS operating systems. They have two data centers in New York.

Jane chooses a Metro Mirror recovery solution because her priority is to protect the system from regional disasters. Jane decides to use Metro Mirror solution because her company has two data centers that are located near each other. Jane realizes that because she uses a z/OS operating system, CKD, and System Storage DS8000 hardware, she can also use a HyperSwap solution. Using Metro Mirror Failover/Failback with HyperSwap, Jane can minimize the effects on any applications, while she maintains the failover process to the secondary site. Jane decides Metro Mirror Failover/Failback with HyperSwap is the best solution.

After installing and configuring Tivoli Storage Productivity Center for Replication on z/OS, Jane starts the Tivoli Storage Productivity Center for Replication GUI. She adds the Tivoli Storage Productivity Center for Replication storage devices that she intends to use on all sites. From the Sessions page, Jane uses the Create Session wizard, and selects the **Metro Mirror Failover/Failback** session type. She selects the **Manage H1H2 with HyperSwap** option. After finishing the wizard, Jane clicks **Launch Add Copy Sets Wizard**. She completes this wizard, and issues a **Start H1->H2** command. After the initial copy is completed, Jane can safely replicate the data between both sites. She can also issue the **HyperSwap** command between sites 1 and 2 to switch sites with minimal effect on the application during either a disaster or maintenance period.

Performing a planned HyperSwap

Jane's company used Metro Mirror Failover/Failback with HyperSwap sessions for the past three months. However, Jane must perform maintenance on an H1 volume. During this time, Jane does not want the applications or replication to be interrupted. To prevent this interruption, before the maintenance is scheduled to begin, Jane uses the Tivoli Storage Productivity Center for Replication GUI to perform a HyperSwap operation to the H2 volumes. This process transitions the applications so that the data is written to H2. To perform a planned HyperSwap operation, Jane issues a **HyperSwap** command.

Understanding what happens when an unplanned HyperSwap occurs

Several weeks after the planned maintenance at Jane's company is completed, an incident occurs at the H1 site. A disk controller fails, causing one of the H1 volumes to encounter a permanent I/O error. Jane's data is safe because she used Metro Mirror Failover/Failback with HyperSwap, and the H2 volume is an exact duplicate of the H1 volume. When the permanent I/O error is detected, a HyperSwap is triggered. The application transitions to write data to the H2 volumes. The applications are not interrupted.

Jane configured a Simple Network Management Protocol (SNMP) listener to alert her to any events, so she receives the SNMP event that indicates

that a HyperSwap occurred. Jane investigates the cause of the HyperSwap process and uses the z/OS console to identify the volume that triggered the HyperSwap process. Jane replaces the faulty disk controller. Then, to recover from the unplanned HyperSwap process, Jane issues the **Start H2->H1** command.

Global Mirror

Global Mirror is asynchronous replication that operates between two sites that are over 300 km apart. The source volumes are on one storage system and the target volumes are on another storage system. Global Mirror replication maintains identical data in both the source and target volumes.

When a write operation is issued to a source volume, the changes are typically propagated to the target volume a few seconds after the data is written to the source volume. However, changes can occur on the source volume before the target volume verifies that it received the change.

Because consistent copies of data are formed on the secondary site at set intervals, data loss is determined by the amount of time since the last consistency group was formed. If the system fails, Global Mirror might lose some data that was transmitted when the failure occurred.

Global Mirror Single Direction

The Global Mirror Single Direction session type enables you to run Global Mirror replication from the primary site.

For TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 storage systems, each copy set in the single direction session consists of two host volumes and a journal volume. The following figure illustrates the data copy flow for Global Mirror Single Direction sessions on these storage systems.

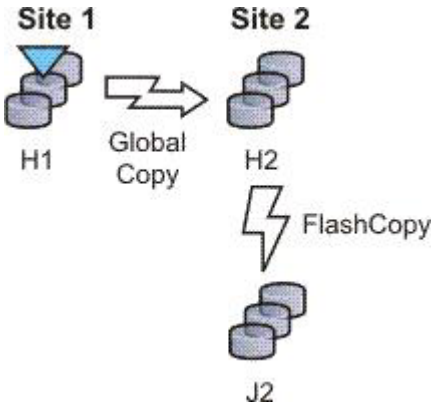

For SAN Volume Controller, Storwize V3700, Storwize V7000, and Storwize V7000 Unified storage systems, each copy set in the Global Mirror Single Direction session consists of two host volumes. The following figure illustrates the data copy flow for Global Mirror Single Direction sessions on these storage systems.

Global Mirror Either Direction with Two-Site Practice

Global Mirror Either Direction with Two-Site Practice enables you to run Global Mirror replication from either the primary or secondary site. This session type provides volumes on the primary and secondary site that you can use to practice for disaster recovery without losing your disaster recovery capability.

This session type is available only for TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 storage systems.

The session consists of two host volumes, two intermediate volumes, and two journal volumes. The following figure illustrates the data copy flow for Global Mirror Either Direction with Two-Site Practice sessions on these storage systems.

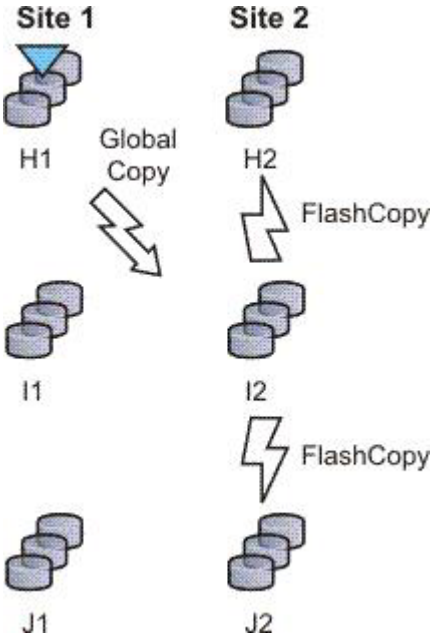

Global Mirror Failover/Failback

The Global Mirror Failover/Failback with Practice session type combines Global Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

For TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 storage systems, each copy set in the Global Mirror Failover/Failback session consists of two host volumes and a journal volume. The following figure illustrates the data copy flow for Global Mirror Failover/Failback sessions on these storage systems.

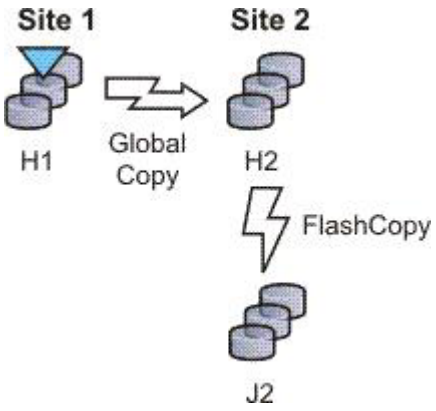

For SAN Volume Controller, Storwize V3700, Storwize V7000, and Storwize V7000 Unified storage systems, each copy set in the Global Mirror Failover/Failback session consists of two host volumes. The following figure illustrates the data copy flow for Global Mirror Failover/Failback sessions on these storage systems.

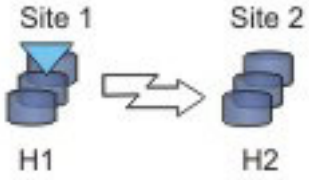

Global Mirror Failover/Failback with Change Volumes

Global Mirror Failover/Failback with Change Volumes replication provides the same capabilities as the Global Mirror Failover/Failback session type. However, Global Mirror Failover/Failback with Change Volumes sessions also provide the option of enabling or disabling the use of change volumes. Change volumes are available in SAN Volume Controller and Storwize storage systems.

Change volumes are composed of a source change volume and a target change volume that contain a point-in-time image of the data from the source and target volumes, respectively.

A FlashCopy operation occurs between the source volume and the source change volume. The frequency of the FlashCopy operation is determined by the cycle period. For information about the cycle period and how it is set, see "Global Mirror Failover/Failback with Change Volumes session properties" on page 777. The data on the source change volume is then replicated to the target volume, and finally to the target change volume.

Because the data that is replicated between sites contains point-in-time changes rather than all changes, a lower bandwidth link is required between the sites. However, when you use change volumes, your data exposure can increase. Therefore, you might want to include or exclude change volumes depending on your network traffic or business requirements.

This session type is available only for SAN Volume Controller, Storwize V3700, Storwize V7000, and Storwize V7000 Unified storage systems. The following figure illustrates the data copy flow for Global Mirror Failover/Failback with Change Volumes sessions on these storage systems.

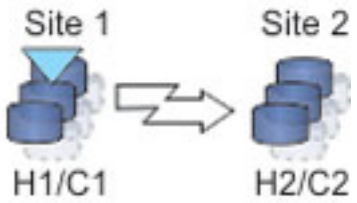

Global Mirror Failover/Failback with Practice

The Global Mirror Failover/Failback with Practice session type combines Global Mirror Failover/Failback and FlashCopy capabilities to provide a point-in-time copy of the data on the secondary site. This session type provides volumes that you can use to practice for disaster recovery without losing your disaster recovery capability.

For TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 storage systems, each copy set in the Global Mirror Failover/Failback with Practice session consists of two host volumes, an intermediate volume, and a journal volume. The following figure illustrates the data copy flow for Global Mirror Failover/Failback with Practice sessions on these storage systems.

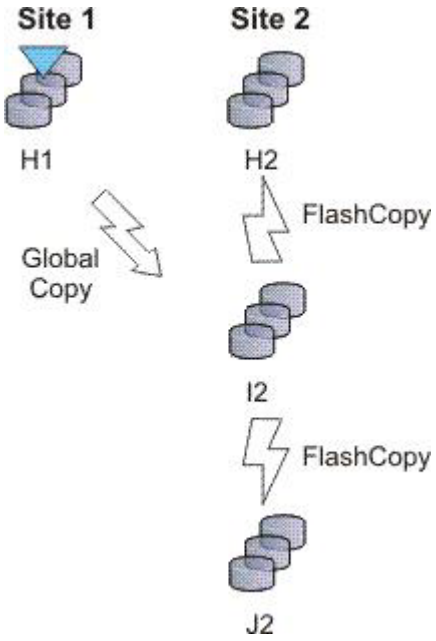

For SAN Volume Controller, Storwize V3700, Storwize V7000, and Storwize V7000 Unified storage systems, each copy set in the Global Mirror failover/failback with Practice session consists of two host volumes and an intermediate volume. The following figure illustrates the data copy flow for Global Mirror Failover/Failback with Practice sessions on these storage systems.

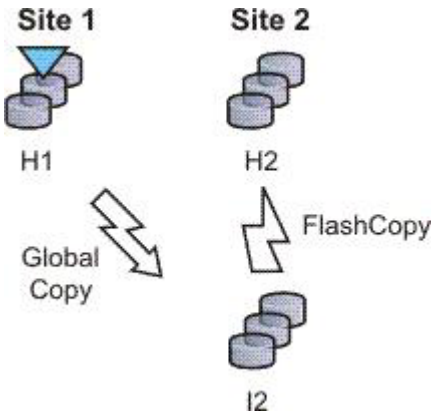

Examples

Read the following scenarios for information about using Global Mirror for asynchronous, remote data replication between two sites.

Global Mirror Single Direction

Although Jane's FlashCopy and Metro Mirror copies were both planned, Jane realizes that sometimes a failure can occur, and she wants to ensure that the data is safe. Because Jane works in San Francisco, she wants her other site to be far away if a disaster occurs locally. Her other site is based in Houston. A minor earthquake occurs in San Francisco and power and communications both fails. Jane arranged for the data on customer accounts that recently opened or closed to be asynchronously copied in Houston, using Global Mirror. Jane risks losing the bytes of data that were being processed when the tremor disrupted the processing in San Francisco. However, she understands that it is minor inconvenience when weighed next to the value of backing up her data to an area that is not prone to earthquakes.

Global Mirror with Practice

Jane wants to run a Global Mirror with practice from San Francisco to Houston. She plans to verify her recovery procedure for the Houston site, but she cannot stop running the Global Mirror session while she takes time to practice a recovery. By using a Global Mirror with Practice session, Jane can practice her disaster recovery scenario in Houston while the Global Mirror session runs uninterrupted. When she practices running the applications at the Houston site, Jane is better prepared to recover the data if a disaster ever strikes the San Francisco site.

Global Mirror Either Direction with Two-Site Practice

Jane wants to run a Global Mirror with Practice from San Francisco to Houston. She plans to verify her recovery procedure for the Houston site, but she cannot stop running the Global Mirror session while she takes time to practice a recovery. By using a Global Mirror Either Direction with Two-Site Practice session, Jane can practice her disaster recovery scenario in Houston while the Global Mirror session runs uninterrupted. When she practices running the applications at the Houston site, Jane is better prepared to recover the data if a disaster ever strikes the San Francisco site.

Jane can use the Global Mirror Either Direction with Two-Site Practice session to run asynchronous consistent data replication from either the San Francisco site or the Houston site. She can practice the disaster-recovery tasks at the target site, regardless of the location of the current production

site. Jane's business can run a consistent Global Mirror session from its Houston site back to San Francisco while it runs a production site at Houston.

Setting up Global Mirror for Resource Groups on System Storage DS8000

If resource groups are defined on a System Storage DS8000, Global Mirror session IDs might be defined for some users. To determine which session IDs are valid, you must modify the `rmserver.properties` file and add the following property:

`gm.master.sessionid.gm_role,session_name = xx`

where *gm_role* is the role that has the master volume (for example, H1 in a Global Mirror failover/failback session), *session_name* is the name of the session that uses the session ID, and *xx* is the decimal number for the session ID.

Tip: System Storage DS8000 represents session IDs as a two-digit hexadecimal number. Use the decimal version of that number. For example, if you want a Global Mirror Failover/Failback session to use a session ID of 0F, the decimal number is 15, as shown in the following example:

`gm.master.sessionid.H2.11194_wprac=15`

Metro Global Mirror

Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that are varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source.

This replication method is available on only TotalStorage Enterprise Storage Server Model 800 and System Storage DS8000 storage systems. You can select ESS storage systems in only the H1 volume role. All other volume roles must use DS8000 volumes.

You can use both ESS and DS8000 volumes in the H1 volume role. If ESS andDS8000 storage systems are both used in the H1 role, the DS8000 storage system performs Incremental Resync processing, and the ESS storage system performs a full copy operation. Because you cannot use the Incremental Resync function with ESS, a full copy is required when you change from H1->H2->H3 to H1>H3 and from H2->H1->H3 to H2->H3.

Attention: If you have Tivoli Storage Productivity Center for Replication for System z sessions that contain Metro Global Mirror relationships, ensure that the session does not contain system volumes (such as paging volumes) unless you select the **Manage H1-H2 with HyperSwap** or the **Enable Hardened Freeze** option for the session. By using these options, z/OS Input/Output Supervisor (IOS) manages freeze operations for the volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself.

Metro Global Mirror maintains a consistent copy of data at the remote site, with minimal effect on applications at the local site. This remote mirroring function works in combination with FlashCopy to provide a disaster-recovery solution that includes the following features:

- Fast failover and failback
- Rapid reestablishment of three-site mirroring, without production outages

- Data currency at the remote site with minimal time lag at the local site, an average of only 3 - 5 seconds for many environments
- Quick resynchronization of mirrored sites using only incremental changes

If Tivoli Storage Productivity Center for Replication runs on a z/OS operating system, you can configure a Metro Global Mirror session to control the Metro Mirror relationship between the primary and secondary site by using the HyperSwap feature. With HyperSwap enabled, a failure on the primary storage system causes an automatic HyperSwap operation, which transparently redirects application I/O to the auxiliary storage system. The Global Mirror relationship continues to run uninterrupted throughout this process. With this configuration, you can achieve almost a zero data loss at larger distances.

Using synchronous mirroring, you can switch from local site H1 to remote site H2 during a planned or unplanned outage. It also provides continuous disaster recovery protection of the H2 and H3 sites. You do not have to configure H2, if a switch from site H1 occurs. With this configuration, you can reestablish H2->H1->H3 recoverability while production continues to run at site H2. Additionally, this setup can reduce the workload on site H1.

Important: In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS storage system, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not available on the ESS storage system, full copies are performed in many of the transitions.

Metro Global Mirror

A Metro Global Mirror session with Practice combines Metro Mirror, Global Mirror, and FlashCopy across three sites to provide a point-in-time copy of the data on the third site.

The following figure illustrates how a Metro Global Mirror session works.

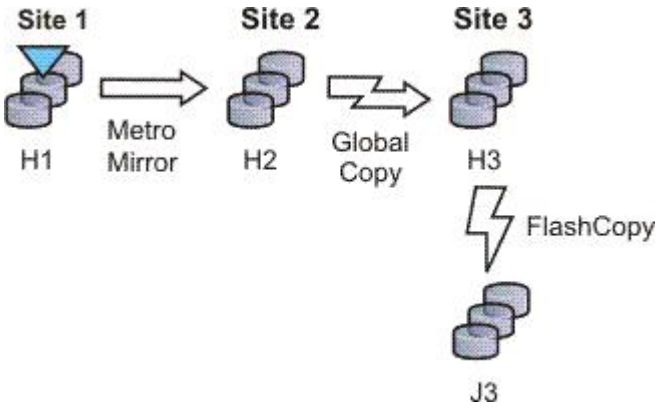

Metro Global Mirror with Practice

A Metro Global Mirror session with Practice combines Metro Mirror, Global Mirror, and FlashCopy across three sites to provide a point-in-time copy of the data on the third site. You can use this session to practice what you might do if a disaster occurred without losing your disaster recovery capability.

The session consists of three host volumes, an intermediate volume, and a journal volume. The following figure illustrates how a Metro Global Mirror with Practice session works.

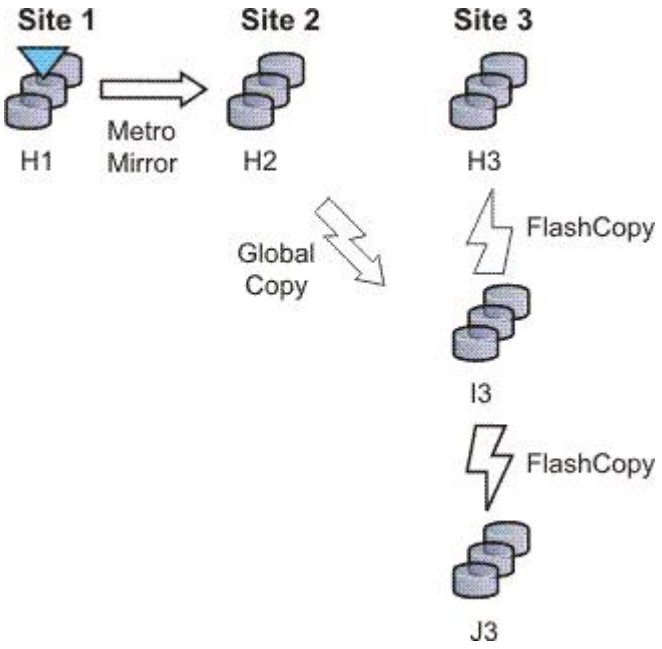

Note: In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS storage system, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not available on the ESS storage system, full copies are performed in many of the transitions.

Examples

Read the following scenarios for information on using Metro Global Mirror for continuous, remote data replication between three sites.

Metro Global Mirror

Although Jane works in San Francisco, she wants to run the business from either Oakland, the secondary site, or Houston, the tertiary site. Jane can use Metro Global Mirror with Failover/Failback to switch the direction of the data flow, so that she can run the business from either Oakland or Houston. Using Metro Global Mirror, Jane has zero data loss at the secondary site and minimal data loss at the tertiary site.

Metro Global Mirror with Practice

Jane wants to run a Metro Global Mirror with Practice from San Francisco to Houston. She plans to verify her recovery procedure for the Houston site. However, she cannot stop running the Metro Global Mirror session while she takes time to practice a recovery. By using a Metro Global Mirror with Practice session, Jane can practice her disaster recovery scenario in Houston while the Metro Global Mirror session runs uninterrupted. When she practices running the applications at the Houston site and is prepared to run the applications at the Oakland site, if necessary, Jane is prepared to recover data if a disaster occurs at the San Francisco site.

Jane can use Metro Global Mirror with Practice to switch the direction of the data flow, so that she can run the business from either Oakland or

Houston. Using Metro Global Mirror, Jane has zero data loss at the secondary site and minimal data loss at the tertiary site.

Managing a session with HyperSwap and Open HyperSwap replication

HyperSwap and Open HyperSwap provide high availability of data if a primary disk storage system failure occurs. When a failure occurs in writing I/O to the primary storage system, the failure is detected by IOS, and IOS automatically swaps the I/O to the secondary site with no user interaction and minimal application effect.

Sessions that can be enabled for HyperSwap or Open HyperSwap:

You can create sessions that enable swapping, which provides a session with a highly available business continuity solution.

Sessions that can enable HyperSwap

The following session types can enable HyperSwap processing:

- Basic HyperSwap
- Metro Mirror with Failover/Failback
- Metro Global Mirror
- Metro Global Mirror with Practice

To enable HyperSwap processing, the following conditions must apply:

- The volumes that are managed by the session are only TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000 storage systems volumes.
- The volumes are count key data (CKD) volumes that are attached to an IBM z/OS system.
- The z/OS system is listed on the Tivoli Storage Productivity Center for Replication Host Systems page and is in the Connected state.

Sessions that can enable Open HyperSwap

Only the Metro Mirror with Failover/Failback session type can enable Open HyperSwap processing.

To enable Open HyperSwap processing, the following conditions must apply:

- The volumes in the session are System Storage DS8000 5.1 or later volumes.
- The volumes in the session are fixed block and mounted to IBM AIX 5.3 or AIX 6.1 hosts with the following modules installed:
  - Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
  - Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)
- The connections between the AIX host systems and the Tivoli Storage Productivity Center for Replication server are established.

Setting up the environment for HyperSwap:

You must set up an environment in which you can run HyperSwap processing before you can enable HyperSwap for a Tivoli Storage Productivity Center for Replication session.

To enable HyperSwap processing, the following conditions must exist:

- The HSIB and IOSHSAPI address spaces are started on the IBM z/OS system to which the storage systems are connected.
- All reserve volumes are converted to global enqueues (ENQs).
- The storage system volumes are count key data (CKD) volumes that are attached to a z/OS system. The volumes can be attached to the same z/OS system or different systems.
- If the storage system volumes are not connected to a z/OS system on which Tivoli Storage Productivity Center for Replication is installed, the required Resource Access Control Facility (RACF) settings are set on the z/OS system and the **SOCKPORT** parameter is defined for the HyperSwap management address space IOSHMCTL. This set up enables Tivoli Storage Productivity Center for Replication to connect to the z/OS system by using an IP connection. For the required RACF settings, see the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

Complete the following tasks to set up an environment in which you can run HyperSwap processing:

1. Add a connection for each z/OS system in the Tivoli Storage Productivity Center for Replication GUI. To add a connection, open the Host Systems page and click **Add Host Connection**. If the z/OS system is the system on which Tivoli Storage Productivity Center for Replication is installed, a native connection is automatically added in the GUI. If the z/OS system does not have Tivoli Storage Productivity Center for Replication installed, you must you use an IP connection to connect to the z/OS system.
2. Add a connection to the storage systems that are attached to the z/OS systems. To add a connection, open the Storage Systems page and click **Add Storage Connection**. You can connect to a storage system by using a Hardware Management Console (HMC), direct, or z/OS connection.
3. Create Basic HyperSwap, Metro Mirror with Failover/Failback, Metro Global Mirror, or Metro Global Mirror with Practice sessions and enable HyperSwap. To create a session, open the Sessions page and click **Create Session**.

Setting up the environment for Open HyperSwap:

You must set up an environment in which you can run Open HyperSwap processing before you can enable Open HyperSwap for a IBM Tivoli Storage Productivity Center for Replication session.

Complete the following steps before you enable Open HyperSwap.

1. Ensure that the IBM AIX hosts and IBM System Storage DS8000 meet the following hardware and software requirements:

   AIX requirements

   You can find the supported AIX version for each Tivoli Storage Productivity Center for Replication release in the support matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446. Click the link for the applicable release under Agents, Servers, and GUI.

   You must have the following AIX modules installed:

   - Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
   - Multi-Path Input/Output (MPIO) module (the version that is provided with AIX)

System Storage DS8000 hardware requirements
The Open HyperSwap function requires System Storage DS8000 5.1 or later.

You cannot use Open HyperSwap with High Availability Cluster Multi-Processing (HACMP™).

2. Create connections from Tivoli Storage Productivity Center for Replication to the AIX hosts by following the instructions at "Adding a host system connection" on page 659.
3. Assign copy set volumes from the storage device to the host by using the System Storage DS8000 command-line interface (CLI) or the graphical user interface (GUI).
4. Run the AIX **cfgmgr** command to discover the volumes that are assigned to the host.

Considerations for Open HyperSwap and the AIX host:

- A single session that has the Open HyperSwap function enabled can manage multiple hosts; however, each host can be associated with only one session. Multiple hosts can share the session.
- For AIX 5.3, a single host can manage a maximum of 1024 devices that are enabled for Open HyperSwap processing on the host. A maximum of eight logical paths can be configured for each copy set in the session. For AIX 6.1, a single host can manage a maximum of 1024 devices that are enabled for Open HyperSwap processing on the host. A maximum of 16 logical paths can be configured for each copy set in the session.
- If an application on the host opened a device, a Tivoli Storage Productivity Center for Replication session for that device cannot be ended. The **Terminate** command fails. To end the session, you must either close the application or remove the copy sets from the session. If you remove copy sets from the session, you must ensure that the application writes the data to the correct volume when the copy set relationship is restored.
- Open HyperSwap processing can fail on a subset of hosts for the session and continue to run on the remaining hosts for the same session. In this situation, you must determine the best action to take if the application is writing data to volumes on the source system and to volumes on the target system.
- To enable Open HyperSwap processing on the host, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

Configuring timers to support Open HyperSwap:

You can configure timeout values for the storage system, Tivoli Storage Productivity Center for Replication, and AIX hosts systems. These values can affect the operation of Open HyperSwap.

The following various timeout values can affect Open HyperSwap:

Storage system quiesce timeout value
This value is the time when the storage system starts a quiesce operation. When the timer value expires, I/O is resumed on the primary device. The default timeout value is 2 minutes, but the value can be set from 30 to 600 seconds. To set the quiesce timeout value, use the **chdev** command in the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

Storage system long busy timeout value
This value is the time in seconds that the logical subsystem (LSS) consistency group volume stays in the long busy state after a remote mirror and copy error is reported.

Timeout values for the applications that are on the host
The various applications that are running on the host have timeout values. The timeout values vary depending on the application.

Considerations for setting timers

Consider the following information for setting timers:

- If the host quiesce timer is set to a shorter value than the Tivoli Storage Productivity Center for Replication response timer, an I/O swap failure can occur. If a storage system triggers an unplanned failover and if the storage system quiesce timer expires before Tivoli Storage Productivity Center for Replication responds, the host attempts to write I/O to the primary volume where the loss of access occurred. If the hardware condition that caused the loss of access continues, the write I/O operation fails again and an unplanned Open HyperSwap is not performed.
- If the host quiesce timer is set to a longer value than the Tivoli Storage Productivity Center for Replication response timer, an application timeout might occur if Open HyperSwap takes too long to complete.

Enabling HyperSwap or Open HyperSwap for a session:

Enabling HyperSwap or Open HyperSwap for a session provides business recovery and business continuity.

Ensure that you can use the HyperSwap or Open HyperSwap functions in your environment. See "Setting up the environment for HyperSwap" on page 697 or "Setting up the environment for Open HyperSwap" on page 698.

To enable HyperSwap or Open HyperSwap for a session, complete the following steps:

1. In the Tivoli Storage Productivity Center for Replication navigation tree, select **Sessions**.
   a. If you are creating a session, click **Create Session**.
   b. If you are modifying a section, click **View/Modify Properties** in the **Select Action** menu, and then click **Go**.
2. On the properties page, select the following options:

   System or sysplex
   Select the z/OS system that is connected to the storage system or the sysplex that contains the system. A system or sysplex is required to use the **Manage H1-H2 with HyperSwap** option.

   Important: Tivoli Storage Productivity Center for Replication must be installed on or connected to one or more host systems to select a system or sysplex and to use the **Manage H1-H2 with HyperSwap** option.

   Manage H1-H2 with HyperSwap
   Select this option to trigger a HyperSwap operation, which redirects application I/O to the target volumes when there is a failure on the host accessible volumes. Tivoli Storage Productivity Center for

Replication uses HyperSwap to manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session.

On Configuration Error:

Partition the system(s) out of the sysplex
Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

Disable HyperSwap
Select this option to prevent a HyperSwap operation from occurring.

On Planned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)
Select this option to partition out the failing system and continue the swap processing on any remaining systems.

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

On Unplanned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)
Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and the HyperSwap operation does not complete.

Requirement: If you select this option, you must restart the system .

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

Manage H1-H2 with Open HyperSwap
Select this option to trigger an Open HyperSwap operation for volumes that are attached to an IBM AIX host. This option redirects application I/O to the target volumes when there is a failure on the host accessible volumes.

Tivoli Storage Productivity Center for Replication uses Open HyperSwap to manage the H1-H2 sequence of a Metro Mirror session.

Only volumes that are attached to host systems that are defined in the Tivoli Storage Productivity Center for Replication Host Systems panel are eligible for Open HyperSwap.

Restarting an AIX Host System that is enabled for Open HyperSwap:

When an AIX host system is restarted, the host automatically opens any volumes for I/O that were open before the system was restarted. If Open HyperSwap was

enabled for a set of volumes on the host system, the host must determine which storage system is the primary system before the host can allow the volumes to be opened.

If the Metro Mirror relationship for the set of volumes is in Prepared or Suspended state and the host has connectivity to both the primary and secondary storage systems, the host can determine through the hardware which storage system is the primary system. In this situation, the host automatically opens the volumes.

If the Metro Mirror relationship for the set of volumes is in Prepared state and the host has connectivity to only the secondary storage system, all I/O to the volumes might be blocked on the host system. The I/O might be blocked until the host can verify the primary volume in the relationship. The AIX `varyonvg` command fails to open the volumes for I/O to prevent the application from writing to an incorrect volume. If the host can determine which volume is the primary volume in the relationship and connectivity to the primary storage system is still lost, a Hyperswap event is triggered. This event causes all I/O to be automatically opened and directed to the secondary storage system.

If the Metro Mirror relationship for the set of volumes is in Target Available state or the host system does not have connectivity and is unable to determine which site is the primary site, all I/O to the volumes are blocked on the host system. The Target Available state occurs after a Hyperswap or a Recover command is issued for the session. The `varyonvg` command fails to open the volumes for I/O to prevent the application from writing to an incorrect volume.

Unblocking I/O on the host system after a host system restart

When any of the previous scenarios cause I/O to be blocked, manual actions might be necessary to remove the block.

If the relationships are in Target Available state on the hardware, issue the **Start** command to the session in the required direction of the relationship. This action defines the primary storage system for the host. The host system can allow the volumes to be opened for I/O processing.

If the relationships cannot be restarted, or the host cannot determine the primary storage system, it might be necessary to manually decouple the volumes on the host system.

To decouple the volumes, complete one of the following tasks:

- Stop the session or remove the copy set. This action requires a full copy when the relationships are restarted.
- Remove Object Data Manager (ODM) entries by using the following command:

  ```
  odmdelete -o Volume_Equivalency
  ```

  Restriction: This command deletes copy set information; use it only for this scenario.

Planned and unplanned swaps:

After a session is enabled for HyperSwap or Open HyperSwap processing and reaches Prepared state, Tivoli Storage Productivity Center for Replication loads the configuration of volumes that can be swapped to a z/OS or AIX operating system.

When the load is complete, the session can start a planned or unplanned swap. The H1-H2 role pair of the session shows a type of HS. An "H" is displayed over the connection in the dynamic image for that role pair, as shown in the following figure.

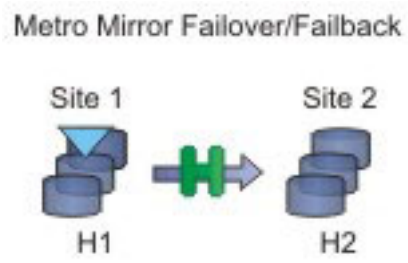

Completing a Planned Swap

When the session configuration is loaded on z/OS for HyperSwap or AIX for Open HyperSwap, the session is considered swap capable. There may be cases such as a planned maintenance or a migration from the primary storage, in which a planned swap might be required. When the session is in swap capable state, a planned swap can be completed by issuing the **HyperSwap** command for the session.

After you run a planned swap for z/OS HyperSwap and Open HyperSwap, the session transitions to Target Available state and all the H1-H2 pairs are in Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session has a status of Normal, which indicates that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session has a status of SEVERE. The active host on the session is shown as H2.

All I/O is redirected to the H2 volumes. After a successful swap to site 2, you cannot enable the copy function to site 2. The only way to restart the copy processing is by issuing the **Start H2->H1** command. To have the volumes protected with high availability and disaster recovery again, correct the error that caused the swap to occur, and manually restart the session to begin copying data to the other site.

The following figure illustrates a planned swap operation.

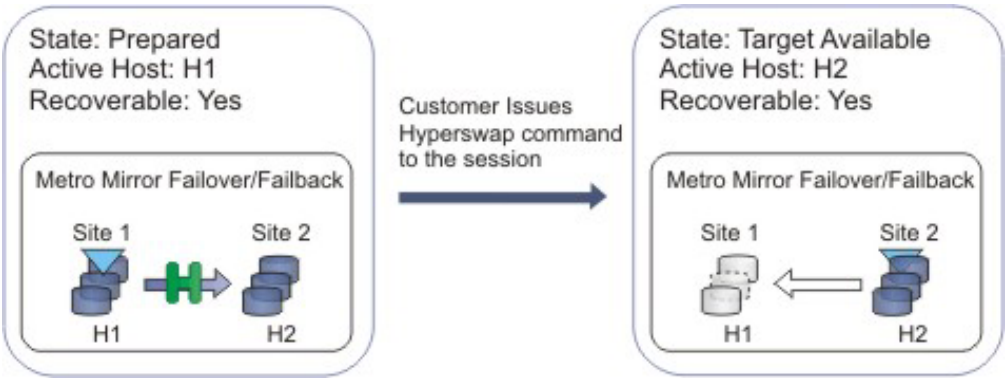

Unplanned swaps

When the session configuration is loaded on z/OS for HyperSwap processing or on AIX for Open HyperSwap processing, the session is considered swap capable. If a primary I/O error occurs, a swap occurs automatically. For HyperSwap, z/OS completes the swap and then alerts Tivoli Storage Productivity Center for Replication that a swap occurred. For Open HyperSwap, Tivoli Storage Productivity Center for Replication and the AIX host work together to complete the swap.

When an unplanned swap occurs for HyperSwap and Open HyperSwap, the session transitions to Target Available state and all the H1-H2 pairs are in Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session has a status of Normal and indicates that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session might show a status of SEVERE. The active host on the session is shown as H2.

All I/O is redirected to the H2 volumes. After a successful swap to site 2, you cannot enable the copy function to site 2. Therefore, it is not possible to issue a **Start H1->H2** command. The only way to restart the copy processing is by issuing a **Start H2->H1** command. To have the volumes protected with high availability and disaster recovery again, correct the error that caused the swap to occur, and manually restart the session to begin copying data to the other site.

The following figure illustrates an unplanned swap.

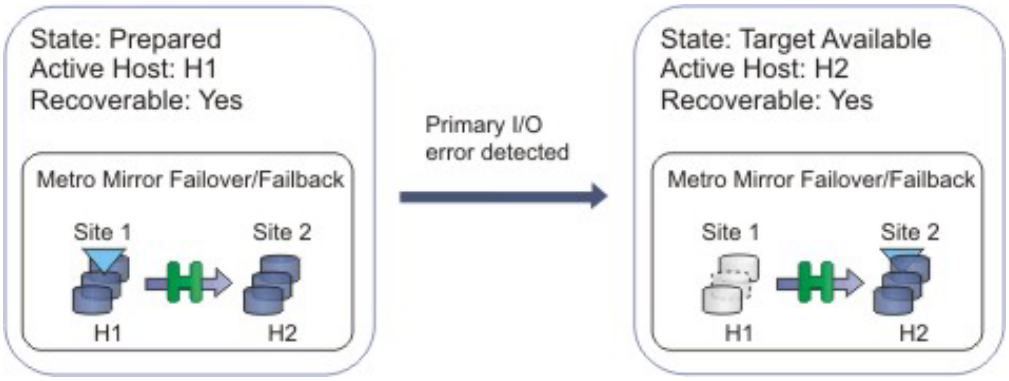

Verifying a session for a planned or unplanned swap:

You can verify whether a session is capable of a planned or unplanned swap from the z/OS console on HyperSwap or the AIX host on Open HyperSwap.

Verify the status of HyperSwap from the z/OS console:

1. View the overall status of the HyperSwap session. Issue the `ds hs,status` command. For example:

```
15.03.06 SYSTEM1 d hs,status
15.03.06 SYSTEM1 STC00063 IOSHM0303I HyperSwap Status 531
Replication Session: SR_HS
HyperSwap enabled
New member configuration load failed: Disable
Planned swap recovery: Disable
Unplanned swap recovery: Disable
FreezeAll: No
Stop: No
```

2. Verify all the volumes in the configuration. Issue the `ds hs,config(detail,all)` command. For example:

```
15.03.51 SYSTEM1 d hs,config(detail,all)
15.03.51 SYSTEM1 STC00063 IOSHM0304I HyperSwap Configuration 534
Replication Session: SR_HS
Prim. SSID UA DEV# VOLSER Sec. SSID UA DEV# Status
06 02 00F42 8K3602 06 04 00FA2
06 01 00F41 8K3601 06 03 00FA1
06 00 00F40 8K3600 06 02 00FA0
```

Verify the status of Open HyperSwap from the AIX host:

1. View the session association and the path that the I/O is being routed to. The path is indicated by an asterisk. Issue the `pcmpath query device` command. For example:

```
host1> pcmpath query device 14

  DEV#:   14  DEVICE NAME: hdisk14  TYPE: 2107900  ALGORITHM:  Load Balance
  SESSION NAME:     session1
  OS Direction:     H1<-H2
  ==========================================================================
  PRIMARY   SERIAL: 25252520000
  ----------------------------
  Path#      Adapter/Path Name          State     Mode     Select     Errors
     0           fscsi0/path0          CLOSE   NORMAL       6091          0
     1           fscsi0/path2          CLOSE   NORMAL       6300          0
     2           fscsi1/path4          CLOSE   NORMAL       6294          0
     3           fscsi1/path5          CLOSE   NORMAL       6187          0

  SECONDARY SERIAL: 34343430000 *
  ----------------------------
  Path#      Adapter/Path Name          State     Mode     Select     Errors
     4           fscsi0/path1          CLOSE   NORMAL      59463          0
     5           fscsi0/path3          CLOSE   NORMAL      59250          0
     6           fscsi1/path6          CLOSE   NORMAL      59258          0
     7           fscsi1/path7          CLOSE   NORMAL      59364          0
```

Temporarily disabling HyperSwap or Open HyperSwap:

In some situations, it might be necessary to temporarily disable the HyperSwap or Open HyperSwap function for a session.

You might want to disable HyperSwap or Open HyperSwap in the following circumstances:

- If you are performing maintenance on the system
- If one sysplex member cannot communicate with one or more volumes

Complete these steps to disable Open HyperSwap for a specific session:

1. In the Tivoli Storage Productivity Center for Replication navigation tree, select **Sessions**.
2. Select the sessions for which you want to disable the HyperSwap or Open HyperSwap function.
3. Select **View/Modify Properties** from the **Select Actions** list, and click **Go**.
4. Select **Disable HyperSwap** or **Disable Open HyperSwap**, and click **OK**.

   Tip: On management servers that run on a z/OS operating system, you can disable the HyperSwap function from an MVS command prompt window by entering the **SETHS DISABLE** command.

Using active and standby Tivoli Storage Productivity Center for Replication servers with HyperSwap or Open HyperSwap:

You can set up an active and standby management server configuration with HyperSwap or Open HyperSwap.

Active and standby servers with HyperSwap

When the storage system is set up to connect through the z/OS interface, the connection information is sent to the standby server and a connection is attempted. The connection might fail if the standby server does not run on a z/OS operating system or does not have access to the same volumes. If the connection fails and the standby server resumes the processing, the standby server cannot manage the HyperSwap processing.

On the z/OS system, if the session configuration was successfully loaded before the HyperSwap processing, the z/OS system can manage the HyperSwap processing. If the z/OS system swaps the volumes but cannot communicate with the Tivoli Storage Productivity Center for Replication server, the session is set to Suspended or Severe state. To set the session to Target Available state, clear the **Manage H1-H2 with Hyperswap** option, and issue the **Recover** command for the session.

Active and standby servers with Open HyperSwap

When there is an active and standby management server configuration and a host system connection is added to the active server, a connection is attempted between the host system and the standby server. After AIX loads the session configuration, Open HyperSwap processing is possible only if there is communication between AIX and the Tivoli Storage Productivity Center for Replication server. If a standby server resumes the processing and is unable to connect to the host system that manages the swap, the session cannot run Open HyperSwap processing. You must activate communication to the host system before the session can run Open HyperSwap processing.

Related tasks:

"Managing management servers" on page 634
This section provides information about how to set up active and standby management servers, restore a lost connection between the management servers, or perform a takeover on the standby management server.

Session commands

The commands that are available for a session depend on the session type.

Commands are issued synchronously to IBM Tivoli Storage Productivity Center for Replication sessions. Any subsequent command issued to an individual session is not processed until the first command has completed. Some commands can take an extended amount of time to complete, such as the Start command because it sets up the hardware. The GUI continues to allow you to issue commands to other sessions and not hold up functionality. When a command has completed, the console displays the results of the command.

Basic HyperSwap commands

Use this information to learn about commands available for Basic HyperSwap sessions.

Note: Individual suspend and recover commands are not available in HyperSwap.

Table 100. Basic HyperSwap commands

Command	Action
HyperSwap	Triggers a HyperSwap where I/O is redirected from the source volume to the target volume, without affecting the application using those volumes. You can use this command if you want to perform maintenance on the original source volumes.
Start H1->H2	Starts copying data synchronously from H1 to H2 in a Metro Mirror session. **Note:** A session might go into a Severe state with error code 1000000 before the session returns to Normal/Prepared State and HyperSwap Capable. The duration of the Severe state depends on how large of a session is running.
Start H2->H1	Starts copying data synchronously from H2 to H1 in a Metro Mirror session. You can issue this command only after the session has been swapped and the production site is H2. To enable data protection when the H1 volumes are available again, start I/O to the H2 volumes, and issue this command to replicate data from the H2 volumes to H1 volumes.
Stop	Suspends updates to all the targets of pairs in a session. You can issue this command at any time during an active session. **Note:** After you issue the stop command, targets might not be consistent.
Terminate	Removes all physical copies and relationships from the hardware during an active session.

FlashCopy commands

Use this information to learn about commands available for FlashCopy sessions.

Table 101. FlashCopy commands

Command	Action
Start	Performs any steps necessary to define the relationship before performing a FlashCopy operation. For ESS, DS6000, and DS8000, this command is not an option. Issue this command to put the session in the prepared state for the following storage systems: • SAN Volume Controller • Storwize V3500 • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified

Table 101. FlashCopy commands (continued)

Command	Action
Flash	Performs the FlashCopy operation using the specified options. Issue the Flash command to create a data consistent point-in-time copy for a FlashCopy Session with volumes on the following storage systems: • SAN Volume Controller • Storwize V3500 • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified For a FlashCopy session containing ESS, DS6000, and DS8000 volumes, the Flash command by itself is not sufficient to create a consistent copy. To create a consistent copy using the ESS, DS6000, and DS8000 Flash commands, you must quiesce application I/O before issuing the Flash command.
InitiateBackgroundCopy	Copies all tracks from the source to the target immediately, instead of waiting until the source track is written to. This command is valid *only* when the background copy is not already running.
Terminate	Removes all active physical copies and relationships from the hardware during an active session. If you want the targets to be data consistent before removing their relationship, you must issue the **InitiateBackgroundCopy** command if NOCOPY was specified, and then wait for the background copy to complete by checking the copying status of the pairs.

Snapshot commands

Use this information to learn about commands that are available for Snapshot sessions and groups. A snapshot group is a grouping of snapshots of individual volumes in a consistency group at a specific point in time.

Table 102. Snapshot session commands

Command	Action
Create Snapshot	Creates a snapshot of the volumes in the session
Restore	Restores the H1 volumes in the session from a set of snapshot volumes. You must have at least one snapshot group to restore from. When you issue this command in the Tivoli Storage Productivity Center for Replication graphical user interface (GUI), you are prompted to select the snapshot group.

Table 103. Snapshot group commands

Command	Action
Delete	Deletes the snapshot group and all the individual snapshots that are in the group from the session and from the XIV system. If the deleted snapshot group is the last snapshot group that is associated with the session, the session returns to the Defined state.

Table 103. Snapshot group commands (continued)

Command	Action
Disband	Disbands the snapshot group. When a snapshot group is disbanded, the snapshot group no longer exists. All snapshots in the snapshot group become individual snapshots that are no longer associated to the consistency group or the session. After a snapshot group is disbanded, it is no longer displayed in or managed by Tivoli Storage Productivity Center for Replication. If the disbanded snapshot group is the last snapshot group that is associated with the session, the session returns to the Defined state.
Duplicate	Duplicates the snapshot group. When a snapshot group is duplicated, a new snapshot group is created with new snapshots for all volumes that are in the duplicated group. The name of the duplicated snapshot group is generated automatically by the XIV system.
Lock	Locks a snapshot group. If the snapshot group is locked, write operations to the snapshots that are in the snapshot group are prevented. By default, a snapshot group is locked when it is created. This action is valid only if the snapshot group is unlocked.
Overwrite	Overwrites the snapshot group to reflect the data that is on the H1 volume.
Rename	Renames the snapshot group to a name that you provide. The name can be a maximum of 64 alphanumeric characters.
Restore	Restores the contents of a snapshot group by using another snapshot group in the session. Both of the snapshot groups must contain the same subset of volumes.
Set Priority	Sets the priority in which a snapshot group is deleted. The value can be the number 1 - 4. A value of 1 specifies that the snapshot group is deleted last. A value of 4 specifies that the snapshot group is deleted first.
Unlock	Unlocks a snapshot group. If the snapshot group is unlocked, write operations to the snapshots that are in the snapshot group are enabled and the snapshot group is displayed as modified. This action is valid only if the snapshot group is locked.

Metro Mirror commands

Use this information to learn about commands available for Metro Mirror sessions.

Table 104. Metro Mirror commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H2->H1** command becomes available.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H1->H2** command becomes available.

Table 104. Metro Mirror commands (continued)

Command	Action
HyperSwap	Triggers a HyperSwap where I/O is redirected from the source volume to the target volume, without affecting the application using those volumes. You can use this command if you want to perform maintenance on the original source volumes. This command is available only for Tivoli Storage Productivity Center for Replication for System z.
Start	Establishes a single-direction session with the hardware and begins the synchronization process between the source and target volumes.
Start H1->H2	Establishes Metro Mirror relationships between the H1 volumes and the H2 volumes, and begins data replication from H1 to H2.
Start H2->H1	Establishes Metro Mirror relationships between the H2 volumes and the H1 volumes and starts data replication from H2 to H1. Indicates direction of a failover and failback between two hosts in a Metro Mirror session. If the session has been recovered such that the production site is now H2, you can issue the **Start H2->H1** command to start production on H2 and start data replication.
Stop	Inconsistently suspends updates to all the targets of pairs in a session. This command can be issued at any point during an active session. **Note:** Targets after the suspend are not considered to be consistent.
StartGC	Establishes Global Copy relationships between the H1 volumes and the H2 volumes, and begins asynchronous data replication from H1 to H2. While in the Preparing state, it will not change to the Prepared state unless you switch to Metro Mirror.
Suspend	Causes all target volumes to remain at a data-consistent point and stops all data that is moving to the target volumes. This command can be issued at any point during a session when the data is actively being copied. **Note:** It is recommended that you avoid using the same LSS pairs for multiple Metro Mirror sessions. Metro Mirror uses a freeze command on ESS, DS6000, and DS8000 storage systems to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions are also suspended. When a **Suspend** command is issued to a source volume in an LSS that has source volumes in another active Metro Mirror session, the other source volumes are affected only if they have the same target LSS. The primary volumes are suspended, but volumes in the same source LSS that have target volumes in a different LSS are not affected because they use a different PPRC path connection.
Recover	Issues the **Recover** command to suspended sessions. This command performs the steps necessary to make the target available as the new primary site. Upon completion of this command, the session becomes Target Available.

Table 104. Metro Mirror commands (continued)

Command	Action
Terminate	Removes all copy relationships from the hardware during an active session. If you want the targets to be data consistent before removing their relationship, you must issue the **Suspend** command, then the **Recover** command, and then the **Terminate** command.

Metro Mirror with Practice commands

Use this information to learn about commands available for Metro Mirror with Practice sessions.

Table 105. Metro Mirror with Practice commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H2->H1** command becomes available.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H1->H2** command becomes available.
Flash	Creates a FlashCopy image from I2 volumes to H2 volumes. The amount of time for this to occur will vary depending on the number of copy sets in the session. **Note:** For ESS, DS6000, and DS8000 storage systems, the **Flash** command uses the freeze and thaw processing to create a data consistent point for the FlashCopy. If there is another Metro Mirror session overlapping on one or more of the same LSS pairs, that session will be suspended. It is also possible that the suspension of the other session might cause the Metro Mirror session to remain suspended after the flash command is issued instead of returning to Prepared state. Avoid using the same LSS pairs for multiple Metro Mirror sessions if possible.
Start H1->H2	Establishes a Metro Mirror relationship from the H1 volumes to the I2 volumes, and begins data replication.
Start H2->H1	Establishes a Metro Mirror relationship from H2 to H1 and begins data replication.
StartGC_H1H2	Distinguishes when the session is in the Preparing state from H1 to I2 and begins the asynchronous process between the source and target volumes. While in the Preparing state the session will not change to the Prepared state unless you switch to Metro Mirror.
StartGC_H2H1	Distinguishes when the session is in the Preparing state from H2 to H1 and begins the asynchronous process between the source and target volumes. While in the Preparing state the session will not change to the Prepared state unless you switch to Metro Mirror.

Table 105. Metro Mirror with Practice commands (continued)

Command	Action
Suspend	Causes all target volumes to remain at a data-consistent point and stops all data that is moving to the target volumes. This command can be issued at any point during a session when the data is actively being copied. **Note:** The Metro Mirror command uses a freeze command on the ESS, DS6000, and DS8000 devices to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions will also become suspended. Avoid using the same LSS pairs for multiple Metro Mirror sessions. When a **Suspend** command is issued to a source volume in an LSS that has source volumes in another active Metro Mirror session, the other source volumes are affected only if they have the same target LSS. The primary volumes are suspended, but volumes in the same source LSS that have target volumes in a different LSS are not affected because they use a different PPRC path connection.
Stop	Inconsistently suspends updates to all the targets of pairs in a session. This command can be issued at any point during an active session. **Note:** Targets after the suspend are not considered to be consistent.
Terminate	Terminates all copy relationships on the hardware.
Recover	Takes a point-in-time copy of the data on I2 to the H2 volumes, enabling the application to be attached and run from the H2 volumes on site 2. **Note:** The point-in-time copy is performed when the session is in a recoverable state to avoid that previous consistent data on H2 are overwritten by inconsistent data upon Recover. You can issue the **Flash** command if you want to force a point-in-time copy from I2 to JH2 volumes afterwards.

Global Mirror commands

Use this information to learn about commands available for Global Mirror sessions.

Table 106. Global Mirror commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H2->H1** command becomes available.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H1->H2** command becomes available.

Table 106. Global Mirror commands (continued)

Command	Action
Recover	Issue this command to recover the session to the target site. This command performs the steps necessary to make the target host volumes consistent and available for access as the new primary site. Upon completion of this command, the session state becomes Target Available. Do not access H2 volumes until the **Recover** command is completed and the session displays Target Available and Recoverable. A Recover to H2 also establishes a point-in-time copy to J2 to preserve the last consistent data.
Start	Establishes all relationships in a single-direction session and begins the process necessary to start forming consistency groups on the hardware.
Start H1->H2	Starts copying data from H1 to H2 in a Global Mirror Failover/Failback session. Establishes the necessary relationships in the session and begins the process necessary to start copying data from the H1 site to the H2 site and to start forming consistency groups.
Start H2->H1	Starts copying data from H2 to H1 in a failover and failback session for TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000. If a recover was performed on a session such that the production site is now H2, you can issue a **Start H2->H1** command to start moving data back to Site 1. However, this start does not provide consistent protection as it copies only asynchronously back because of the long distance. A Global Copy relationship is used. When you are ready to move production back to Site 1, issue a suspend to the session. This action puts the relationships into a synchronized state and suspends them consistently. Sessions are consistent when a H2->H1 copy occurs for the following storage systems: • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified
StartGC H1->H2 This command is available only for Global Mirror Failover/Failback and Global Mirror Failover/Failback with Practice sessions.	Establishes Global Copy relationships between site 1 and site 2 and begins asynchronous data replication from H1 to H2. To change the session state from Preparing to Prepared, you must issue the **Start H1->H2** command and the session must begin to form consistency groups. There is no disaster recovery protection for Global Copy relationships. If a disaster such as the loss of a primary storage system or a link failure between the sites occurs, the session might be inconsistent when you issue the **Recover** command. This command is available for the following storage systems: • TotalStorage Enterprise Storage Server Model 800 • System Storage DS6000 • System Storage DS8000

Table 106. Global Mirror commands (continued)

Command	Action
Suspend	Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H1 to H2 Global Copy pairs. **If you are using a System Storage DS8000:** The **Suspend** command invokes a pause command for System Storage DS8000. The pause command that is invoked depends on the System Storage DS8000 microcode level. If the following conditions are true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command **pausegmir -withsecondary**: • All primary volumes that are in the session are on a System Storage DS8000 storage system with a microcode level that includes the pause with secondary consistency option. To determine whether this option is available for your storage system, refer to the System Storage DS8000 documentation for the microcode level that you are using. • The pause with secondary consistency option is enabled on the storage system. If these conditions are not true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command **pausegmir**. Both commands temporarily pause the formation of consistency groups after the current consistency group is formed. However, the command for a pause with secondary consistency creates a consistent data set on the secondary volumes. If you do not want to suspend the Global Copy pairs when a Global Mirror session is suspended, open the *TPC_install_directory*\wlp\usr\servers\replicationServer\properties\rmserver.properties file and add the following property: `csm.server.sus_gc_pairs_on_gm_pause = false` **Tip:** The pause with secondary consistency option requires that the Global Copy pairs are suspended. If you do not want to use the pause with secondary consistency option for storage systems on which the option is enabled, set the **csm.server.sus_gc_pairs_on_gm_pause** property to **false**.
Terminate	Removes all physical copies and relationships from the hardware during an active session. If you want the targets to be data consistent before you remove their relationship, you must issue the **Suspend** command, the **Recover** command, and then the **Terminate** command.

Global Mirror with Practice commands

Use this information to learn about commands available for Global Mirror with Practice sessions.

Table 107. Global Mirror with Practice commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H2->H1** command becomes available.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H1->H2** command becomes available.
Flash	The Flash command ensures that all I2s are consistent, and then flashes the data from I2 to the H2 volumes. After the flash is complete, the Global Mirror session is automatically restarted, and the session begins forming consistency groups on I2. You can then use the H2 volumes to practice your disaster recovery procedures. For Global Mirror with Practice sessions for TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000, the Flash command temporarily pauses the formation of consistency groups.
Recover	Restores consistent data on I2 volumes and takes a point-in-time copy of the data on I2 to the H2 volumes, enabling the application to be attached and run from the H2 volumes on site 2. The I2 volumes continue to hold the consistent data and can be flashed again to H2 by using the **Flash** command.
Start H1->H2	Starts copying data from H1 to H2. After the first pass of the copy is complete for all pairs, the session establishes the I2->J2 FlashCopy pairs. The session starts the Global Mirror master so that the hardware begins to form consistency groups, to ensure that consistent data is at site 2.
Start H2->H1	Starts copying data from H2 to H1 in a failover and failback session. If a recover was performed on a session such that the production site is now H2, you can issue a **Start H2->H1** command to start moving data back to Site 1. However, this start does not provide consistent protection as it copies only asynchronously back because of the long distance. **Note:** TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 volumes are not consistent for the **Start H2->H1** command.A Global Copy relationship is used. When you are ready to move production back to Site 1, issue a suspend command to the session. The relationships are put into a synchronized state and are suspended consistently.

Table 107. Global Mirror with Practice commands (continued)

Command	Action
StartGC H1->H2 This command is available only for Global Mirror Failover/Failback and Global Mirror Failover/Failback with Practice sessions.	Establishes Global Copy relationships between site 1 and site 2 and begins asynchronous data replication from H1 to I2. To change the session state from Preparing to Prepared, you must issue the **Start H1->H2** command and the session must begin to form consistency groups. There is no disaster recovery protection for Global Copy relationships. If a disaster such as the loss of the primary Tivoli Storage Productivity Center for Replication server occurs, the session might be inconsistent when you issue the **Recover** command. This command is available for the following storage systems: • TotalStorage Enterprise Storage Server Model 800 • System Storage DS6000 • System Storage DS8000

Table 107. Global Mirror with Practice commands (continued)

Command	Action
Suspend	Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H1 to H2 Global Copy pairs. **If you are using a System Storage DS8000:** The **Suspend** command invokes a pause command for System Storage DS8000. The pause command that is invoked depends on the System Storage DS8000 microcode level. If the following conditions are true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command **pausegmir -withsecondary**: • All primary volumes that are in the session are on a System Storage DS8000 storage system with a microcode level that includes the pause with secondary consistency option. To determine whether this option is available for your storage system, refer to the System Storage DS8000 documentation for the microcode level that you are using. • The pause with secondary consistency option is enabled on the storage system. If these conditions are not true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command **pausegmir**. Both commands temporarily pause the formation of consistency groups after the current consistency group is formed. However, the command for a pause with secondary consistency creates a consistent data set on the secondary volumes. If you do not want to suspend the Global Copy pairs when a Global Mirror session is suspended, open the *TPC_install_directory*\wlp\usr\servers\replicationServer\properties\rmserver.properties file and add the following property: `csm.server.sus_gc_pairs_on_gm_pause = false` **Tip:** The pause with secondary consistency option requires that the Global Copy pairs are suspended. If you do not want to use the pause with secondary consistency option for storage systems on which the option is enabled, set the **csm.server.sus_gc_pairs_on_gm_pause** property to **false**.
Terminate	Removes all physical copies and relationships that are on the hardware.

Metro Global Mirror commands

Use this information to learn about commands available for Metro Global Mirror sessions.

Table 108. Metro Global Mirror commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H2->H1->H3** command becomes available.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the **Start H1->H2->H3** command becomes available.
HyperSwap	Causes a site switch, equivalent to a suspend and recover for a Metro Mirror with failover and failback individual suspend and recover commands are not available. This command is available only for Tivoli Storage Productivity Center for Replication for System z.
RecoverH1	Specifying H1 makes the H1 volume TargetAvailable. Metro Global Mirror (when running H2->H1->H3) can move production to either the H1 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore, the site designation is added to the **Recover** command so IBM Tivoli Storage Productivity Center for Replication can set up for the failback.
RecoverH2	Specifying H2 makes the H2 volume TargetAvailable. Metro Global Mirror (when running H1->H2->H3) can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing is different depending on the recovery site. Therefore the site designation is added to the **Recover** command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.
RecoverH3	Specifying H3 makes the H3 volume TargetAvailable. Metro Global Mirror (when H1->H2->H3 is run) can then move production to the H3 set of volumes. Because Tivoli Storage Productivity Center for Replication processing differs depending on the recovery site, the site designation is added to the **Recover** command so that Tivoli Storage Productivity Center for Replication can prepare for the failback. This command prepares H3 so that you can start the application on H3. H3 becomes the active host, and you then have the option start H3->H1->H2 to perform a Global Copy copy back. The recovery establishes point-in-time copy to J3 volumes to preserve the last consistent data.
Re-enable Copy to Site 1	After you issue a **RecoverH1** command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
Re-enable Copy to Site 2	After you issue a **RecoverH2** command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
Re-enable Copy to Site 3	After you issue a **RecoverH3** command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.

Table 108. Metro Global Mirror commands (continued)

Command	Action
Start H1->H2->H3 This command is the Metro Global Mirror initial start command.	Establishes Metro Mirror relationships between H1 and H2, and Global Mirror relationships between H2 and H3. For Metro Global Mirror, this relationship includes the J3 volume to complete the Global Mirror configuration. (The J3 volume role is the journal volume at site 3.) The **Start H1->H2->H3** command can be used from some Metro Global Mirror configurations to transition back to the starting H1->H2->H3 configuration. This command is valid only when the session is in a defined, preparing, prepared, target available, or suspended state.
Start H1->H3	From the H1->H2->H3 configuration, this command changes the session configuration to a Global Mirror-only session between H1 and H3, with H1 as the source. Use this command in case of an H2 failure with transition bitmap support provided by incremental resynchronization. It can be used when a session is in preparing, prepared, and suspended states because there is not a source host change involved. You can use this command to bypass the H2 volume in case of an H2 failure and copy only the changed tracks and tracks in flight from H1 to H3. After the incremental resynchronization is performed, the session is running Global Mirror from H1 to H3 and loses the near-zero data loss protection that is achieved with Metro Mirror when H1->H2->H3 is run. However, data consistency is still maintained at the remote site with the Global Mirror solution. From H2->H1->H3 configuration, this command changes the session configuration to a Global Mirror-only session configuration between H1 and H3, with H1 as the source. Use this command when the source site has a failure and production is moved to the H1 site. For example, for an unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when you restart the H1->H3 configuration or from the TargetAvailable H2->H1->H3 state.
Start H2->H3	From the H1->H2->H3 configuration, this command moves the session configuration to a configuration between H2 and H3, with H2 as the source. Use this command when the source site has a failure and production is moved to the H2 site. For example, for an unplanned HyperSwap. The Global Mirror session is continued. This session is a host-volume change so this command is valid only when you restart the H1->H3 configuration or from the TargetAvailable H2->H1>H3 state. From the H2->H1->H3 configuration, this command changes the session configuration to a configuration between H2 and H3 with H2 as the source. Use this command in case of an H1 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source-host change involved, this command can be used when the session is in the preparing, prepared, and suspended states. The **Start H2->H1->H3** command can be used to transition back to the starting H2->H1->H3 configuration.

Table 108. Metro Global Mirror commands (continued)

Command	Action
Start H2->H1->H3 This command is the Metro Global Mirror start command.	This configuration completes the HyperSwap processing. This command creates Metro Mirror relationships between H2 and H1 and Global Mirror relationships between H1 and H3. For Metro Global Mirror, the J3 volume completes the Global Mirror configuration.
Start H3->H1->H2	After a recovery to H3 is completed, this command sets up the hardware to allow the application to begin writing to H3, and the data is copied back to H1 and H2. However, issuing this command does not guarantee consistency in the case of a disaster because only Global Copy relationships are established to cover the long-distance copy back to site 1. To move the application back to H1, you can issue a suspend while in this state to drive all the relationships to a consistent state and then issue a freeze to make the session consistent. You can then issue a **Rcover** followed by a **Start H1->H2->H3** to go back to the original configuration.

Table 108. Metro Global Mirror commands (continued)

Command	Action
SuspendH2H3	Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H2 to H3 Global Copy pairs. **If you are using a System Storage DS8000:** The **Suspend** command invokes a pause command for System Storage DS8000. The pause command that is invoked depends on the System Storage DS8000 microcode level. If the following conditions are true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command **pausegmir -withsecondary**: • All primary volumes that are in the session are on a System Storage DS8000 storage system with a microcode level that includes the pause with secondary consistency option. To determine whether this option is available for your storage system, refer to the System Storage DS8000 documentation for the microcode level that you are using. • The pause with secondary consistency option is enabled on the storage system. If these conditions are not true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command **pausegmir**. Both commands temporarily pause the formation of consistency groups after the current consistency group is formed. However, the command for a pause with secondary consistency creates a consistent data set on the secondary volumes. If you do not want to suspend the Global Copy pairs when a Global Mirror session is suspended, open the *TPC_install_directory*\wlp\usr\servers\replicationServer\properties\rmserver.properties file and add the following property: `csm.server.sus_gc_pairs_on_gm_pause = false` **Tip:** The pause with secondary consistency option requires that the Global Copy pairs are suspended. If you do not want to use the pause with secondary consistency option for storage systems on which the option is enabled, set the **csm.server.sus_gc_pairs_on_gm_pause** property to **false**.
SuspendH1H3	Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H1 to H3 Global Copy pairs. The information that is specific to System Storage DS8000 in the previous description for the **SuspendH2H3** command is also applicable to this command.

Metro Global Mirror with Practice commands

Use this information to learn about commands available for Metro Global Mirror with Practice sessions.

Table 109. Metro Global Mirror with Practice commands

Command	Action
Enable Copy to Site 1	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session.
Enable Copy to Site 2	Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session.
Flash	This command is available in the following states: • Target Available state when the active host is H3 **Note:** Use this command if the FlashCopy portion of the Recover command from I3 to H3, fails for any reason. The problem can be addressed; and a **Flash** command can be issued to complete the flash of the consistent data from I3 to H3. • Prepared state when the active host is H1 and data is copying H1 to H2 to I3, or the active host is H2 and data is copying H2 to H1 to H3. • Prepared state when the active host is H2 and data is copying H2 to I3. • Prepared state when the active host is H1 and data is copying H1 to I3. Use this command if the FlashCopy portion of the Recover command from I3 to H3, fails for any reason. The problem can be addressed; and a **Flash** command can be issued to complete the flash of the consistent data from I3 to H3. Issuing a **Flash** command on a Global Mirror Practice session for ESS, DS6000, and DS8000 temporarily pauses the formation of consistency groups. Ensure that all I3s are consistent, and then flash the data from I3 to the H3 volumes. After the flash is complete, the Global Mirror session will be automatically restarted, and the session will begin forming consistency groups on I3. You can then use the H3 volumes to practice your disaster recovery procedures.
HyperSwap	Causes a site switch, equivalent to a suspend and recover for a Metro Mirror with failover and failback individual suspend and recover commands are not available. This command is available only for Tivoli Storage Productivity Center for Replication for System z.
RecoverH1	Specifying H1 makes the H1 volume TargetAvailable. When H1->H2->H3 is run, Metro Global Mirror can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore the site designation is added to the **Recover** command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback. The FlashCopy creates a consistent copy of the data on the H3 volumes so that an application can recover to those volumes and begin writing I/O. When the FlashCopy is complete, the session reaches a Target Available state, and you can attach your volumes on Site 3.

Table 109. Metro Global Mirror with Practice commands (continued)

Command	Action
RecoverH2	Specifying `H2` makes the H2 volume TargetAvailable. When H1->H2->H3 is run, Metro Global Mirror can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore the site designation is added to the **Recover** command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.
RecoverH3	Specifying `H3` makes the H3 volume the TargetAvailable. When H1->H2->H3 is run, Metro Global Mirror can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site; therefore, the site designation is added to the **Recover** command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.
Re-enable Copy to Site 1	After you issue a **RecoverH1** command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
Re-enable Copy to Site 2	After you issue a **RecoverH2** command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
Re-enable Copy to Site 3	After you issue a **RecoverH3** command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.
Start H1->H2->H3 This command is the Metro Global Mirror initial start command.	This command creates Metro Mirror relationships between H1 and H2, and Global Mirror relationships between H2 and H3. For Metro Global Mirror, this relationship includes the J3 volume to complete the Global Mirror configuration. (The J3 volume role is the journal volume at site 3.) The **Start H1->H2->H3** command can be used from some Metro Global Mirror configurations to return to the starting H1>H2>H3 configuration. This command is valid only when the session is in a defined, preparing, prepared, target available, or suspended state.

Table 109. Metro Global Mirror with Practice commands (continued)

Command	Action
Start H1->H3	From the H1->H2->H3 configuration, this command changes the session configuration to a Global-Mirror-only session between H1 and H3, with H1 as the source. Use this command in case of an H2 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source host change involved, it can be used when a session is in preparing, prepared, and suspended states. You can use this command to bypass the H2 volume in case of an H2 failure and copy only the changed tracks and tracks in flight from H1 to H3. After the incremental resynchronization is completed, the session is running Global Mirror from H1 to H3 and loses the near-zero data loss protection that is achieved with Metro Mirror when H1->H2->H3 is run. However, data consistency is still maintained at the remote site with the Global Mirror solution. From H2->H1->H3 configuration, this command changes the session configuration to a Global-Mirror-only session configuration between H1 and H3, with H1 as the source. Use this command when the source site has a failure and production is moved to the H1 site. For example, for an unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when you restart the H1->H3 configuration or from the TargetAvailable H2->H1->H3 state.
Start H2->H3	From the H1->H2->H3 configuration, this command moves the session configuration to a configuration between H2 and H3, with H2 as the source. Use this command when the source site has a failure and production is moved to the H2 site. For example, for an unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when you restart the H1>H3 configuration or from the TargetAvailable H2->H1->H3 state. From the H2->H1->H3 configuration, this command changes the session configuration to a configuration between H2 and H3 with H2 as the source. Use this command in case of an H1 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source-host change involved, this command can be used when the session is in the preparing, prepared, and suspended states. The **Start H2->H1->H3** command can be used to return to the starting H2->H1->H3 configuration.
Start H2->H1->H3 This command is the Metro Global Mirror start command.	This configuration completes the HyperSwap processing. This command creates Metro Mirror relationships between H2 and H1 and Global Mirror relationships between H1 and H3. For Metro Global Mirror, the J3 volume completes the Global Mirror configuration.

Table 109. Metro Global Mirror with Practice commands (continued)

Command	Action
Start H3->H1->H2	After a recovery to H3 is completed, this command sets up the hardware to allow the application to begin writing to H3, and the data is copied back to H1 and H2. However, issuing this command does not guarantee consistency in the case of a disaster because only Global Copy relationships are established to cover the long-distance copy back to site 1. To move the application back to H1, you can issue a suspend while in this state to drive all the relationships to a consistent state and then issue a freeze to make the session consistent. You can then issue a **Rcover** followed by a **Start H1->H2->H3** to go back to the original configuration.
SuspendH2H3	Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H2 to H3 Global Copy pairs. **If you are using a System Storage DS8000:** The **Suspend** command invokes a pause command for System Storage DS8000. The pause command that is invoked depends on the System Storage DS8000 microcode level. If the following conditions are true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command **pausegmir -withsecondary**: • All primary volumes that are in the session are on a System Storage DS8000 storage system with a microcode level that includes the pause with secondary consistency option. To determine whether this option is available for your storage system, refer to the System Storage DS8000 documentation for the microcode level that you are using. • The pause with secondary consistency option is enabled on the storage system. If these conditions are not true, the **Suspend** command invokes a command that is equivalent to the System Storage DS8000 command **pausegmir**. Both commands temporarily pause the formation of consistency groups after the current consistency group is formed. However, the command for a pause with secondary consistency creates a consistent data set on the secondary volumes. If you do not want to suspend the Global Copy pairs when a Global Mirror session is suspended, open the *TPC_install_directory*\wlp\usr\servers\replicationServer\properties\rmserver.properties file and add the following property: `csm.server.sus_gc_pairs_on_gm_pause = false` **Tip:** The pause with secondary consistency option requires that the Global Copy pairs are suspended. If you do not want to use the pause with secondary consistency option for storage systems on which the option is enabled, set the **csm.server.sus_gc_pairs_on_gm_pause** property to **false**.

Table 109. Metro Global Mirror with Practice commands (continued)

Command	Action
SuspendH1H3	Pauses the Global Mirror master session, which causes the session to stop forming consistency groups, and suspends the H1 to H3 Global Copy pairs. The information that is specific to System Storage DS8000 in the previous description for the **SuspendH2H3** command is also applicable to this command.
Terminate	This command terminates all copy relationships on the hardware.

Site awareness

You can associate a location with each storage system and each site in a session. This *site awareness* ensures that only the volumes whose location matches the location of the site are allowed for selection when you add copy sets to the session. This prevents a session relationship from being established in the wrong direction.

Note: To filter the locations for site awareness, you must first assign a site location to each storage system.

IBM Tivoli Storage Productivity Center for Replication does not perform automatic discovery of locations. Locations are user-defined and specified manually.

You can change the location associated with a storage system that has been added to the IBM Tivoli Storage Productivity Center for Replication configuration. You can choose an existing location or add a new one. Locations are deleted when there is no longer a storage system with an association to that location.

When adding a copy set to a session, a list of candidate storage systems is presented, organized by location. Storage systems that do not have a location are displayed and available for use when you create a copy set.

You can also change the location for any site in a session. Changing the location of a session does not affect the location of the storage systems that are in the session.

Changing the location of a storage system might have consequences. When a session has a volume role with a location that is linked to the location of the storage system, changing the location of the storage system could change the session's volume role location. For example, if there is one storage system with the location of A_Location and a session with the location of A_Location for its H1 role, changing the location of the storage system to a different location, such as B_Location, also changes the session's H1 location to Site 1. However, if there is a second storage system that has the location of A_Location, the session's role location is not changed.

Important: Location matching is enabled only when adding copy sets. If you change the location of a storage system or volume role, IBM Tivoli Storage Productivity Center for Replication does not audit existing copy sets to confirm or deny location mismatches.

Preserve Mirror option

This topic presents recommendations for using the Preserve Mirror option in FlashCopy and Metro Mirror sessions.

When the source of the FlashCopy relationship is a source of a Metro Mirror relationship, and the target of the FlashCopy relationship is the source of a Metro Mirror relationship, the Preserve Mirror option attempts to preserve the consistency of the Metro Mirror relationship at the target of the FlashCopy relationship, preventing a full copy from being performed over the Metro Mirror link. Instead, parallel flashes will be performed (if possible) on both sites. If the consistency cannot be preserved, the Flash for the FlashCopy relationships will fail, and the data of the Metro Mirror relationship at the target of the FlashCopy relationship will not be changed.

Note: This option is available only on DS8000 storage devices with the required code levels installed.

However, in some instances, the Preserve Mirror option can cause a Metro Mirror session to go into a Preparing state, or even a Suspended state. This topic describes the recommended usage of the Preserve Mirror feature. Using this feature in other ways might lead to a Metro Mirror session going into a Preparing or Suspended state.

FlashCopy session

You can use the Preserve Mirror option in FlashCopy sessions in two different methods:

Perform an incremental resynchronization
: To perform an incremental resynchronization, select the **Incremental** and **Persistent** options in the FlashCopy session: do *not* select the **No Copy** option.

Perform a single full copy
: To perform a single full copy, ensure that the **Incremental**, **Persistent** and **No Copy** options are *not* selected before you issue a **Flash** command. If you use the **No Copy** option, issue either an **Initiate Background Copy** command or **Terminate** command before you issue the Flash command.

Refer to your DS8000 documentation for more information about the Preserve Mirror function.

Metro Mirror session

You can set up your Metro Mirror pairs in two different ways, depending on the level of consistency you need, and your preferences.

Note: For the examples in this section, the source pair is H1a->H2a and the target pair is H1b->H2b. The source pair will contain volumes that will be the source of the FlashCopy relationship and the target pair will always contain volumes that will be the target of the FlashCopy relationship.

Create one Metro Mirror session, and add the Metro Mirror pairs as copy sets to that session
: The benefit to this approach is that you do not need to worry about whether the host considers the H1a->H2a and H1b->H2b volumes to be consistent with one another. IBM Tivoli Storage Productivity Center for Replication will ensure that all of the volumes remain consistent.

  A drawback to this approach is that when using the **Attempt to preserve Metro Mirror consistency, but fail FlashCopy if Metro Mirror target consistency cannot be preserved** option (Preserve Mirror Required), there

is a chance that the target pair (H1b->H2b) might suspend unexpectedly: this causes all other pairs in the Metro Mirror session to suspend (including H1a->H2a). This can occur when a FlashCopy establish or withdraw fails unexpectedly on the remote (H1b->H2b) site. If the host requires the H1a->H2a and H1b->H2b volumes to be consistent, then you should suspend all other volumes.

Create one Metro Mirror session for the H1a->H2a volumes, and another Metro Mirror session for the H1b->H2b volumes

Use this option when the hosts and applications do not require the H1a->H2a and H1b->H2b volumes to be consistent with one another. In this case, you should create one Metro Mirror session for all of the H1a->H2a volumes, and another Metro Mirror session for the H1b->H2b volumes. The H1a->H2a pair is added to the first session, while the H1b->H2b pair is added to the second Metro Mirror session. As long as the host does not require consistency between the H1a and H1b volumes, this option benefits you when you use the **Attempt to preserve Metro Mirror consistency, but fail FlashCopy if Metro Mirror target consistency cannot be preserved** option (Preserve Mirror Required). The benefit is that if one pair is suspended (such as H1a->H2a), the pairs in the other session will not be affected, since it is in a different Metro Mirror session. Using this method, you can avoid the situation in which a critical application is writing to the source pair (H1a->H2a), while a batch job is writing to the target pair (H1b->H2b), and both pairs are in the same IBM Tivoli Storage Productivity Center for Replication session. These factors cause both applications to receive extended long busy signals, instead of just the batch job.

Creating sessions and adding copy sets

A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

Creating a FlashCopy session and adding copy sets

FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy session was established.

When you create a FlashCopy session for Global Mirror or Metro Global Mirror space-efficient target volumes, you must select **No Copy** for the FlashCopy session. With space-efficient volumes, you can use your FlashCopy repository more efficiently. Instead of requiring an equal amount of space to write data to, you can set aside a smaller amount of space in which to write data, where only the tracks that are changed are recorded. When your pool of storage is full, you can no longer perform a FlashCopy operation, and your session goes into a Severe state.

1. Follow these steps to create a FlashCopy session:
   a. In the navigation tree, select **Sessions**.
   b. Click **Create Session**.
   c. In the Create Session wizard, for the **Choose Hardware Type** list, select the item that shows the type of storage system for the session. If you want to add any of the following storage systems, click **Storwize Family**:
      - Storwize V3500
      - Storwize V3700

- Storwize V7000

d. In the **Choose Session Type** list, select **FlashCopy** and click **Next**.

e. On the Properties page, enter a session name and a description. Select the options that you want, as described in "Viewing session properties" on page 756. The options that are displayed depend on the storage system type. After you select the options that you want, click **Next**.

f. On the Site Locations page, select a location for Site 1 and click **Next**.

g. On the Results page, verify that the session was added successfully.

2. Follow these steps to add copy sets to the session:

a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.

b. In the Add Copy Sets wizard, complete the following information. The field names that are displayed depend on the storage system type. When you complete the information, click **Next**.

Storage system
Select a storage system. If the role has a location assigned to it, acceptable values for the storage system list are storage systems assigned to the same location as the role, and storage systems assigned to no location. In this case, the storage systems are grouped under different headings. If the role does not have a location, any storage system is acceptable.

Logical storage system or I/O Group
Select a logical storage system (LSS) or I/O group.

Volume
Select one volume or all volumes. The volumes are limited to the volumes within the LSS or I/O group that you selected.

The use of extent space-efficient volumes in volume roles is dependent on the capabilities of the storage system. Extent space-efficient volumes must be fixed block (FB). You cannot use count key data (CKD) volumes.

Session image
Shows an image that represents the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles are in the session and how the roles are distributed between the sites.

Volume Details
Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space-efficient.

Use a CSV file to import copy sets
Select this option to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file.

c. On the Choose Target page, select the target storage system, LSS or I/O group, and volume. Click **Next**.

d. On the Select Copy Sets page, select the copy sets that you want to add. You can click **Select All** to select all copy sets, **Deselect All** to select none of the copy sets, or **Add More** to add more copy sets to this session. Click **Next**.

e. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.

f. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

Related concepts:

"Sessions" on page 665
A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

"Session types" on page 53
Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

Creating a Snapshot session and adding copy sets

A Snapshot session is a session type that creates a point-in-time copy of a volume or set of volumes without having to define a specific target volume. Snapshot sessions are available only for the XIV system.

1. Follow these steps to create a Snapshot session:
   a. In the navigation tree, expand **Sessions**.
   b. Click **Create Session**.
   c. In the Create Session wizard, for the **Choose Hardware Type** list, select **XIV**.
   d. In the **Choose Session Type** list, select **Snapshot** and click **Next**.
   e. On the Properties page, type a session name and description, and click **Next**.
   f. On the Site Locations page, select a location for Site 1 and click **Next**.
   g. On the Results page, verify that the session was added successfully.
2. Follow these steps to add copy sets to the session:
   a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.
   b. From the **Host1 storage system** list, select the storage system that contains the volumes that you want to add.

      If the H1 role has an assigned location, only those storage systems that have the same location as the H1 role or storage systems that do not have a set location are displayed for selection.

      If the H1 role does not have an assigned location, all storage systems are displayed for selection.

      Storage systems that are assigned to a location are listed under the location name. Storage systems that are not assigned to a location are listed under the **None.** column.
   c. From the **Host1 pool** list, select the pool that contains the volumes.
   d. From the **Host1 volume** list, select the volumes. To select multiple volumes, press Ctrl or Shift and click the volumes in the list.
   e. If you want to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file. Click **Next**.
   f. On the Matching Results page, click **Next** if the match was successful.

g. On the Select Copy Sets page, select from the following options and click **Next**.

Select All
Click this button to select all of the copy sets in the table.

Deselect All
Click this button to clear all of the copy sets in the table.

Add More
Click this button to add another copy set to the list of copy sets to be created.

When you click **Add More**, you are returned to the Choose Host1 page of the wizard. The **Host1 storage system** and **Host1 pool** lists are populated with the values from the previously selected copy set. When you select the volumes for the various roles, the volumes are matched together and added to the list of copy sets on this page.

Selection check boxes
Select one or more copy sets that you want to create.

Host 1 Lists the volume IDs that are associated with the Host1 role. You can click the link to display information about the volume, including the full name, type, capacity, and whether the volume is protected and space efficient.

Copy Set
Displays the copy set information for the specified copy sets and any warning or error messages that are associated with the copy set.

A warning or icon next to the **Show** button indicates that you cannot create a copy set for the H1 volume. Click **Show** to view the message.

h. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.

i. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

Related concepts:

"Sessions" on page 665
A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

"Session types" on page 53
Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

Creating a Metro Mirror session and adding copy sets

A Metro Mirror session is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart.

1. Follow these steps to create a Metro Mirror session:
   a. In the navigation tree, select **Sessions**.
   b. Click **Create Session**.

c. In the Create Session wizard, for the **Choose Hardware Type** list, select the item that shows the type of storage system for the session. If you want to add any of the following storage systems, click **Storwize Family**:
   - Storwize V3500
   - Storwize V3700
   - Storwize V7000

d. In the **Choose Session Type** list, select a Metro Mirror session type and click **Next**.

e. On the Properties page, enter a session name and description. Depending on the session type and storage system, additional options are displayed. Select the options that you want, as described in "Viewing session properties" on page 756. After you select the options that you want on the Properties page, click **Next**.

f. On the Site Locations pages, select a location for Sites 1 and 2, and click **Next**.

g. On the Results page, verify that the session was added successfully.

2. Follow these steps to add copy sets to the session:

a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.

b. In the Add Copy Sets wizard, enter the following information for the Host pages. The field names that are displayed depend on the storage system type. When you complete the information on each page, click **Next**.

Storage system
Select a storage system. If the role has a location assigned to it, acceptable values for the storage system list are storage systems assigned to the same location as the role, and storage systems assigned to no location. In this case, the storage systems are grouped under different headings. If the role does not have a location, any storage system is acceptable.

Logical storage system, I/O Group, or Pool
Select a logical storage system (LSS), I/O group, or pool.

Volume
Select one volume or all volumes. The volumes are limited to the volumes within the LSS, I/O group, or pool that you selected.

The use of extent space-efficient volumes in volume roles is dependent on the capabilities of the storage system. Extent space-efficient volumes must be fixed block (FB). You cannot use count key data (CKD) volumes.

Session image
Shows an image that represents the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles are in the session and how the roles are distributed between the sites.

Volume Details
Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space-efficient.

Use a CSV file to import copy sets
Select this option to import copy sets from a comma-separated

value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file.

c. On the Select Copy Sets page, select the copy sets that you want to add. You can click **Select All** to select all copy sets, **Deselect All** to select none of the copy sets, or **Add More** to add more copy sets to this session. Click **Next**.

d. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.

e. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

Related tasks:

"Managing a session with HyperSwap and Open HyperSwap replication" on page 697

HyperSwap and Open HyperSwap provide high availability of data if a primary disk storage system failure occurs. When a failure occurs in writing I/O to the primary storage system, the failure is detected by IOS, and IOS automatically swaps the I/O to the secondary site with no user interaction and minimal application effect.

Creating a Global Mirror session and adding copy sets

A Global Mirror session is a method of asynchronous, remote data replication between two sites that are over 300 KM apart.

1. Follow these steps to create a Global Mirror session:
   a. In the navigation tree, select **Sessions**.
   b. Click **Create Session**.
   c. In the Create Session wizard, for the **Choose Hardware Type** list, select the item that shows the type of storage system for the session. If you want to add any of the following storage systems, click **Storwize Family**:
      - Storwize V3500
      - Storwize V3700
      - Storwize V7000
   d. In the **Choose Session Type** list, select a Global Mirror session type and click **Next**.
   e. On the Properties page, enter a session name and description. Depending on the session type and storage system, additional options are displayed. Select the options that you want, as described in "Viewing session properties" on page 756. After you select the options that you want on the Properties page, click **Next**.
   f. On the Site Locations pages, select a location for Sites 1 and 2, and click **Next**.
   g. On the Results page, verify that the session was added successfully.
2. Follow these steps to add copy sets to the session:
   a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.
   b. In the Add Copy Sets wizard, complete the following information on the host, journal, intermediate, or change volume pages. The field names that are displayed depend on the storage system type. When you complete the information on each page, click **Next**.

      Storage system
      Select a storage system from this list. If the volume role has a site location that is assigned to it, you can select a storage system that is

assigned to the same location as the role or a storage system that is not assigned to a location. Storage systems that have a location are listed under the location. Storage systems that do not have a location are listed under **None**.

If the role does not have a site location, you can select any storage system.

Logical storage system, I/O Group, or Pool
Depending on the storage system that you selected, this field displays a logical storage system (LSS), I/O group, or pool. Select the LSS, I/O group, or pool that contains the volumes that you want to add to the copy set.

Volume
Select a volume or all volumes from this list. The volumes are limited to the volumes within the LSS, I/O group, or pool that you selected.

The use of extent space-efficient volumes in volume roles is dependent on the capabilities of the storage system. Extent space-efficient volumes must be fixed block (FB).

You cannot use count key data (CKD) volumes.

Session image
Shows an image that represents the session, in which the role for which you are selecting volumes is highlighted. This image shows how many roles there are in the session and how the roles are distributed between the sites.

Volume Details
Shows information about the selected volume including the volume name, full name, type, capacity, and whether the volume is protected and space efficient.

Use a CSV file to import copy sets
Select this option to import copy sets from a comma-separated value (CSV) file. Type the full path name of the CSV file in the text box, or click **Browse** to open a file dialog box and select the CSV file. A CSV file is a text file that is created in a spreadsheet program such as Microsoft Excel.

The first row of the CSV file is a comma-separated list of the volume roles for the session. The following rows specify each of the copy sets that you want to add to the session, with each volume separated by a comma. The following example file shows copy sets for a FlashCopy session:

```
H1,T1
DS8000:2107.FRLL1:VOL:1004,DS8000:2107.FRLL1:VOL:1104
DS8000:2107.FRLL1:VOL:1011,DS8000:2107.FRLL1:VOL:1101
DS8000:2107.FRLL1:VOL:1005,DS8000:2107.FRLL1:VOL:1105
```

Blank lines or lines that begin with # are ignored.

For further examples of CSV files, you can export files for existing sessions. On the Session page, click the session that you want, and then select **Export Copy Sets** from the **Select Action** menu.

c. On the Select Copy Sets page, select the copy sets that you want to add. You can click **Select All** to select all copy sets, **Deselect All** to select none of the copy sets, or **Add More** to add more copy sets to this session. Click **Next**.

d. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.

e. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

Related concepts:

"Sessions" on page 665
A *session* completes a specific type of data replication for a specific set of volumes. During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

"Session types" on page 53
Tivoli Storage Productivity Center for Replication provides several methods to replicate data. The type of data replication that is associated with a session is known as the *session type*.

"Copy sets" on page 665
During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a *copy set*.

Creating a Metro Global Mirror session and adding copy sets

A Metro Global Mirror session is a method of continuous, remote data replication that operates between three sites of varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source.

1. Follow these steps to create a Metro Global Mirror session:
   a. In the navigation tree, select **Sessions**.
   b. Click **Create Session**.
   c. In the Create Session wizard, for the **Choose Hardware Type** list, select the item that shows the type of storage system for the session. If you want to add any of the following storage systems, click **Storwize Family**:
      - Storwize V3500
      - Storwize V3700
      - Storwize V7000
   d. In the **Choose Session Type** list, select a Metro Global Mirror session type and click **Next**.
   e. On the Properties page, enter a session name and a description. Select the options that you want, as described in "Viewing session properties" on page 756. After you select the options that you want, click **Next**.
   f. On the Site Locations pages, select a location for Sites 1, 2, and 3, and click **Next**.
   g. On the Results page, verify that the session was added successfully.
2. Follow these steps to add copy sets to the session:
   a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.
   b. In the Add Copy Sets wizard, complete the following information for the Host and Journal pages. The field names that are displayed depend on the storage system type. When you complete the information on each page, click **Next**.

      Storage system
      Select a storage system. If the role has a location assigned to it,

acceptable values for the storage system list are storage systems assigned to the same location as the role, and storage systems assigned to no location. In this case, the storage systems are grouped under different headings. If the role does not have a location, any storage system is acceptable.

Logical storage system
Select a logical storage system (LSS).

Volume
Select a volume or all volumes from this list. The volumes are limited to the volumes within the LSS that you selected.

The use of extent space-efficient volumes in volume roles is dependent on the capabilities of the storage system. Extent space-efficient volumes must be fixed block (FB). You cannot use count key data (CKD) volumes.

Session image
Shows an image that represents the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles are in the session and how the roles are distributed between the sites.

Volume Details
Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space-efficient.

Use a CSV file to import copy sets
Select this option to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file.

c. On the Select Copy Sets page, select the copy sets that you want to add. You can click **Select All** to select all copy sets, **Deselect All** to select none of the copy sets, or **Add More** to add more copy sets to this session. Click **Next**.

d. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.

e. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

Using the Metro Mirror heartbeat

This topic provides information about Metro Mirror heartbeat, including how to enable and disable the heartbeat.

Metro Mirror heartbeat

The *heartbeat* is a Metro Mirror function. When the Metro Mirror heartbeat is disabled, data consistency across multiple storage systems is not guaranteed if the IBM Tivoli Storage Productivity Center for Replication management server cannot communicate with one or more storage systems. The problem occurs as a result of the Hardware Freeze Timeout Timer function within the storage system. If the controlling software loses connection to a storage system, the Metro Mirror relationships that it is controlling remains established and there is no way to *freeze* those pairs to create consistency across the multiple storage systems. When the freeze times out, dependent I/O is written to the target storage systems, which might corrupt data consistency. Freeze refers to a Metro Mirror (peer-to-peer remote copy [PPRC]) freeze function.

When determining whether to use the Metro Mirror heartbeat, analyze your business needs. Disabling the Metro Mirror heartbeat might result in data inconsistency. If you enable the Metro Mirror heartbeat and a freeze occurs, your applications will be unable to write during the freeze.

Metro Mirror heartbeat is disabled by default.

Metro Mirror heartbeat is not available for Metro Mirror with HyperSwap or Metro Global Mirror with HyperSwap.

There are two cases where lost communication between the coordination software (controller) and one or more storage systems can result in data consistency loss:

Freeze event not detected by a disconnected storage system

Consider a situation with four storage system machines in a primary site and four in a secondary site. One of the four storage systems on the primary loses the connection to the target site. This causes the affected storage system to prevent any writes from occurring, for a period determined by the Freeze timeout timer. At the same time, the affected storage controller loses communication with the controlling software and cannot communicate the Freeze event to the software.

Unaware of the problem, the controlling software does not issue the Freeze command to the remaining source storage systems. The freeze will stop dependent writes from being written to connected storage systems. However, once the Freeze times out and the long-busy is terminated, dependent write I/Os continue to be copied from the storage systems that did not receive the Freeze command. The Metro Mirror session remains in a state where one storage system has suspended copying while the other three storage systems are still copying data. This state causes inconsistent data on the target storage systems.

Freeze event detected, but unable to propagate the Freeze command to all storage systems

Consider a situation with four storage system machines in a primary site and four in a secondary site. One of the four storage systems on the primary loses the connection to the target site. This causes the affected storage system to issue long-busy to the applications for a period determined by the Freeze timeout timer. At the same time, one of the remaining three source systems loses communications with the controlling software.

The storage system that had an error writing to its target cannot communicate the Freeze event to the controlling software. The controlling software issues the Freeze command to all but the disconnected storage system (the one that lost communication with the software). The long-busy stops dependent writes from being written to the connected storage systems.

However, once the Freeze times out on the frozen storage system and the long-busy is terminated, dependent write I/Os continue to the target storage system from the source storage system that lost communication and did not receive the Freeze command. The Metro Mirror session remains in a state where three storage systems have suspended copying and one storage system is still copying data. This state causes inconsistent data on the target storage systems.

Before IBM Tivoli Storage Productivity Center for Replication V3.1, if the controlling software within a Metro Mirror environment detected that a managed storage system lost its connection to its target, the controlling software stopped all the other source systems to ensure consistency across all the targets. However, if the controlling software lost communication with any of the source subsystems during the failure, it could not notify those storage systems of the freeze event or ensure data consistency. The Metro Mirror heartbeat helps to overcome this problem. In a high-availability configuration, the Metro Mirror heartbeat is continued by the standby server after the Takeover command is issued on the standby, enabling you to perform actions on the standby server without causing a freeze.

IBM Tivoli Storage Productivity Center for Replication registers with the managed ESS 800, DS6000 or DS8000 storage systems within a Metro Mirror session when the start command is issued to the session. After this registration occurs, a constant heartbeat is sent to the storage system. If the storage system does not receive a heartbeat from the IBM Tivoli Storage Productivity Center for Replication management server within the allotted time (a subset of lowest LSS timeout value across all the source LSSs), the storage system initiates a freeze. If IBM Tivoli Storage Productivity Center for Replication did not successfully communicate with the storage system, it initiates a freeze on the remaining storage system after the allotted time is expired.

Note: It is recommended that you avoid using the same LSS pairs for multiple Metro Mirror sessions. Metro Mirror uses a freeze command on ESS, DS6000, and DS8000 storage systems to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions are also suspended.

When you are using the Metro Mirror heartbeat, be aware that:

- The Metro Mirror heartbeat can cause a single point of failure: if an error occurs on just the management server and not the storage system, a freeze might occur.
- When the Metro Mirror heartbeat timeout occurs, the storage system remains in a long busy state for the duration of the LSS freeze timeout.

Note: If Metro Mirror heartbeat is enabled for storage systems that are connected through a HMC connection, a connection loss might cause lost Metro Mirror heartbeats, resulting in Freeze actions with application I/O impact for configured Extended Long Busy timeout.

The Metro Mirror heartbeat is supported on storage systems connected though a TCP/IP (direct connect or HMC) connection. It is not supported on storage systems connected though a z/OS connection. Enabling the Metro Mirror heartbeat with a z/OS connection does not fail; however, a warning message is displayed specifying that the Metro Mirror heartbeat function does not work unless you have an IP connection.

If Metro Mirror heartbeat is enabled for storage systems that are connected through a TCP/IP (either direct connect or HMC) connection and z/OS connection, and the TCP/IP connection fails, IBM Tivoli Storage Productivity Center for Replication suspends the Metro Mirror session because there is no heartbeat through the z/OS connection.

If Metro Mirror heartbeat is enabled for storage systems that are connected through a TCP/IP connection and z/OS connection and you remove all TCP/IP

connections, IBM Tivoli Storage Productivity Center for Replication suspends the Metro Mirror sessions and the applications using those volume will be in Extended Long Busy timeout until the storage system's internal timeout timer expires. Ensure that you disable the Metro Mirror heartbeat for all Metro Mirror sessions before removing the last TCP/IP connection to avoid the Extended Long Busy timeout.

Enabling and disabling the Metro Mirror heartbeat

The Metro Mirror heartbeat guarantees data consistency across multiple storage systems when the IBM Tivoli Storage Productivity Center for Replication management server cannot communicate with one or more storage systems. The Metro Mirror heartbeat is disabled by default.

To enable the Metro Mirror heartbeat, perform the following steps:

1. In the navigation tree, select **Advanced Tools**.
2. To enable the Metro Mirror heartbeat, click **Enable Heartbeat**.
3. To disable the Metro Mirror heartbeat, click **Disable Heartbeat**.

Exporting copy set data

You can export data about all copy sets in a specific session, to maintain a backup copy that can be used to recover if you lose your session or upgrade to a different server.

Perform these steps to export the copy sets in a specific session:

1. In the navigation tree, select **Sessions**. The Sessions panel is displayed
2. Select the session for which you want to export copy sets.
3. Select **Export Copy Sets** from the **Actions** list, and click **Go**. The Export Copy Set wizard displays the status of the export and a link to the exported file.
4. Click that link and save the file to the local system.

   Important: You must save the file to your local system. After you close the panel, the data will be lost.
5. Click **Finish**.

Importing copy set data

You can import copy set data that was previously exported to a comma separated value (CSV) file.

Perform the following steps to import copy sets into an existing session:

1. In the navigation tree, select **Sessions**. The Session panel is displayed.
2. Select the session for which you want to import copy sets.
3. Select **Add Copy Sets** from the **Actions** list, and click **Go**. The Add Copy Sets wizard is displayed.
4. Select **Use a CSV file to import copy sets**.
5. Type the location and name of the CSV file to import, or use **Browse** to select the file. Then, click **Next**.
6. Verify that the matching results were successful, and then click **Next**.
7. Select the copy sets you want to add, and then click **Next**.
8. Confirm the number of copy sets that you want to create, and click **Next**. A progress bar displays.
9. Click **Next**.
10. Verify the matches, and click **Finish**.

Exporting historical data for Global Mirror role pairs

You can export data for a Global Mirror role pair that is in a TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, or System Storage DS8000 session to a comma-separated value (CSV) file. You can then use the data in the CSV file to analyze trends in your storage environment that affect your recovery point objective (RPO).

Complete these steps to export data to a CSV file:

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, click **Sessions**.
2. Click the Global Mirror session for which you want to export data.
3. Click **Export Global Mirror Data** from the **Select Action** menu, and click **Go**. The Export Historical Data for Global Mirror wizard is displayed.
4. Complete the following information, and then click **Next**.

   Select the role pair
   Select the role pair for the data that you want to show in the CSV file.

   Select the type of data
   Select the type of data that you want to export to a CSV file. Depending on the data type that you select, the CSV file contains data about the RPO or data about logical subsystem (LSS) out-of-sync tracks.

   To better analyze trends, you can create a CSV file with one data type and then create a file with the other data type. For example, the file that contains data for the RPO might show that the RPO threshold is often exceeded on a particular day and time. You can then view the file that contains data for logical subsystem (LSS) out-of-sync tracks to see whether a particular LSS or set of LSSs have high out-of-sync track values for that day and time.

   RPO Select this option to export data for the RPO. The data that is shown in the export file includes the average RPO for the dates that you select and information related to the formation of consistency groups.

   LSS Out-of-Sync Tracks
   Select this option to export data for the out-of-sync tracks that are in the LSSs.

   Start date
   Select the start date for the data that you want to show in the export file.

   By default, the date range maximum for RPO files is 31 days of data and the maximum for LSS out-of-sync track files is 7 days of data.

   End date
   Select the end date for the data that you want to show in the export file.

If the export was successful, a link to the CSV file is provided on the Results page. You can save the CSV file to your local system.

Modifying the location of session sites

You can change the location associated with each site of in session.

Prerequisites: You must have Administrator privileges to modify the location of a site.

Changing the location of a site in a session does not affect the location of the storage systems associated with that site.

Perform these steps to modify the location of a site:

1. In the navigation tree, select **Sessions**. The Sessions panel is displayed.
2. Select the session whose site locations you want to change.
3. Click **Modify Site Location(s)** from the actions list, and then click **Go**. The Modify Site Locations wizard is displayed.
4. Select the location for the site from the drop-down list, and then click **Next**. Repeat for each available site.

   To disable site awareness, set the location to `None`.

   Note: You can select only locations that have already been associated with one or more storage systems.
5. Click **Next**.
6. Click **Finish**.

Removing sessions

This topic describes how to remove sessions.

Important: You can remove only sessions that are in the Defined state.

Perform these steps to remove a session:

1. In the navigation tree, select **Sessions**.
2. Click the radio button next to the session you want to remove.
3. Select **Remove Session** from the **Actions** drop-down menu, and click **Go**.

Removing copy sets

This topic describes how to remove copy sets.

Perform these steps to remove a copy set:

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **Sessions**. Click the radio button next to the session that you want to remove copy sets from.
2. From the **Select Action** menu, select **Remove Copy Sets** and click **Go**. This starts the Remove Copy Sets wizard.
3. From the drop-down menus in the Remove Copy Sets wizard, select the Host 1 storage system, logical storage subsystem, and volume or select the all option. If you select all for a filter, the lower-level filter or filters are disabled. Click **Next**.
4. Select the copy sets that you want to remove and click **Next**.
5. The number of copy sets to be removed is displayed. Select the following options for removing the copy sets and click **Next**:
   - **Do you want to keep the base relationships on the hardware, but remove the copy sets from the session?**
     - **Yes**. This option specifies that the base relationships remain on the hardware, but the copy sets are removed from the Tivoli Storage Productivity Center for Replication session. This option supports scenarios

in which it might be best to leave the relationship on the hardware to avoid performing a full copy. For example, when you are migrating from one session type to another.

Only the base relationships (Metro Mirror, Global Copy, Snapshot, and FlashCopy) remain on the hardware. The relationships are removed from any consistency groups that are defined on the storage system.

- **No**. This option specifies that all relationships for the copy sets are removed from the hardware as well as the Tivoli Storage Productivity Center for Replication session. This option is the default.

- **If there are errors removing relationships on the hardware, do you want to force the copy sets to be removed from the session?**
  - **Yes**. This option forces the removal of copy sets despite any errors that occur when removing the relationships from the storage system. Once a forced removal is complete, any relationships that remain on the storage system for that copy set must be removed manually using the storage system interface.
  - **No**. This option does not force the removal of copy sets. This option enables you to correct the errors and try to remove the copy sets again. This option is the default.

6. After the copy sets are removed, click **Finish**.

Important: If an application on the host has opened a device, the copy sets in a Tivoli Storage Productivity Center for Replication session for that device are removed, but the copy sets remain coupled on the host. To decouple the copy sets, see the Troubleshooting and support section of theTivoli Storage Productivity Center for Replication Information Center at http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp.

Related concepts:

"Copy sets" on page 665
During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a *copy set*.

Migrating an existing configuration to Tivoli Storage Productivity Center for Replication

This topic describes how to convert an existing DS6000 or DS8000 hardware relationship for a Metro Mirror, Global Mirror, or Metro Global Mirror session. This topic also describes how to assimilate pairs in a relationship for a Metro Mirror or Global Mirror session with SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000 Unified or Storwize V7000 volumes.

You can either complete this action manually or use the data migration utility for Tivoli Storage Productivity Center for Replication. You can download the data migration utility from the Data Migration Utility for IBM Tivoli Storage Productivity Center for Replication website at http://www-01.ibm.com/support/docview.wss?uid=swg24000625. The data migration utility can handle all session types, and produces the Tivoli Storage Productivity Center for Replication CLI command script files and DSCLI script files necessary to migrate the relationships to Tivoli Storage Productivity Center for Replication.

Metro Mirror

To manually assimilate pairs in a relationship for a Metro Mirror session, complete the following steps:

1. Determine which existing Metro Mirror pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
2. Create a Metro Mirror session in Tivoli Storage Productivity Center for Replication.
3. Using the Add Copy Set Wizard, choose the same source and target volumes that you identified in step 1.
4. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
5. Issue a **Start** command from Tivoli Storage Productivity Center for Replication. At this time, the session automatically assimilates the relationships on the hardware (for example, if they are already in a Prepared state, Tivoli Storage Productivity Center for Replication will show them as Prepared).

   Note: The **Start** command does not fully resynchronize data; instead, it detects the existing relationships and adopts them.

Global Mirror

To manually assimilate pairs in a relationship for a Global Mirror session, complete the following steps:

1. Terminate the Global Mirror session on the hardware using the IBM System Storage DS command-line interface (DSCLI) or another application.
2. Remove the volumes from the Global Mirror session using the DSCLI or another application
3. Close the session on each LSS upon which it is open using the DSCLI or another application.
4. Determine which existing Global Mirror pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
5. Create a Global Mirror session in Tivoli Storage Productivity Center for Replication.
6. Using the Add Copy Set Wizard, choose the same source and target volumes that you identified in step 4.
7. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
8. Issue a **Start** command from Tivoli Storage Productivity Center for Replication. At this time, the session automatically assimilates the relationships on the hardware.

   Note: The **Start** command does not fully resynchronize data; instead, it detects the existing relationships and adopts them.

Assimilating Metro Mirror pairs into a Three Site Metro Global Mirror session

To manually assimilate a set of Metro Mirror pairs into a Three Site Metro Global Mirror session, complete the following steps:

1. Determine which existing Metro Mirror pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
2. Create a Metro Global Mirror session in Tivoli Storage Productivity Center for Replication.

3. Using the Add Copy Set Wizard, choose the same source and target volumes for the H1 and H2 roles of the Metro Global Mirror copy set, that you identified in step 1.
4. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
5. Issue a **Start** command from Tivoli Storage Productivity Center for Replication. At this time, the session automatically assimilates the Metro Mirror relationships on the hardware (for example, if they are already in a Prepared state, Tivoli Storage Productivity Center for Replication will show them as Prepared). In addition, the **Start** command will establish the Global Copy relationships.

   Note: The **Start** command does not fully resynchronize data on the Metro Mirror pairs assimilated; instead, it detects the existing relationships and adopts them. However, a full synchronization will be required for the Global Copy pairs.

Assimilating Three Site pairs into a Three Site Metro Global Mirror session

To assimilate a full Three Site set of pairs into a Three Site Metro Global Mirror session, complete the following steps:

1. Terminate the Global Mirror session on the hardware using the IBM System Storage DS command-line interface (DSCLI) or another application.
2. Remove the volumes from the Global Mirror session using the DSCLI or another application
3. Close the session on each LSS upon which it is open using the DSCLI or another application.
4. Determine which existing Metro Mirror and Global Copy pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
5. Create a Metro Global Mirror session in Tivoli Storage Productivity Center for Replication.
6. Using the Add Copy Set Wizard, choose the same source and target volumes for the H1, H2 and H3 roles of the Metro Global Mirror copy set, that you identified in step 4.
7. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
8. Issue a Start command from IBM Tivoli Storage Productivity Center for Replication. At this time, the session automatically assimilates the Metro Mirror and Global Copy relationship on the hardware.

   Note: The **Start** command does not fully resynchronize data on the Metro Mirror and Global Copy pairs assimilated; instead, it detects the existing relationships and adopts them.

Global Mirror and Metro Mirror assimilation for SAN Volume Controller, Storwize V7000, Storwize V3500, Storwize V3700, Storwize V7000 Unified, or the XIV system

To manually assimilate pairs in a relationship for a Metro Mirror or Global Mirror session with SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000,Storwize V7000 Unified or the XIV system volumes, complete the following steps:

1. Remove all pairs from the consistency group and delete the consistency group on the storage system using CLI for the storage system.
2. Determine which existing Metro Mirror and Global Mirror pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
3. Create a Metro Mirror or Global Mirror session in Tivoli Storage Productivity Center for Replication.
4. Using the Add Copy Set Wizard, choose the same source and target volumes that you identified in step 2.
5. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
6. Issue a **Start** command from Tivoli Storage Productivity Center for Replication. The session automatically assimilates the relationships on the hardware and adds them to a new consistency group on the storage system.

   Note: The **Start** command does not fully resynchronize data; instead, it detects the existing relationships and adopts them placing them in a consistency group managed by the session.

Practicing disaster recovery

You can use practice volumes to test your disaster recovery actions while maintaining disaster recovery capability. Practice volumes are available in Metro Mirror Failover and Failback sessions, Global Mirror Failover and Failback, Global Mirror Either Direction sessions, and Metro Global Mirror with Practice sessions.

Important: You can test your disaster recovery actions without using practice volumes. However, without practice volumes, you cannot maintain disaster recovery capability while continuing to copy the data.

Practice volumes

You can use a *practice volume* to practice what you would do in the event of a disaster, without interrupting current data replication. Practice volumes are available in Metro Mirror, Global Mirror, and Metro Global Mirror sessions.

To use the practice volumes, the session must be in the prepared state. Issuing the Flash command against the session while in the Prepared state creates a usable practice copy of the data on the target site.

Note: You can test disaster-recovery actions without using practice volumes; however, without practice volumes, you cannot continue to copy data changes between volumes while testing disaster-recovery actions.

Practicing disaster recovery for a Metro Mirror Failover/Failback with Practice session

A Metro Mirror Failover and Failback session with Practice combines Metro Mirror and FlashCopy to provide a point-in-time copy of the data on the remote site. You can use this to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

This function is available on the following storage systems:

- System Storage SAN Volume Controller
- Storwize V7000

- TotalStorage Enterprise Storage Server Model 800
- System Storage DS8000
- System Storage DS6000

Perform these steps to practice disaster recover actions for a Metro Mirror Failover/Failback with Practice session:

1. Start a Metro Mirror with Practice session.
2. When the Metro Mirror session reaches the Prepared state, issue a Flash command to make a point-in-time copy of the data on H2. This creates a consistent point-in-time copy of your data on the H2 volume and then restarts the session so the copying from H1 to I2 continues. This temporarily stops copying of the data from site 1 to site 2, and creates a consistent point-in-time copy. The data replication from H1 to I2 is then restarted.

   Note: For ESS, DS6000, and DS8000 storage systems, the **Flash** command uses the freeze and thaw processing to create a data consistent point for the FlashCopy. If there is another Metro Mirror session overlapping on one or more of the same LSS pairs, that session will be suspended. It is also possible that the suspension of the other session might cause the Metro Mirror session to remain suspended after the flash command is issued instead of returning to Prepared state. Avoid using the same LSS pairs for multiple Metro Mirror sessions if possible.
3. Practice the same actions you would take in an actual disaster, using H2 as your new host volume, while running the real application on H1.

Practicing disaster recovery for a Global Mirror either Direction with two-site Practice session

A Global Mirror (either direction) with two-site Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your first site. You can use this to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

Note: This function is available only on ESS, DS6000, and DS8000 storage systems.

Perform these steps to practice disaster recover actions for a Global Mirror either Direction with two-site Practice session:

1. Start a Global Mirror with Practice session.
2. When the session reaches the Prepared state, issue a Flash command to restore consistent data on I2 and make a point-in-time copy of the data on H2. This creates a consistent point-in-time copy of your data on the H2 volume and then restarts the session so the copying from H1 to I2 continues. This temporarily stops copying of the data from site 1 to site 2, and creates a consistent point-in-time copy. The data replication from H1 to I2 is then restarted.

   Note: FlashCopy must be always a full copy due to limitations of the hardware.
3. Practice the same actions you would take in an actual disaster, using H2 as your new host volume, while running the real application on H1.

   Note: With two directions, you can reverse the direction of your data flow.

Practicing disaster recovery for a Global Mirror Failover/Failback with Practice session

A Global Mirror Failover and Failback with Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your first site. You can use this to practice what you might do if a disaster occurred.

You can do this practice without losing your disaster recovery capability. The number of volumes used for the device varies, but the steps to conduct a Global Mirror Failover and Failback with Practice are the same for both devices.

This function is available on the following storage systems:

- System Storage SAN Volume Controller
- Storwize V7000
- TotalStorage Enterprise Storage Server Model 800
- System Storage DS8000
- System Storage DS6000

Perform these steps to practice disaster recover actions for a Global Mirror Failover/Failback with Practice session:

1. Start a Global Mirror with Practice session.
2. When the session reaches the Prepared state, issue a Flash command to restore consistent data on I2 and make a point-in-time copy of the data on H2. This creates a consistent point-in-time copy of your data on the H2 volume and then restarts the session so the copying from H1 to I2 continues. This temporarily stops copying of the data from site 1 to site 2, and creates a consistent point-in-time copy. The data replication from H1 to I2 is then restarted.

   Note: FlashCopy must be always a full copy due to limitations of the hardware.
3. Practice the same actions you would take in an actual disaster, using H2 as your new host volume, while running the real application on H1.

Practicing disaster recovery for a Metro Global Mirror Failover/Failback with Practice session

A Metro Global Mirror Failover/Failback with Practice session combines Metro Mirror, Global Mirror and FlashCopy across three sites to provide a point-in-time copy of the data on the third site. You can use this to practice what you might do if a disaster occurred without losing your disaster recovery capability.

Note: This function is available on ESS, DS6000, and DS8000 storage systems.

The intermediate volume is on the third site (I3). This maintains disaster recovery capability while a copy is kept on the H3 volume for practice purposes.

Perform these steps to practice disaster recover actions for a Metro Global Mirror Failover/Failback with Practice session:

1. Start a Metro Global Mirror with Practice session.
2. When the session reaches the Prepared state, issue a Flash command to take a point-in-time copy of the data that is on I3, on H3. This creates a consistent point-in-time copy of your data on the H3 volume, and then automatically restarts the session so that copying from H1 to H2 to I3 continues. The Flash

command temporarily stops copying the data from site 2 to site 3, in order to create a consistent point-in-time copy on I3, while maintaining disaster recovery capabilities on site 2 using the Metro Mirror portion of the session. Then, data replication from H2 to I3 is restarted.

Note: FlashCopy must be always a full copy due to hardware limitations.

3. Practice the same actions you would take in an actual disaster, using H3 as your practice host volume, while you run the real application on H1. This enables you to use the same scripts and commands to run on H3 that you would use in an actual disaster.

Monitoring health and status

Viewing the health summary

Use the Health Overview panel to view overall health and status of sessions, storage systems, host systems, and management servers.

The Health Overview panel is the first panel that you see after you log on. You can display this panel by selecting **Health Overview** in the Tivoli Storage Productivity Center for Replication navigation tree. This panel provides the following information:

Overall session status
Indicates session status, which can be normal, warning, or severe. The status can also be inactive, if all sessions are defined or if no sessions exit.

Overall storage system status
Indicates the connection status of storage systems.

Overall host system status
Indicates the connection status of host systems.

Management server status
Indicates the status of the standby server if you are logged on to the local server. If you are logged on to the standby server, this status indicates the status of the local server. Management server status is not available if you are using Tivoli Storage Productivity Center for Replication Basic Edition for System z.

Health information is always shown as a mini-panel under the navigation tree.

Viewing SNMP alerts

You can view Tivoli Storage Productivity Center for Replication SNMP trap descriptions from the Alert Log node in the Tivoli Storage Productivity Center stand-alone GUI.

In the stand-alone GUI navigation tree, expand **Alerting** > **Alert Log** and click **Replication**.

Viewing sessions

This section describes how to view sessions details and session properties.

Session status icons

The IBM Tivoli Storage Productivity Center for Replication GUI uses icons to represent the status of each session.

The following table describes each session status icon.

Table 110. Session status icons

Icon	Meaning	Description
	Inactive	The session is in a defined state, with no activity on the hardware.
	Normal	A consistent copy of the data either exists or is being maintained.
	Warning	For Metro Mirror, Global Mirror, and Metro Global Mirror, the session might have volumes that are being synchronized or are about to be synchronized, with no suspended volumes. For FlashCopy, the warning status is valid only after the start command is issued and before the flash. This warning status means that the session is either preparing or is ready for a flash command but targets do not yet have a consistent copy. If a HyperSwap session is degraded, which means it is enabled on one or more sysplex members and disabled on at least one sysplex member, then the session is in a warning state.
	Severe	One or more errors must be dealt with immediately. Possible causes include the following: • One or more volumes are suspended • A session is suspended • A volume is not copying correctly

Session images

The IBM Tivoli Storage Productivity Center for Replication GUI provides a visual aid to help you create and manage your sessions. The visual aid shows the number of volume roles in the session and how the roles are distributed between the sites. It also shows the copy method and direction.

Volume role symbols

The volume role symbols represent the replication status on the volumes.

Table 111. Volume role symbols

Symbol	Description	Meaning
	Active host volumes	This symbol represents volumes that contain the source of updated tracks to which the application is actively issuing read and write input/output (I/O).
	Recoverable volumes	This symbol represents volumes that contain a consistent copy of the data.
	Inconsistent volumes	This symbol represents the volumes that do not contain a consistent copy of the data.

Table 111. Volume role symbols (continued)

Symbol	Description	Meaning
	Selected volumes	This symbol represents the volumes that are selected for an operation (for example, changing location or displaying role pair information).

Data copying symbols

The data copying symbols are arrows indicate the type of copy that occurs between the volume roles. The direction of the arrow indicates the direction of the copy.

Table 112. Data copying symbols

Symbol	Description	Meaning
	FlashCopy copying (the lightning bolt indicates the direction of the FlashCopy)	This symbol represents a FlashCopy relationship in which data is being copied from the host to the target.
	FlashCopy copying with errors (the lightning bolt indicates the direction of the FlashCopy)	This symbol represents a FlashCopy relationship in which data is being copied from the host to the target, but there are errors on one or more pairs.
	FlashCopy inactive (the lightning bolt indicates the direction of the FlashCopy)	This symbol represents an inactive FlashCopy relationship.
	FlashCopy inactive with errors (the lightning bolt indicates the direction of the FlashCopy)	This symbol represents an inactive FlashCopy relationship in which there are errors on one or more pairs.
	Metro Mirror copying	This symbol represents a copying Metro Mirror relationship.
	Metro Mirror copying with errors	This symbol represents a Metro Mirror relationship that is copying, but with errors on one or more pairs.
	Metro Mirror inactive	This symbol represents an inactive Metro Mirror relationship.

Table 112. Data copying symbols (continued)

Symbol	Description	Meaning
	Metro Mirror inactive with errors	This symbol represents an inactive Metro Mirror relationship with errors on one or more pairs.
	Global Copy copying	This symbol represents a copying Global Copy relationship.
	Global Copy copying with errors	This symbol represents a copying Global Copy relationship with errors on one or more pairs.
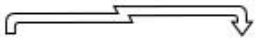	Global Copy inactive	This symbol represents an inactive Global Copy relationship.
	Global Copy inactive with errors	This symbol represents an inactive Global Copy relationship with errors on one or more pairs.
	HyperSwap or Open HyperSwap	This symbol HyperSwap or Open HyperSwap for a session. If a failure occurs when input/output (I/O) is being written to the primary storage system, HyperSwap or Open HyperSwap automatically swap the I/O to the secondary site with no user interaction and little or no application impact.

Session states

You can view the health and status of a session in the Tivoli Storage Productivity Center for Replication GUI.

Attention: Use only the Tivoli Storage Productivity Center for Replication graphical user interface (GUI) or command-line interface CLI to manage session relationships, such as volume pairs and copy sets. Do not modify session relationships through other interfaces such as the System Storage DS CLI. If you modify relationships through other interfaces, a loss of consistency can occur across the relationships that are managed by the session. The exceptions to this requirement are failover operations that are managed by external applications for certain session and storage system types as described in "Failover and failback operations" on page 671.

The `Refresh States` command is used to refresh the states of the session. Issue this command to query the states of the copy sets on the hardware. You do not need to

run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable Tivoli Storage Productivity Center for Replication to update itself. Because this command triggers multiple queries on the hardware (having an adverse impact on hardware performance if you run it too often), you can only execute this command every few minutes in each session.

The following table describes each session state.

Table 113. Session states

State	Session type	Description
Defined	All	The session exists but is inactive.
Flashing	All	In a Metro Mirror or Global Mirror session, data copying is temporarily suspended while a consistent practice copy of data is being prepared on site 2. In a Metro Global Mirror session, data copying is temporarily suspended while a consistent practice copy of data is being prepared on site 3.
Prepared	All	The source to target data transfer is active. In a Metro Mirror, Global Mirror, or Metro Global Mirror session, the data written to the source is transferred to the target, and all volumes are consistent and recoverable. In a FlashCopy session, the volumes are not yet consistent, but the flash is ready to begin. **Important:** For sessions on the following storage systems, do not alter the relationships on the hardware that you established with Tivoli Storage Productivity Center for Replication: • SAN Volume Controller • Storwize V3500 (use only FlashCopy sessions for this storage system) • Storwize V3700 • Storwize V7000 • Storwize V7000 Unified For example, if a Metro Mirror session with one copy set is in the Prepared state, and you stop the role pair, the session is still displayed in the Prepared state.

Table 113. Session states (continued)

State	Session type	Description
Preparing	All	The volumes are initializing, synchronizing, or resynchronizing. In a Metro Mirror, Global Mirror, or Metro Global Mirror session, synchronization occurs after the first **Start** command is issue on a session. Resynchronization occurs when a volume was prepared and then suspended. The hardware records the changed tracks so that on the next startup, only the changed tracks are copied. In a FlashCopy session, the volumes are initializing. The preparing state for FlashCopy sessions applies only to the SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, or Storwize V7000 Unified.
Recovering	All	The session is in the process of recovering.
Suspended	All	Data copying has temporarily stopped. **Important:** The suspended state applies only to Global Mirror, Metro Mirror, and Metro Global Mirror sessions.
SuspendedH1H2 MGM	Metro Global Mirror	Data copying between site 1 and site 2 is suspended.
SuspendedH1H3 MGM	Metro Global Mirror	Data copying between site 1 and site 3 is suspended.
Suspending	All	The session is transitioning into a Suspended state. **Important:** The Suspending state applies only to Global Mirror and Metro Global Mirror sessions and does not apply to SAN Volume Controller, Storwize V3700, Storwize V7000, or Storwize V7000 Unified.
Target available	All	Target volumes are available for application updates.
Terminating	FlashCopy	The session is being terminated because you issued a Terminate action under the following conditions: • You permitted the target to be Metro Mirror or Global Copy. • You set the Require or Attempt to Preserve Mirror option. The session displays as Terminating until the FlashCopy background copy is complete and no longer exists on the hardware.

Role pair status and progress

In a session, a *role pair* is the association of two volume roles that take part in a copy relationship. Tivoli Storage Productivity Center for Replication provides detailed role pair status and progress messages for sessions.

The role pair status and progress messages are updated to provide a message that indicates what the session is doing at the time. By hovering over a progress bar, you can see specific information about the action running on the session. Some status messages might include an estimated time-to-completion for the action in hours and minutes.

Role pair status is not provided for the XIV system Snapshot sessions because role pairs are not used for these sessions.

The status messages are displayed in the Session Details and Role Pair Details page.

Table 114. Detailed status messages for Participating and Non-Participating role pairs

Supported session type	Status message
FlashCopy Metro Mirror Global Copy Global Mirror	Starting *role_pair_name* relationships on the hardware
Metro Mirror Global Mirror	Waiting for all pairs in the role pair *role_pair_name* to reach state of *state*
FlashCopy Metro Mirror Global Copy Global Mirror	Terminating all pairs in role pair *role_pair_name*
FlashCopy Metro Mirror Global Mirror	Recovering all pairs in role pair *role_pair_name*
Metro Mirror Global Copy Global Mirror	Suspending all pairs in role pair *role_pair_name*
FlashCopy	Background copy is running for role pair *role_pair_name*
Global Copy	Waiting for all pairs in role pair *role_pair_name* to become consistent
Global Copy	Waiting for all pairs in role pairs *role_pair_name* to complete the initial copy
FlashCopy	Waiting for all pairs in role pairs *role_pair_name* to complete FRR
Global Mirror	Waiting for all pairs in role pairs *role_pair_name* to join the Global Mirror session

Viewing session details

You can view detailed information about a session, including role pairs, error count, whether the session is recoverable, copying progress, session type, and the timestamp.

Perform these steps to view session details:

1. In the navigation tree, select **Sessions**.
2. Select the session that you want to view.
3. Select **View Details** from the **Actions** drop-down menu, and click **Go**.

Viewing additional details for Global Mirror sessions

Additional detail information is available for TotalStorage Enterprise Storage Server Model 800, System Storage DS6000, and System Storage DS8000 Global Mirror sessions. This information includes details about the Global Mirror master, consistency groups that have been formed, and data exposure.

Perform these steps to view additional details for Global Mirror sessions:

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **Sessions**.
2. Select the Global Mirror session that you want to view.
3. Select **View Details** from the **Actions** drop-down menu, and click **Go**.
4. Click the **Global Mirror Info** tab. The following information is displayed on the tab:

   Global Mirror Master
   Shows the name of the storage system acting as the Global Mirror master.

   Last Master Consistency Group Time
   Shows the time that the last consistency group was formed

   Master Time During Last Query
   Shows the time on the master storage device when the query was performed,

   Data Exposure
   Shows the average exposure to potential data loss in seconds over the query interval.

   Session ID
   Shows the Global Mirror session ID.

   Master State
   Shows the state of the master session on the hardware.

   Unsuccessful CGs since last successful CG
   Shows the number of consistency groups that have failed to form since the last successful consistency group was formed.

   CG Interval Time
   Shows the interval time between attempts to form a consistency group.

   Max Coordination Interval
   Shows the extended distance consistency maximum coordination interval.

   Max CG Drain Time
   Shows the maximum time the consistent set of data is allowed to drain at the remote site before failing consistency group formation.

   Unsuccessful CGS/Previous Query
   Shows the number of consistency groups and percentage of consistency groups that were unsuccessful since the previous query.

   Unsuccessful CGS/Total
   Shows the total number of unsuccessful consistency groups and percentage of consistency groups that have failed.

Successful CGS/Previous Query
Shows the number of consistency groups and percentage of consistency groups that were successful since the previous query.

Successful CGS/Total
Shows the total number of successful consistency groups and percentage of consistency groups that have been successful.

Consistency Group Failure Messages
Shows the failure messages that have occurred on the Global Mirror session that prevented the formation of a consistency group.

Data Exposure chart
Shows the data exposure values in seconds for the last 15 minutes or 24 hours.

Highlight Data Exposure
Use the following fields to define a value in seconds for which you want data exposure to tracked in the **Data Exposure** chart.

Show Data Exposure over
Data exposure that is over the value that is entered in this field is shown in the **Data Exposure** chart.

Show Data Exposure under
Data exposure that is under the value that is entered in this field is shown in the **Data Exposure** chart.

Viewing session properties

You can view and modify the description and options for a session. The options that are presented depend on the storage system type and the session type that you select.

To view or modify the properties for a session, complete the following steps:

1. In the navigation tree, click **Sessions**. The Sessions page is displayed.
2. Select the session with the properties that you want to view.
3. Click **View/Modify Properties** on the **Select Action** menu and click **Go**. The properties that are displayed depend on the session type as described in the following sections.

TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 sessions

The properties for IBM TotalStorage Enterprise Storage Server Model 800, IBM System Storage DS8000, and IBM System Storage DS6000 sessions depend on the session type.

The following sections describe the properties for each session type.

FlashCopy session properties:

Use the View/Modify Properties page to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 FlashCopy sessions.

The following options are on this page.

Description
Enter a description for the session.

ESS/DS FlashCopy Options

Select the options for the session.

Incremental

Select this option to apply incremental changes to the target volume. After the initial FlashCopy operation, only data that changed on the source volume since the last FlashCopy operation was performed is copied to the target volume.

If you select this option, a persistent FlashCopy relationship is created regardless of whether you select the **Persistent** check box.

Persistent

Select this option to keep the FlashCopy relationship established on the hardware after all source tracks are copied to the target volume. If you do not select this option, the local replication relationship ends after the target volume contains a complete point-in-time image of the source volume.

No Copy

Select this option if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the target volume until the blocks or tracks of the source volume are modified.

This option is required for space-efficient volumes.

Allow FlashCopy target to be Metro Mirror source

Select this option to enable the FlashCopy operation if the target volume of the FlashCopy relationship is also the source volume of a Metro Mirror relationship. If this option is not selected, the FlashCopy operation fails.

Requirement: This option requires that the IBM Remote Pair FlashCopy option is available for your IBM System Storage DS8000 storage system. To determine whether you can use the IBM Remote Pair FlashCopy option with your System Storage DS8000 storage system, refer to the System Storage DS8000 documentation for microcode level that you are using.

Select one of the following options to specify whether you want to maintain consistency, if possible:

Do not attempt to preserve Metro Mirror consistency

Click this option if you want the FlashCopy operation to complete without preserving consistency of the Metro Mirror relationship on the remote site. The FlashCopy operation does not occur on the remote site.

Attempt to preserve Metro Mirror consistency, but allow FlashCopy even if Metro Mirror target consistency cannot be preserved

Click this option to preserve the consistency of the Metro Mirror relationship at the target of the FlashCopy relationship when both the source and target of the FlashCopy relationship are the source of a Metro Mirror relationship. If the consistency cannot be preserved, a full copy of the Metro Mirror relationship at the target of the FlashCopy relationship is performed. To preserve consistency, parallel FlashCopy operations are performed on both sites if possible.

Attempt to preserve Metro Mirror consistency, but fail FlashCopy if Metro Mirror target consistency cannot be preserved
Click this option to prevent a full copy from being performed over the Metro Mirror link. Instead, if possible, parallel FlashCopy operations are performed on both sites. If the consistency cannot be preserved, the flash for the FlashCopy relationships fails, and the data of the Metro Mirror relationship at the target of the FlashCopy relationship is not changed.

Metro Mirror Single Direction session properties:

Use the View/Modify Properties page to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Mirror Single Direction sessions.

The following options are on this page.

Description
Enter a description for the session.

ESS/DS Metro Mirror Options
Select the options for the session.

Reset Secondary Reserves
Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

Attention: This option causes the session to overwrite all data that is on the target volume.

Fail MM/GC if target is online (CKD only)
Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

Metro Mirror Suspend Policy
Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume

Hold I/O after Suspend
Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication **Release I/O** command for the session. If you

do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

Release I/O after Suspend
Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the affect on the applications that are writing to the source volume. This option is enabled by default.

Attention: This option can cause the data on the source volume to be different from the data on target volume.

System or sysplex
Select the z/OS system or sysplex that is attached to the volumes that are contained in the session. This system or sysplex is used to enable the **Enabled Hardened Freeze** option.

Important: Tivoli Storage Productivity Center for Replication must be installed on or connected to one or more host systems to select a system or sysplex and to use the **Enabled Hardened Freeze** option.

Enable Hardened Freeze
Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

Requirement: This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

Metro Mirror Failover/Failback session properties:

Use the View/Modify Properties page to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Mirror Failover/Failback sessions.

The following options are on this page.

Description
Enter a description for the session.

ESS/DS Metro Mirror Options

Select the options for the session.

Reset Secondary Reserves

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

Attention: This option causes the session to overwrite all data that is on the target volume.

Fail MM/GC if target is online (CKD only)

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

Metro Mirror Suspend Policy

Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume

Hold I/O after Suspend

Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication **Release I/O** command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

If the **Manage H1-H2 with Hyperswap** option is selected, you must determine what caused the relationships to be suspended and then take the following actions:

If the source volumes are accessible.

Issue the IBM z/OS **SETHS RESUMEIO** command and the **Release I/O** command. If you do not issue the **SETHS RESUMEIO** command, the z/OS extended long busy timer value must expire before you can issue the **Release I/O** command.

There are situations in which an application must access some of the HyperSwap managed volumes, such as the paging volumes, and you might not be able to issue the **SETHS RESUMEIO** command or the command might not finish. In these situations, restart the application to enable write operations to the H1 site to resume. Restart the application after you issue the **Release I/O** command in Tivoli

Storage Productivity Center for Replication or after the extended long busy timer value expires.

If one or more of the source volumes are no longer accessible, for example, in the case of a disaster.
Issue the Tivoli Storage Productivity Center for Replication **Recover** command and complete an initial program load (IPL) on the H1 site by using the storage data from the H2 site. This data might not be the latest data, but the data should be consistent.

Release I/O after Suspend
Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the affect on the applications that are writing to the source volume. This option is enabled by default.

Attention: This option can cause the data on the source volume to be different from the data on target volume.

System or sysplex
Select the z/OS system or sysplex that is attached to the volumes that are contained in the session. This system or sysplex is used to enable the **Enabled Hardened Freeze** and **Manage H1-H2 with HyperSwap** options.

Important: Tivoli Storage Productivity Center for Replication must be installed on or connected to one or more host systems to select a system or sysplex and to use the **Enabled Hardened Freeze** and **Manage H1-H2 with HyperSwap** options.

Enable Hardened Freeze
Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

Requirement: This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

If you select the **Manage H1-H2 with HyperSwap** option for the session, this option is ignored. HyperSwap includes IOS for

managing freeze operations. The **Enable Hardened Freeze** option ensures data integrity if Tivoli Storage Productivity Center for Replication freezes and HyperSwap is not enabled for a session.

Manage H1-H2 with HyperSwap
Select this option to trigger a HyperSwap operation, which redirects application I/O to the target volumes when there is a failure on the host accessible volumes. Tivoli Storage Productivity Center for Replication uses HyperSwap to manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session.

Disable HyperSwap
Select this option to prevent a HyperSwap operation from occurring.

On Configuration Error:

Partition the system(s) out of the sysplex
Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

Disable HyperSwap
Select this option to prevent a HyperSwap operation from occurring.

On Planned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)
Select this option to partition out the failing system and continue the swap processing on any remaining systems.

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

On Unplanned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)
Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and the HyperSwap operation does not complete.

Requirement: If you select this option, you must restart the system.

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

Manage H1-H2 with Open HyperSwap
Select this option to trigger an Open HyperSwap operation for

volumes that are attached to an IBM AIX host. This option redirects application I/O to the target volumes when there is a failure on the host accessible volumes.

Tivoli Storage Productivity Center for Replication uses Open HyperSwap to manage the H1-H2 sequence of a Metro Mirror session.

Only volumes that are attached to host systems that are defined in the Tivoli Storage Productivity Center for Replication Host Systems panel are eligible for Open HyperSwap.

Disable Open HyperSwap
Select this option to prevent an Open HyperSwap operation from occurring while keeping the configuration on the host system and all source and target volumes coupled.

Metro Mirror Failover/Failback with Practice session properties:

Use the View/Modify Properties page to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Mirror Failover/Failback with Practice sessions.

The following options are on this page.

Description
Enter a description for the session.

ESS/DS Metro Mirror Options
Select the options for the session.

Reset Secondary Reserves
Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

Attention: This option causes the session to overwrite all data that is on the target volume.

Fail MM/GC if target is online (CKD only)
Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

Metro Mirror Suspend Policy
Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume

Hold I/O after Suspend
Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source

volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication **Release I/O** command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

Release I/O after Suspend
Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the affect on the applications that are writing to the source volume. This option is enabled by default.

Attention: This option can cause the data on the source volume to be different from the data on target volume.

System or sysplex
Select the z/OS system or sysplex that is attached to the volumes that are contained in the session. This system or sysplex is used to enable the **Enabled Hardened Freeze** option.

Important: Tivoli Storage Productivity Center for Replication must be installed on or connected to one or more host systems to select a system or sysplex and to use the **Enabled Hardened Freeze** option.

Enable Hardened Freeze
Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

Requirement: This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

DS FlashCopy Options for Role pair H2-I2
The following option is available only for IBM System Storage DS8000 version 4.2 or later.

Persistent
Select this option to keep FlashCopy pairs persistent on the hardware.

Global Mirror Single Direction and Failover/Failback session properties:

Use the View/Modify Properties page to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Global Mirror Single Direction and Global Mirror Failover/Failback sessions.

The following options are on this page.

Description
Enter a description for the session.

ESS/DS Global Mirror Options
Select the options for the session.

Consistency group interval time (seconds)
Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

Recovery Point Objective Alerts
Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

Warning level threshold (seconds)
Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

Severe level threshold (seconds)
Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

Reset Secondary Reserves
Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

Attention: This option causes the session to overwrite all data that is on the target volume.

Fail MM/GC if target is online (CKD only)
Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

Global Mirror Failover/Failback with Practice session properties:

Use the View/Modify Properties page to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Global Mirror Failover/Failback with Practice sessions.

The following options are on this page.

Description
Enter a description for the session.

ESS/DS Global Mirror Options
Select the options for the session.

Consistency group interval time (seconds)
Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

Recovery Point Objective Alerts
Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

Warning level threshold (seconds)
Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

Severe level threshold (seconds)
Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

Reset Secondary Reserves
Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

Attention: This option causes the session to overwrite all data that is on the target volume.

Fail MM/GC if target is online (CKD only)
Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

DS FlashCopy Options for Role pair H2-I2
The following option is available only for IBM System Storage DS8000 version 4.2 or later.

Persistent
Select this option to keep FlashCopy pairs persistent on the hardware.

No Copy
Select this option if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the H2 volume until the blocks or tracks of the I2 volume are modified.

DS FlashCopy Options for Role pair I2-J2
The following option is available only for System Storage DS8000 version 4.2 or later.

Reflash After Recover
Select this option if you want to create a FlashCopy replication between the I2 and J2 volumes after the recovery of a Global Mirror Failover/Failback with Practice session. If you do not select this option, a FlashCopy replication is created only between the I2 and H2 volumes.

Global Mirror Either Direction with Two-Site Practice session properties:

Use the View/Modify Properties page to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Global Mirror Either Direction with Two-Site Practice sessions.

The following options are on this page.

Description
Enter a description for the session.

ESS/DS Global Mirror Options
Select the options for the session.

Consistency group interval time (seconds)
Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the

session to attempt to create consistency groups more frequently, which can increase network traffic.

Recovery Point Objective Alerts
Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

Warning level threshold (seconds)
Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

Severe level threshold (seconds)
Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

Reset Secondary Reserves
Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

Attention: This option causes the session to overwrite all data that is on the target volume.

Fail MM/GC if target is online (CKD only)
Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

Metro Global Mirror session properties:

Use the View/Modify Properties page to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Global Mirror sessions.

The following options are on this page.

Description
Enter a description for the session.

ESS/DS Metro Global Mirror Options

Select the options for the session.

Consistency group interval time (seconds)

Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

Recovery Point Objective Alerts

Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

Warning level threshold (seconds)

Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

Severe level threshold (seconds)

Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

Reset Secondary Reserves

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

Attention: This option causes the session to overwrite all data that is on the target volume.

Fail MM/GC if target is online (CKD only)

Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

Metro Mirror Suspend Policy

Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is

suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume

Hold I/O after Suspend
Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for Replication **Release I/O** command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

If the **Manage H1-H2 with Hyperswap** option is selected, you must determine what caused the relationships to be suspended and then take the following actions:

If the source volumes are accessible.
Issue the IBM z/OS **SETHS RESUMEIO** command and the **Release I/O** command. If you do not issue the **SETHS RESUMEIO** command, the z/OS extended long busy timer value must expire before you can issue the **Release I/O** command.

There are situations in which an application must access some of the HyperSwap managed volumes, such as the paging volumes, and you might not be able to issue the **SETHS RESUMEIO** command or the command might not finish. In these situations, restart the application to enable write operations to the H1 site to resume. Restart the application after you issue the **Release I/O** command in Tivoli Storage Productivity Center for Replication or after the extended long busy timer value expires.

If one or more of the source volumes are no longer accessible, for example, in the case of a disaster.
Issue the Tivoli Storage Productivity Center for Replication **Recover** command and complete an initial program load (IPL) on the H1 site by using the storage data from the H2 site. This data might not be the latest data, but the data should be consistent.

Release I/O after Suspend
Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the affect on the applications that are writing to the source volume. This option is enabled by default.

Attention: This option can cause the data on the source volume to be different from the data on target volume.

System or sysplex
Select the z/OS system or sysplex that is attached to the volumes that are contained in the session. This system or sysplex is used to enable the **Enabled Hardened Freeze** and **Manage H1-H2 with HyperSwap** options.

Important: Tivoli Storage Productivity Center for Replication must be installed on or connected to one or more host systems to select a system or sysplex and to use the **Enabled Hardened Freeze** and **Manage H1-H2 with HyperSwap** options.

Enable Hardened Freeze
Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

Requirement: This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

If you select the **Manage H1-H2 with HyperSwap** option for the session, this option is ignored. HyperSwap includes IOS for managing freeze operations. The **Enable Hardened Freeze** option ensures data integrity if Tivoli Storage Productivity Center for Replication freezes and HyperSwap is not enabled for a session.

Manage H1-H2 with HyperSwap
Select this option to trigger a HyperSwap operation, which redirects application I/O to the target volumes when there is a failure on the host accessible volumes. Tivoli Storage Productivity Center for Replication uses HyperSwap to manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session.

Disable HyperSwap
Select this option to prevent a HyperSwap operation from occurring.

On Configuration Error:

Partition the system(s) out of the sysplex
Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

Disable HyperSwap
Select this option to prevent a HyperSwap operation from occurring.

On Planned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)
Select this option to partition out the failing system and continue the swap processing on any remaining systems.

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

On Unplanned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)
Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and the HyperSwap operation does not complete.

Requirement: If you select this option, you must restart the system.

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

Metro Global Mirror with Practice session properties:

Use the View/Modify Properties page to view or modify properties for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 Metro Global Mirror with Practice sessions.

The following options are on this page.

Description
Enter a description for the session.

ESS/DS Metro Global Mirror Options
Select the options for the session.

Consistency group interval time (seconds)
Type how often, in seconds, the Global Mirror session attempts to form a consistency group. A lower value possibly reduces the data exposure of the session. However, a lower value also causes the session to attempt to create consistency groups more frequently, which can increase network traffic.

Recovery Point Objective Alerts
Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is

exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 65535 seconds. The default is 0 seconds, which specifies that no alerts are generated.

Warning level threshold (seconds)
Type the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Consistency group interval time (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

Severe level threshold (seconds)
Type the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

Reset Secondary Reserves
Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets when a Start command is issued for the session.

Attention: This option causes the session to overwrite all data that is on the target volume.

Fail MM/GC if target is online (CKD only)
Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

Metro Mirror Suspend Policy
Select one of the following options to specify the policy for holding or releasing input/output (I/O) after a Metro Mirror relationship is suspended. When a relationship is in a Suspended state, write operations from the source volume are no longer mirrored to the target volume

Hold I/O after Suspend
Click this option to prevent the source volume from receiving new data after the relationship is suspended. Use this option if you want to determine the scope of the suspension before new data is written to the source volume. This option helps to ensure that the data on the source and target volumes remains the same.

To enable write operations to the source volume to continue, issue the Tivoli Storage Productivity Center for

Replication **Release I/O** command for the session. If you do not issue this command, write operations are enabled when the hardware timeout value on the storage system expires.

If the **Manage H1-H2 with Hyperswap** option is selected, you must determine what caused the relationships to be suspended and then take the following actions:

If the source volumes are accessible.
Issue the IBM z/OS **SETHS RESUMEIO** command and the **Release I/O** command. If you do not issue the **SETHS RESUMEIO** command, the z/OS extended long busy timer value must expire before you can issue the **Release I/O** command.

There are situations in which an application must access some of the HyperSwap managed volumes, such as the paging volumes, and you might not be able to issue the **SETHS RESUMEIO** command or the command might not finish. In these situations, restart the application to enable write operations to the H1 site to resume. Restart the application after you issue the **Release I/O** command in Tivoli Storage Productivity Center for Replication or after the extended long busy timer value expires.

If one or more of the source volumes are no longer accessible, for example, in the case of a disaster.
Issue the Tivoli Storage Productivity Center for Replication **Recover** command and complete an initial program load (IPL) on the H1 site by using the storage data from the H2 site. This data might not be the latest data, but the data should be consistent.

Release I/O after Suspend
Click this option to enable the source volume to automatically receive data after the relationship is suspended. Use this option if you want to limit the affect on the applications that are writing to the source volume. This option is enabled by default.

Attention: This option can cause the data on the source volume to be different from the data on target volume.

System or sysplex
Select the z/OS system or sysplex that is attached to the volumes that are contained in the session. This system or sysplex is used to enable the **Enabled Hardened Freeze** and **Manage H1-H2 with HyperSwap** options.

Important: Tivoli Storage Productivity Center for Replication must be installed on or connected to one or more host systems to select a system or sysplex and to use the **Enabled Hardened Freeze** and **Manage H1-H2 with HyperSwap** options.

Enable Hardened Freeze
Select this option to enable the z/OS Input/Output Supervisor (IOS) to manage freeze operations.

Requirement: This option requires the z/OS address spaces Basic HyperSwap Management and Basic HyperSwap API. For instructions about how to start these address spaces, see the information about preparing to use Basic HyperSwap from z/OS in the *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.

If you select this option, IOS can freeze volumes regardless of whether the Tivoli Storage Productivity Center for Replication server is started or stopped.

In addition, this option enables you to include z/OS system volumes such as paging, database, and IBM WebSphere Application Server hierarchical file system (HFS) as Metro Mirror volumes in the session. When you select this option, IOS manages the freeze operations for all Metro Mirror volumes in the session, which prevents Tivoli Storage Productivity Center for Replication from freezing the volumes and possibly freezing itself. This option does not enable IOS to manage freeze operations for Global Mirror volumes.

If you select the **Manage H1-H2 with HyperSwap** option for the session, this option is ignored. HyperSwap includes IOS for managing freeze operations. The **Enable Hardened Freeze** option ensures data integrity if Tivoli Storage Productivity Center for Replication freezes and HyperSwap is not enabled for a session.

Manage H1-H2 with HyperSwap
Select this option to trigger a HyperSwap operation, which redirects application I/O to the target volumes when there is a failure on the host accessible volumes. Tivoli Storage Productivity Center for Replication uses HyperSwap to manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session.

Disable HyperSwap
Select this option to prevent a HyperSwap operation from occurring.

On Configuration Error:

Partition the system(s) out of the sysplex
Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

Disable HyperSwap
Select this option to prevent a HyperSwap operation from occurring.

On Planned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)
Select this option to partition out the failing system and continue the swap processing on any remaining systems.

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

On Unplanned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)
Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and the HyperSwap operation does not complete.

Requirement: If you select this option, you must restart the system.

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

ESS / DS FlashCopy Options for Role pair H3-I3
The following option is available only for IBM System Storage DS8000 version 4.2 or later.

No Copy
Select this option if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the H3 volume until the blocks or tracks of the I3 volume are modified.

SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified sessions

The properties for IBM System Storage SAN Volume Controller, IBM Storwize V3500, IBM Storwize V3700, IBM Storwize V7000, and IBM Storwize V7000 Unified sessions depend on the session type.

The following sections describe the properties for each session type with the exception of session types that include only the **Description** property.

FlashCopy session properties:

Use the View/Modify Properties page to view or modify properties for SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified FlashCopy sessions.

The following options are on this page.

Description
Enter a description for the session.

SAN Volume Controller / Storwize Family / Storwize V7000 Unified FlashCopy Options
Select the options for the session.

Incremental
Select this option to apply incremental changes to the target volume. After the initial FlashCopy operation, only data that changed on the source volume since the last FlashCopy operation was performed is copied to the target volume.

Background Copy Rate
Type the copy rate that the storage system uses to perform the background copy of the FlashCopy role pair. You can specify a percentage in the range of 0 - 100. The default is 50.

Specify 0 if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the target volume until the blocks or tracks of the source volume are modified.

You can modify this value at any time during the session. If the session is performing a background copy when you change the option, IBM Tivoli Storage Productivity Center for Replication immediately modifies the background copy rate of the consistency group on the storage system. The storage system consistency group immediately starts by using this new rate to complete the background copy.

Metro Mirror and Global Mirror Failover/Failback with Practice session properties:

Use the View/Modify Properties page to view or modify properties for SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified Metro Mirror Failover/Failback with Practice and Global Mirror Failover/Failback with Practice sessions.

The following options are on this page.

Description
Enter a description for the session.

SAN Volume Controller / Storwize Family / Storwize V7000 Unified Metro Mirror Options
Select the options for this session.

Incremental
Select this option to set up the relationship for recording changes to the practice volume (H2). All subsequent FlashCopy operations between the intermediate volume and the host volume copy only the data that changed since the previous FlashCopy operation.

Background Copy Rate for H2-I2
Type the copy rate that the storage system uses to perform the background copy of the FlashCopy role pair. You can specify a percentage in the range of 0 - 100. The default is 50.

Global Mirror Failover/Failback with Change Volumes session properties:

Use the View/Modify Properties page to view or modify properties for SAN Volume Controller, Storwize V3700, Storwize V7000, and Storwize V7000 Unified Global Mirror Failover/Failback with Change Volumes sessions.

The following options are on this page.

Description
Enter a description for the session.

SAN Volume Controller / Storwize Family / Storwize V7000 Unified Global Mirror Options
Select the options for this session.

Enable change volumes
Select this option to enable the use of change volumes in the copy sets for the session. If this option is not selected, the session functions the same as a Global Mirror Failover/Failback session.

Change volumes are denoted as C*x*, where *x* identifies the site. These volumes contain point-in-time images that are copied from the H1 and H2 volumes. The C1 volume stores changes from the H1 volume. These changes are sent from the C1 volume to the H2 volume, and then to the C2 volume.

Because the data that is replicated between sites contains point-in-time changes rather than all changes, a lower bandwidth link is required between the sites when change volumes are used. However, the use of change volumes can result in an increase to data exposure. Therefore, you might want to enable or disable this option depending on your network traffic or business requirements.

When you select this option, it is the equivalent of setting the **-cyclingmode** parameter to `multi` on a SAN Volume Controller or Storwize storage system. When you clear this option, it is the equivalent of setting the **-cyclingmode** parameter to `none` for these storage systems.

You can select this option and enter a value in the **Cycle period (seconds)** field only when the session is in an inactive state. An example of an inactive state is Defined or Suspended.

Cycle period (seconds)
Specify the amount of time, in seconds, in which the change volumes will be refreshed with a consistent copy of the data. If a copy does not complete in the cycle period, the next cycle period will not start until the copy is complete. The range of possible values is 60 - 86400. The default is 300.

Recovery Point Objective Alerts
Specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time, in seconds, of data exposure that is acceptable if a disaster occurs.

Use the following options to set the RPO threshold values. For both options, you can specify an RPO threshold in the range of 0 - 172800 seconds. The default is 0 seconds, which specifies that no alerts are generated.

Warning level threshold (seconds)
Specify the number of seconds that you want to set for the warning level RPO threshold. If the RPO is greater than this value, an alert is generated.

If the value in this field is other than 0, it must be greater than the value in the **Cycle period (seconds)** field and less than the value in the **Severe level threshold (seconds)** field.

Severe level threshold (seconds)
Specify the number of seconds that you want to set for the severe level RPO threshold. If the RPO is greater than this value, an alert is generated and the session status changes to Severe.

If the value in this field is other than 0, it must be greater than the value in the **Warning level threshold (seconds)** field.

XIV system sessions

The properties for the XIV system sessions depend on the session type.

The following sections describe the properties for each session type with the exception of session types that include only the **Description** property.

XIV system Global Mirror Failover/Failback session properties:

Use the View/Modify Properties page to view or modify properties for the XIV system Global Mirror Failover/Failback sessions.

The following options are on this page.

Description
Enter a description for the session.

XIV Global Mirror Options
Select the options for the session.

Recovery point objective threshold (seconds)
Type the number of seconds that you want to set for the recovery point objective (RPO) threshold for the role pair.

If the XIV system determines that the RPO is greater than this threshold value, an alert is generated and the session status changes to Severe. You can specify an RPO in the range of 30 - 86400 seconds. The default is 30 seconds.

Synchronization schedule (HH:MM:SS)
Select an interval to create a synchronization schedule. The schedule is used to form consistency groups. The XIV system attempts to form consistent points of data by taking automatic snapshots of the volumes in the session at this interval. The default is **Minimum Interval**, which is 20 seconds.

If you select **Never**, synchronization is not scheduled and the XIV system does not create consistency groups. When the XIV system determines that the RPO threshold is exceeded, the session state becomes Severe.

Basic HyperSwap session properties

Use the View/Modify Properties page to view or modify properties for Basic HyperSwap sessions.

The following options are on this page.

Description
Enter a description for the session.

Basic HyperSwap Options
Select the options for the session.

Disable HyperSwap
Select this option to prevent a HyperSwap operation from occurring.

System or sysplex
Select the IBM z/OS system that is connected to the storage system or the sysplex that contains the system. A system or sysplex is required to use HyperSwap.

Important: Tivoli Storage Productivity Center for Replication must be installed on or connected to one or more host systems to select a system or sysplex and to use HyperSwap.

On Configuration Error:

Partition the system(s) out of the sysplex
Select this option to partition a new system out of the sysplex when an error occurs because the system cannot be added to the HyperSwap configuration.

Disable HyperSwap
Select this option to prevent a HyperSwap operation from occurring.

On Planned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)
Select this option to partition out the failing system and continue the swap processing on any remaining systems.

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

On Unplanned HyperSwap Error:

Partition out the failing system(s) and continue swap processing on the remaining system(s)
Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and the HyperSwap operation does not complete.

Requirement: If you select this option, you must restart the system.

Disable HyperSwap after attempting backout
Select this option to enable IOS to back out the HyperSwap operation, if possible, if an error occurs during HyperSwap processing. HyperSwap is disabled.

Fail MM/GC if target is online (CKD only)
Select this option to fail any session commands for a Metro Mirror or Global Copy relationship if the target volume is in the Online state. For more information about this state, refer to the documentation for the storage system.

This option applies only to count key data (CKD) volumes.

Viewing copy sets in a session

You can view the copy sets that are in a session. You can also export the copy sets to a comma-separated value (CSV) file for backup purposes.

Follow these steps to view or export the copy sets for a session:

1. In the navigation tree, select **Sessions**.
2. On the Sessions page, select the session that contains the copy sets.
3. Select **View Copy Sets** from the **Select Action** list, and then click **Go**.

Viewing storage system details

You can view detailed information about storage system, including the name, location, type, vendor, and the status of all connections to the storage system.

Perform these steps to view storage system details:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Perform one of these steps to view details for a specific storage system:
   - Click the storage system ID.
   - Select the storage system, click **View storage system details** from the **Actions** list, and then click **Go**.

Viewing storage connection details

You can view storage connection details and a list of all storage systems that are located behind the connection.

Perform these steps to view storage connection details:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Click the **Connections** tab.
3. Perform one of these steps to view details for a specific storage connection:
   - Click the storage connection ID.
   - Select the storage connection, click **View/modify Connection Details** from the actions list, and then click **Go**.

Viewing volume details

You can view information about volumes such as the name of the volume, the capacity of the volume, and the type of volume.

1. In the navigation tree, select **Volumes**.
2. Select a storage.
3. Depending on the type of storage system, do one of the following actions:
   a. Select **All IO Groups** or a specific I/O group.
   b. Select **All Logical Storage Subsystems** or a specific logical storage subsystem.
   c. Select **All Pools** or a specific pool.
4. Click **Perform Query**. Information about the volumes is displayed in a table.

Viewing logical paths

You can view all logical paths that are defined on an IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, or IBM System Storage DS6000 storage system.

Perform one of these procedures to view logical paths:

- From the ESS/DS Paths panel of IBM Tivoli Storage Productivity Center for Replication:
  1. In the navigation tree, select **ESS/DS Paths**. The ESS/DS Paths panel is displayed.
  2. Click the storage system ID to display logical paths for that storage system.
- From the Storage Systems panel:
  1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
  2. Select an ESS, DS6000, or DS8000 storage system for which you want to view logical paths.
  3. Select **View Paths** from the **Select Action** list, and click **Go**. The ESS/DS Paths panel is displayed with a list of defined logical paths.

Viewing host connection details

You can view and modify the connection information for a host system. You can modify the connection information only if the host system in the Disconnected state.

Follow these steps to view or modify connection information for a host system:

1. In the navigation tree, select **Host Systems**.
2. On the Host Systems page, click the link for the host system in the **Host System** column.

Viewing console messages

This topic describes how to view the console and messages.

IBM Tivoli Storage Productivity Center for Replication provides detailed information about actions taken by users, errors that occur during normal operation, and hardware error indications.

From the graphical user interface, you can view console messages by selecting **Console** in the navigation tree. You can then click the link for the specific message code to get more information on the message.

You can also get detailed information and help for specific messages in the IBM Tivoli Storage Productivity Center information center at http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp .

Security

The IBM Tivoli Storage Productivity Center for Replication authentication process uses a configured user registry from either the operating system or Lightweight Directory Access Protocol (LDAP) server. To perform a specific action and manage specific sessions in the IBM Tivoli Storage Productivity Center for Replication GUI or CLI, the user must also have the appropriate authorization. Authorization is granted by assigning a specific role to the user account or user group.

Users and groups

To log on to Tivoli Storage Productivity Center for Replication, user ID and password is required.

The information in this topic is specific to Tivoli Storage Productivity Center for Replication on open systems. For information about defining users and groups for Tivoli Storage Productivity Center for Replication for System z, see the IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide.

Tivoli Storage Productivity Center for Replication does not maintain a directory of user names and passwords. Instead, the application uses one of the following repositories for user authentication:

- An operating system repository
- A Lightweight Directory Access Protocol (LDAP) repository

Tivoli Storage Productivity Center uses the federated repositories feature in WebSphere Application Server and an operating system repository is created by default during installation of Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication.

If you choose to use LDAP authentication, you must add the LDAP repository to the federated framework after the installation of Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication is complete.

For more information about the operating system and LDAP repositories, see "Changing the user authentication configuration" on page 146

Adding Tivoli Storage Productivity Center users and groups to Tivoli Storage Productivity Center for Replication

By default, the common user that was defined during the installation of Tivoli Storage Productivity Center and the user tpcFileRegistryUser are granted access to Tivoli Storage Productivity Center for Replication and assigned to the Administrator role.

The user tpcFileRegistryUser is used only for recovery purposes. For example, if you accidentally delete the repository that you are using for authentication, you can access Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication by using the user tpcFileRegistryUser.

The password for the user tpcFileRegistryUser is the same as the password that was entered for the common user during the installation of Tivoli Storage Productivity Center.

To ensure smooth integration of Tivoli Storage Productivity Center with Tivoli Storage Productivity Center for Replication, complete the following actions:

- Add all Tivoli Storage Productivity Center users and groups (other than those that are assigned by default) to Tivoli Storage Productivity Center for Replication. For example, if you added a TPCSuperuser group to Tivoli Storage Productivity Center for LDAP authentication, add that group to Tivoli Storage Productivity Center for Replication as well.
- Use the same user or group to log on to both applications.

User roles

A *user role* is a set of privileges that is assigned to a user or user group to allow the user or user group to perform certain tasks and manage certain sessions.

To be assigned to a role, each user or group of users must have a valid user ID or group ID in the user registry on the management server.

Both individual users and a group of users can be assigned to a role. All users in a group are assigned the role of the group. If a user is assigned to one role as an individual and a different role as a member of a group, the user has access to the permissions of the role that has greater access.

Restricting access to sessions prevents unwarranted administrative access. This is especially useful in an open environment, where there can be many storage administrators who are responsible for their servers, applications, databases, file systems, and so on.

Tivoli Storage Productivity Center for Replication provides a set of predefined user roles: Monitor, Operator, and Administrator.

Monitor

Monitors can view the health and status in the Tivoli Storage Productivity Center for Replication GUI and CLI; however, they cannot modify or perform any commands or actions.

Monitors can view the following information:

- All storage systems and storage system details
- All connections and connection details
- All sessions and session details
- All path information
- Management server status and details

Operator

Operators can manage sessions to which they have been assigned, including:

- Adding or removing a session. The user ID that created the session is automatically granted access to manage that session.
- Performing actions on an assigned session, such as start, flash, terminate, and suspend.
- Modifying session properties.
- Adding copy sets to a session. The session operator can add volumes to a copy set only when the volume is not protected and not in another session.
- Removing copy sets from a session.
- Adding Peer To Peer Remote Copy (PPRC) paths, and removing paths with no hardware relationships. PPRC paths are a common resource used in Tivoli Storage Productivity Center for Replication sessions and also in an ESS, DS6000, or DS8000 storage system relationship that is established between two common logical subsystems (LSSs).

  Note: The session operator cannot issue a force removal of a path.

  Note: A path can also be auto-generated when starting a session.
- Monitoring health and status, including viewing the following information:
  - All storage systems and storage system details
  - All connections and connection details

- All sessions and session details
- All path information
- Management server status and details

Note: Operators can make changes only to the volumes that they own. They are not able to make changes to volumes being managed by other users.

Administrator

Administrators have unrestricted access. They can manage all sessions and perform all actions associated with Tivoli Storage Productivity Center for Replication, including:

- Granting permissions to users and groups of users.
- Adding or removing a session. The user ID that created the session is automatically granted access manage that session.
- Performing actions on all sessions, such as start, flash, terminate, and suspend.
- Modifying session properties.
- Adding and removing copy sets from a session. The administrator can add volumes to a copy set only when the volume is not protected and not in another session.
- Protecting volumes and removing volume protection.
- Adding or removing storage system connections.
- Modifying connection properties.
- Assigning or changing storage system locations.
- Adding PPRC paths and removing paths with no hardware relationships. PPRC paths are a common resource used in Tivoli Storage Productivity Center for Replication sessions and also in an ESS, DS6000, or DS8000 storage-system relationship that is established between two common logical subsystems (LSSs).

  Note: A path can also be auto-generated when starting a session.
- Managing management servers. The standby management server is a common resource that is available to multiple sessions.
- Packaging program error (PE) log files.
- Monitoring health and status, including viewing the following information:
  - All storage systems and storage system details
  - All connections and connection details
  - All sessions and session details
  - All path information
  - Management server status and details

Note: Administrators cannot revoke their own administrative access rights.

Granting access privileges for a user

You can assign user roles to an IBM Tivoli Storage Productivity Center for Replication user to grant access privileges to individual sessions and tasks.

Perform the following steps to authorize a user:

1. Create the user ID or group ID if it does not already exist in the user registry, either the operating system of the active management server or Lightweight Directory Access Protocol (LDAP) server.

2. Log in to IBM Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
3. In the navigation tree, select **Administration**. The Administration panel is displayed.
4. Click **Add Access**. The Add Access wizard is displayed.
5. Type the name of the user to whom you want to give access, and click **Next**. The Select Users and Groups panel is displayed.

   Tip: You can enter a partial name and use the * wildcard character to represent zero or more characters.
6. Select one or more names from the list of found users.
7. Select the role to associate with this user.
8. If you selected the Operator role, select one or more session that this user can manage, and click **Next**.
9. Click **Next** to confirm this action.
10. Click **Finish**.

Viewing access privileges for a user

You can view a list of all IBM Tivoli Storage Productivity Center for Replication user and their assigned roles. You can also view the assigned sessions for each user.

Perform the following steps to view access privileges for a user:

1. Log in to IBM Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the navigation tree, select **Administration**. The Administration page is displayed with a list of IBM Tivoli Storage Productivity Center for Replication users and user groups and their associated role.
3. Select the user whose access privileges you want to view.
4. Select **View/Modify Access** from the **Actions** drop-down list, and click **Go**. The View/Modify Access panel is displayed. This panel shows the role assigned to the user and lists the sessions that the user can manage.
5. Click **Cancel**.

Modifying access privileges for a user

You can change the user role and assigned sessions for an IBM Tivoli Storage Productivity Center for Replication user.

Prerequisite: You must have Administrator privileges to perform this action.

Perform the following steps to modify the access privileges for a user:

1. Log in to IBM Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the navigation tree, select **Administration**. The Administration panel is displayed with a list of users and user groups and their associated role
3. Select the user whose access privileges you want to view.
4. Select **View/Modify Access** from the **Select Action** drop-down list, and click **Go**. The View/Modify Access panel is displayed. This panel shows the role assigned to the user and lists the sessions that the user can manage.
5. Select the role to associate with this user.

6. If you selected the Operator role, select one or more session that this user can manage and click **Next**.
7. Click **OK**.

Removing access privileges for a user

You can remove access privileges for an IBM Tivoli Storage Productivity Center for Replication user. When you remove access, the user ID cannot access the IBM Tivoli Storage Productivity Center for Replication GUI or run commands from the command line.

Prerequisite: You must have Administrator privileges to perform this action.

Perform the following steps to remove user access:

1. Log in to IBM Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the navigation tree, select **Administration**. The Administration panel is displayed with a list of users and user groups and their associated role.
3. Select the user from which you want to remove access.
4. Select **Remove Access** from the **Actions** list, and click **Go**.

Chapter 7. Reporting

Tivoli Storage Productivity Center provides multiple user interfaces for viewing reports about the storage infrastructure in an enterprise environment.

Tivoli Common Reporting interface
This interface runs in a web browser. Use this interface to view predefined reports and create custom reports about Tivoli Storage Productivity Center. You access reports from the web-based GUI, and work with the reports in Tivoli Common Reporting.

Tivoli Common Reporting includes Cognos reporting software.

Restriction: Cognos Workspace is not supported by Tivoli Storage Productivity Center reports in this release. If you use Tivoli Storage Productivity Center reports to create workspaces, those workspaces might be incompatible with future releases of the product.

Stand-alone GUI
View detailed information about the storage resources in your environment in the stand-alone GUI. These reports are organized into different types and categories and provide both summary and detailed information, depending on your needs. Reports are organized into report types and categories that you can expand to reveal extra subcategories of reporting. Many reports are also accessible through the topology viewer.

vSphere Web Client extension interface
Use the vSphere Web Client extension for Tivoli Storage Productivity Center to view reports on your virtual environment and storage devices. You can view reports that are customized to use information from Tivoli Storage Productivity Center. The reports include fabric connections, storage mapping information, and performance metrics for storage systems. To view Tivoli Storage Productivity Center storage information in block and file storage reports, you must register Tivoli Storage Productivity Center as a VASA provider.

Getting started

You can view predefined reports and create custom reports about Tivoli Storage Productivity Center in Tivoli Common Reporting.

You access reports from the Tivoli Storage Productivity Center web-based GUI, and work with the reports in Tivoli Common Reporting. Tivoli Common Reporting provides the reporting service for Tivoli Storage Productivity Center reports. Tivoli Common Reporting includes Cognos reporting software.

Predefined reports

You can use IBM Tivoli Common Reporting to view predefined reports about Tivoli Storage Productivity Center. Predefined reports provide information about the capacity and performance of your resources.

Predefined reports

You can view the following types of predefined reports:

© Copyright IBM Corp. 2001, 2013

Reports about the capacity and relationships of your resources
Reports about the capacity and relationships of servers, hypervisors, network-attached storage systems, storage virtualizers, storage systems, and storage resource groups. Use these reports to analyze the capacity and relationships of resources. For example, to show the current space statistics for a pool on a storage system, run the Pools Capacity report. To show the pools that are related to a storage system, run the Storage Resource Relationships Summary report.

Reports about performance
Reports about the performance of your monitored resources. You can view reports about storage systems, components of storage systems, switches, and switch ports. Use these reports to analyze the performance of resources. For example, to find out which are the most active volumes on a storage system over a specified period, run the Most Active Volumes report. To view details about the performance of a volume over a specified period, run a report such as the Performance of One Volume report.

Reports about historical capacity
Reports that include the historical information about the capacity of your storage systems, pools, and volumes. Use these reports to determine whether more storage resources are required. For example, to show how space is used over a specified period on volumes on storage systems, run the Volumes Historical Capacity report. You can also use these reports to analyze future storage requirements for your environment.

Reports about storage tiering
Reports that include information about the workload activity of these resources: managed disk (MDisk) groups, virtual disks (VDisks) in MDisk groups, VDisks in storage resource groups, and VDisks in storage virtualizers. Use these reports to determine the best storage-tier configuration for your environment.

Charts are automatically generated for most of the predefined reports. Depending on the type of resource, the charts show statistics for space usage, workload activity, bandwidth percentage, and other statistics.

Tip: When you view a predefined report, click **Help** to find out more about the report.

Related reference:

"Predefined reports listed by package"
The data that the predefined reports use is organized into four packages: Capacity and Relationships, Historical Capacity, Performance, and Storage Tiering. Each report uses data from only one package.

Predefined reports listed by package

The data that the predefined reports use is organized into four packages: Capacity and Relationships, Historical Capacity, Performance, and Storage Tiering. Each report uses data from only one package.

Tip: When you view a predefined report, click **Help** to find out more about the report.

Capacity and Relationships package

You can view the following predefined reports about capacity and relationships:

Disks Capacity
Shows a chart of the disk classes with the most allocated space, and a table with space statistics for all disks by class. Use the report to review space allocation in your disks.

File System to Volume Relationships
Shows the end-to-end relationship from file systems on servers to volumes on storage pools. The report shows the usage of space on file systems over a specified period. Use the report to analyze changes in your file system requirements.

File Systems Capacity
Shows the available and used capacity of file systems, which are grouped by server. You can use this report to identify file systems that might be running out of space.

Hypervisor Data Stores Capacity
Shows a chart of the 20 hypervisor data stores with the most allocated space. This report also shows a table with space statistics for all hypervisor data stores. Use the report to review space allocation in your hypervisor data stores.

Hypervisor Disks Capacity
Shows a chart of the 20 hypervisor disks with the most allocated space, and a table with space statistics for all hypervisor disks. Use the report to review space allocation in your hypervisor disks.

Hypervisors Capacity
Shows a chart of the 20 hypervisors with the most allocated space, and a table with space statistics for all hypervisors. Use the report to review space allocation in your hypervisors.

Managed Disks Capacity
Shows a chart of the 20 managed disks with the most available space, and a table with space statistics for all managed disks. Use the report to review space availability in your managed disks.

Pools Capacity
Shows a chart of the 20 pools with the most allocated space, and a table with space statistics for all pools. Use the report to review space allocation in your pools.

Server Disks Capacity
Shows the available and used disk space for all disks, which are grouped by server. You can use the report to identify disks that have no allocated space.

Servers Capacity
Shows a summary of the available and used disk space, and the available and used file systems, for servers. Use the report to identify servers that are running out of space.

Storage Resource Relationships Summary
Shows the end-to-end relationship from a server to back-end storage. The report shows servers, hypervisors, NAS, storage systems, storage virtualizers, pools, and volumes.

Storage Resource Relationships Summary (Configurable)
Shows the end-to-end relationship from a server to back-end storage. You can configure which resources to show in the report.

Storage Systems Capacity
Shows a chart of the 20 storage systems with the most used space, and a table with details of all storage systems. Use the report to review the space allocation in your storage systems.

Volumes Capacity
Shows a chart of the 20 volumes with the most used space, and a table with details of all volumes. Use the report to review the space allocation in your volumes.

Historical Capacity package

You can view the following predefined reports about historical capacity:

Pools Historical Capacity
Shows the usage of space on storage pools over a specified period. Use the report to analyze changes in your storage pool requirements.

Storage Systems Historical Capacity
Shows the usage of space on storage systems over a specified period. Use the report to analyze changes in your storage system requirements.

Volumes Historical Capacity
Shows the usage of space on volumes over a specified period. Use the report to analyze changes in your volume requirements, particularly the increasing usage of space on thin-provisioned volumes.

Performance package

You can view the following predefined reports about performance:

Compare Performance of Multiple Controllers and Modules
Shows up to four performance metrics for multiple controllers or modules over time. Use the report to compare the performance metrics for controllers or modules that are on storage systems.

Compare Performance of Multiple Disks
Shows up to four performance metrics for multiple local disks over time. Use the report to compare the performance metrics for local disks that are on storage systems.

Compare Performance of Multiple IO Groups
Shows up to four performance metrics for multiple I/O groups over time. Use the report to compare the performance metrics for I/O groups that are on storage systems.

Compare Performance of Multiple Managed Disks
Shows up to four performance metrics for multiple managed disks over time. Use the report to compare the performance metrics for managed disks that are on storage systems.

Compare Performance of Multiple Nodes
Shows up to four performance metrics for multiple nodes over time. Use the report to compare the performance metrics for nodes that are on storage systems.

Compare Performance of Multiple Pools
Shows up to four performance metrics for multiple storage pools over time. Use the report to compare the performance metrics for storage pools that are on storage systems.

Compare Performance of Multiple Ports
Shows up to four performance metrics for multiple ports over time. Use the report to compare the performance metrics for ports that are on storage systems.

Compare Performance of Multiple Storage Systems
Shows up to four performance metrics for multiple storage systems over time. Use the report to compare the performance metrics for storage systems.

Compare Performance of Multiple RAID Arrays
Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

Compare Performance of Multiple Switch Ports
Shows up to four performance metrics for multiple switch ports over time. Use the report to compare the performance metrics for switch ports.

Compare Performance of Multiple Switches
Shows up to four performance metrics for multiple switches over time. Use the report to compare the performance metrics for switches.

Compare Performance of Multiple Volumes
Shows up to four performance metrics for multiple volumes over time. Use the report to compare the performance metrics for volumes that are on storage systems.

Compare Performance of One Pool over Time Ranges
Shows one performance metric on one pool over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Compare Performance of One Storage System over Time Ranges
Shows one performance metric on one storage system over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Compare Performance of One Switch over Time Ranges
Shows one performance metric on one switch over two time periods. For example, you can use the report to compare the I/O rate for this week and last week.

Compare Performance of One Switch Port over Time Ranges
Shows one performance metric on one switch port over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Most Active Controllers or Modules
Shows a chart of the 20 controllers or modules that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of controllers or modules.

Most Active Disks
Shows a chart of the 20 local disks that are most active, and details for all local disks, for a time period that you specify. Use the report to analyze the performance of local disks.

Most Active Host Connections
Shows a chart of the 20 host connections that have the most active volumes. The most active volumes have the greatest aggregate load.

Most Active Hypervisors
Shows a chart of the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load.

Most Active IO Groups
Shows a chart of the 20 I/O groups that are most active, and details for all I/O groups, for a time period that you specify. Use the report to analyze the performance of I/O groups.

Most Active Managed Disks
Shows a chart of the 20 managed disks that are most active, and details for all managed disks, for a time period that you specify. Use the report to analyze the performance of managed disks.

Most Active Nodes
Shows a chart of the 20 nodes that are most active, and details for all nodes, for a time period that you specify. Use the report to analyze the performance of nodes.

Most Active Pools
Shows a chart of the 20 storage pools that are most active, and details for all pools, for a time period that you specify. Use the report to analyze the performance of pools.

Most Active Ports
Shows a chart of the 20 ports that are most active, and details for all ports, for a time period that you specify. Use the report to analyze the performance of ports.

Most Active RAID Arrays
Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

Most Active Servers
Shows a chart of the 20 servers that have the most active volumes. The most active volumes have the greatest aggregate load.

Most Active Storage Systems
Shows a chart of the 20 storage systems that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of storage systems.

Most Active Switch Ports
Shows a chart of the 20 switches that are most active, and details for all switches, for a time period that you specify. Use the report to analyze the performance of switches.

Most Active Volumes
Shows a chart of the 20 volumes that are most active, and details for all volumes, for a time period that you specify. Use the report to analyze the performance of volumes.

Performance Data Export
Exports the performance metrics for one or more resources on a storage system to a spreadsheet file in Microsoft Excel file format. Use the report if you want to analyze performance data in Excel or to send performance data to IBM.

Performance of One Controller or Module
Shows four charts, and a table of performance metrics for a controller or module over a time period. Use the report to analyze the performance of a controller or module.

Performance of One Disk
Shows four charts, and a table of performance metrics for a local disk over a time period. Use the report to analyze the performance of a local disk.

Performance of One IO Group
Shows four charts, and a table of performance metrics for an I/O group over a time period. Use the report to analyze the performance of an I/O group.

Performance of One Managed Disk
Shows four charts, and a table of performance metrics for a managed disk over a time period. Use the report to analyze the performance of a managed disk.

Performance of One Node
Shows four charts, and a table of performance metrics for a node over a time period. Use the report to analyze the performance of a node.

Performance of One Pool
Shows four charts, and a table of performance metrics for a storage pool over a time period that you specify. Use the report to analyze the performance of a storage pool.

Performance of One Port
Shows four charts, and a table of performance metrics for a port over a time period that you specify. Use the report to analyze the performance of a port.

Performance of One RAID Array
Shows four charts, and a table of performance metrics for a RAID array over a time period. Use the report to analyze the performance of a RAID array.

Performance of One Storage System
Shows four charts, and a table of performance metrics for a storage system over a time period that you specify. Use the report to analyze the performance of a storage system.

Performance of One Switch
Shows three charts, and a table of performance metrics for a switch over a time period that you specify. Use the report to analyze the performance of a switch.

Performance of One Switch Port
Shows four charts, and a table of performance metrics for a switch port. Use this report to view multiple performance metric types on the same chart. For example, you can view I/O rate and response time metrics on the same chart.

Performance of One Volume
Shows four charts, and a table of performance metrics for a volume over a time period that you specify. Use the report to analyze the performance of a volume.

Performance of Volumes by Host Connection
Shows the performance metrics for volumes on a host connection. Use the

report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

Performance of Volumes by Hypervisor
Shows the performance metrics for volumes on a hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment.

Performance of Volumes by Server
Shows the performance metrics for volumes on a server. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

Summarized Performance of Volumes by Host Connection
Shows summarized performance metrics for volumes on a host connection. Use the report to identify the aggregate load and the average response time of a particular server.

Summarized Performance of Volumes by Hypervisor
Shows summarized performance metrics for volumes on a hypervisor. Use the report to identify the aggregate load and the average response time of a particular hypervisor.

Summarized Performance of Volumes by Server
Shows summarized performance metrics for volumes on a server. Use the report to identify the aggregate load and the average response time of a particular server.

Storage Tiering package

You can view the following predefined reports about storage tiering:

MDisk Group - VDisk Workload Activity Details
Shows the workload activity of virtual disks (VDisks) in a managed disk (MDisk) group. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

MDisk Group Details
Shows the workload activity of a managed disk (MDisk) group. Use the report to analyze storage services for an MDisk group.

MDisk Groups - VDisk Workload Activity
Use this report to monitor the performance of virtual disks (VDisks) in managed disk (MDisk) groups. The workload activity of each MDisk group is determined by calculating the average and the maximum peak utilization of the VDisks in the MDisk group.

MDisk Groups - Workload Activity
Shows the workload activity of managed disk (MDisk) groups. You can use the information that the report provides to investigate which MDisk groups are being underused or overused.

Storage Resource Group - VDisk Workload Activity
Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDisks in the group.

Storage Resource Group - VDisk Workload Activity Details
Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDisks in the group.

Storage Virtualizer - VDisk Workload Activity
Shows the workload activity of virtual disks (VDisks) in a storage virtualizer. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

VDisk Details
Shows the workload activity of a virtual disk (VDisk) in a managed disk group. Use the report to analyze the performance of a VDisk in a managed disk (MDisk) group.

Custom reports

You can use IBM Tivoli Common Reporting to create custom reports about Tivoli Storage Productivity Center. Use custom reports to analyze the capacity and performance of your resources.

Custom reports

Unlike predefined reports, you can select the information that you want to include in custom reports. Custom reports can contain detailed information about the relationships between monitored resources, the properties of monitored resources, and detailed information about the performance of monitored resources.

To create a report, browse the list of monitored resources in your storage environment, and drag items that are associated with the resources into a report. For example, drag items that show capacity data for storage systems into a report.

If you create custom reports, you can also do the following tasks:

- Copy a predefined report and modify the copy to create a custom report.
- Create custom charts. You can create charts with data that is in the combinations that you want to see. Use these custom charts to troubleshoot and to report on the status of resources.
- Create reports with multiple charts. You can create multiple charts with different scales, and display those charts side-by-side. You can use these charts to compare performance metrics. For example, you can review read I/O rates and the number of bytes transferred in the same report.
- Create a custom folder structure to organize your custom reports.
- Run a report from a URL. You can include options for the report in the URL. For options that are not included in the URL, you can specify default values, or prompt the user to specify a value.
- Share a custom report with your organization, or keep it for only your own use. You can also distribute different sections of a report to different teams.

Advanced formatting for custom reports

You can create reports with basic formatting, or design more complex reports with advanced formatting.

Access related reports
You can configure custom reports so that you can click a resource in the report output to access other related reports. If you identify a problem in a report, you can open another report from within that report that might provide specific information about the problem.

For example, the Most Active Servers predefined report shows data for servers. Click the name of a server in the report to view the Summarized

Performance of Volumes by Server report for the server. This report shows the aggregated data for all of the volumes on the server. Click the name of a volume in the report to view the Performance of Volumes by Server report for the volume.

Use GUI controls in reports

You can configure custom reports to include GUI controls to enable users to select options for the reports. You can specify default values for these options. In addition, you can create copies of reports with different default values.

The GUI controls that you can add to your reports enable users to select dates, type values, and select from lists. You can also create GUI controls whose values depend on what is selected in another GUI control. You can include GUI controls on the report output in HTML format so that users can change options and run the report again. For example, you can include a drop-down list in the HTML report output so that users can change the sort order of the data.

Schedule reports

You can schedule when to run the report. You can specify the following formats for the report: CSV, Excel, HTML, PDF, and XML. You can specify to email the report output and to save the report output to a location on your file system or on a web server. For example, you might require a monthly report that shows the usage of space on storage systems. Management staff might use this report to analyze changes in your storage system requirements. You can create this report and schedule it to run every month. You can specify that the report output is emailed to management staff, or saved to a location from which management staff can access it.

Viewing predefined reports

Use the Tivoli Storage Productivity Center web-based GUI to access predefined reports.

1. Start the web-based GUI for Tivoli Storage Productivity Center.
2. Select **Reporting** > **View predefined reports**.
3. Navigate to the report that you want to see, and then click the name of the report. To view a description of a report or folder, point to the report icon that is next to the report name. To view descriptions of all reports and folders, click the **Details View** icon .

Creating custom reports

Use the Tivoli Storage Productivity Center web-based GUI to access custom reports. You can create basic reports with limited formatting, or design custom reports to add advanced formatting features to reports.

Restriction:

You cannot create custom reports from the following packages:

- Historical Capacity
- Storage Tiering

1. Start the web-based GUI for Tivoli Storage Productivity Center.
2. Depending on the type of report that you want to see, open the appropriate report tool:

- To create a basic report with limited formatting about the capacity and relationships of resources, select **Reporting** > **Create capacity reports**.
- To create a basic report with limited formatting about the performance of resources, select **Reporting** > **Create performance reports**.

3. Optional: To create a report with advanced formatting features, complete the following steps:
   a. Click the **Return** icon .
   b. Select **Launch** > **Report Studio**.

Related tasks:

"Creating custom performance reports" on page 1096
You can create reports that show performance metrics. Performance reports can contain detailed information about the performance of monitored resources and some of the properties of those resources. You access custom performance reports from the Tivoli Storage Productivity Center web-based GUI, and you create the reports in Tivoli Common Reporting.

"Creating custom capacity and relationship reports" on page 1219
You can create custom reports that show capacity and relationship information for resources. Capacity and relationship reports can contain detailed information about the properties of monitored resources, and the available space and capacity of those resources. You access custom capacity and relationship reports from the Tivoli Storage Productivity Centerweb-based GUI, and you create the reports in Tivoli Common Reporting.

"Adding filters to reports" on page 1220
You can use filters with capacity and relationship reports. Filters allow you to show only certain data that complies with your criteria. You add filters to your capacity and relationship reports when you create the reports.

Viewing and creating reports in Tivoli Common Reporting

You access reports from the Tivoli Storage Productivity Center web-based GUI, and work with the reports in Tivoli Common Reporting.

When you view and create reports, you can navigate from predefined reports to custom reports and from custom reports to predefined reports within Tivoli Common Reporting. You do not have to return to the web-based GUI.

1. From the Tivoli Common Reporting user interface, click the **Home** icon , if you are not on the Tivoli Common Reporting home page.
2. Depending on the type of report that you want to work with, complete one of the following actions on the **Public Folders** tab:
   - To view predefined reports, click **IBM Tivoli Productivity Center Predefined Reports**.
   - To create a basic custom report with limited formatting, complete the following steps:
     a. Click **IBM Tivoli Productivity Center Packages**.
     b. Click the package from which you want to create a report.
     c. Select **Launch** > **Query Studio**.
   - To create a report and apply advanced formatting features to the report, complete the following steps:
     a. Click **IBM Tivoli Productivity Center Packages**.
     b. Click the package from which you want to create a report.

c. Select **Launch** > **Report Studio**.

Searching the predefined reports for a property or performance metric

You can search the contents of the predefined reports. You can check which reports contain a property or performance metric. For example, you can list the reports that contain the Storage Pool Name property or the Port Receive I/O Rate performance metric.

Before you can search the content of predefined reports, a search index must be created for the reports.

Tip: If there is no search index, you can search only the titles and descriptions of the reports, not the contents. To search the titles and descriptions of reports, select **Name field**, **Description field**, or **Name or description field** from the search options list.

1. In the search field beside the **Home** icon in Tivoli Common Reporting, type the name of the property or metric in quotation marks. For example, type "Storage Pool Name" or "Port Receive I/O Rate".
2. Select **Full text and all fields** from the search options list.
3. Click the search icon. The list of reports that contain the property or metric is displayed in the search results.
4. Optional: Alternatively, you can search the online help for properties or metrics. To search the online help, complete the following steps:
   a. In Tivoli Common Reporting, click any predefined report.
   b. Click the **Help** link at the upper right of the page.
   c. In the **Search** field, type the name of the property or metric in quotation marks, and then click **Go**. The list of help topics that contains the property or metric is displayed in the search results.

Related tasks:

"Creating a search index of reports, report views, and folders" on page 811
You can create a search index of the reports, report views, and folders in Tivoli Common Reporting.

Related information:

IBM Cognos Business Intelligence 10.2.0 Information Center
For more information about searching for items in Tivoli Common Reporting, see the IBM Cognos Business Intelligence 10.2.0 Information Center at pic.dhe.ibm.com/infocenter/cbi/v10r2m0/index.jsp. Search for "Searching for entries using full text and all fields".

Scenarios

Use these scenarios to learn how to use Tivoli Storage Productivity Center reports to analyze the most active volumes on hypervisors and to monitor the performance of a critical application. You can learn how to investigate a degradation in the performance of a storage pool. You can also learn how to identify the relationships between a storage system and other resources.

Related information:

IBM Tivoli Storage Productivity Center Information Center
For information about the concepts behind the tasks that are associated with these scenarios, see the Tivoli Storage Productivity Center Information Center at http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp.

Analyzing the volumes on the most active hypervisors

In this scenario, you use the predefined reports for hypervisors to analyze the volumes on hypervisors. You can view the most active volumes. Then, you can view summary information for a hypervisor and details for individual volumes. You can also use this scenario to analyze the volumes on servers, agentless servers, and cluster resource groups.

You have 10-20 volumes on the hypervisors in your storage environment. You receive reports that the performance of some volumes is slow. You want to analyze the performance of these volumes as a group, and the performance of individual volumes.

1. Run the Most Active Hypervisors report. Use the chart in the report output to see the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load. You see that the volumes on one hypervisor are considerably more active than the volumes on other hypervisors. You want to investigate that hypervisor.
2. To analyze the activity on this hypervisor, click the name of the hypervisor in the report table. The Summarized Performance of Volumes by Hypervisor report runs for the hypervisor. Use the report to see summarized performance metrics for the volumes on the hypervisor. You can also use the report to see the aggregate load and the average response time of the volumes on the hypervisor. In the charts and the table on the report, you see that the I/O response times for the hypervisor are slower than normal.
3. To investigate the individual volumes on the hypervisor, click the name of the hypervisor in the report table. The Performance of Volumes by Hypervisor report runs for the hypervisor. Use the report to see the performance metrics for individual volumes on the hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment. In the charts and the table on the report, you see that the I/O response times for one volume are slow.

You identified the volume with the slow I/O response times. You can now copy data to the other volumes to more evenly distribute the load of the volumes.

Related reference:

"Most Active Hypervisors report" on page 951
Shows a chart of the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load.

"Summarized Performance of Volumes by Hypervisor report" on page 953
Shows summarized performance metrics for volumes on a hypervisor. Use the report to identify the aggregate load and the average response time of a particular hypervisor.

"Performance of Volumes by Hypervisor report" on page 955
Shows the performance metrics for volumes on a hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment.

Monitoring daily the performance of volumes

In this scenario, you create a report view from a predefined report to run a daily report about the performance of volumes. You can also use this scenario to monitor any other resources and performance metrics.

You have applications that are critical to your business. The applications access volumes in your storage environment. You want to monitor daily the performance of the volumes to confirm that they are functioning correctly. To monitor the volumes, you must create a report view to run a daily report about the performance of volumes.

1. In the Tivoli Storage Productivity Center web-based GUI, select **Reporting** > **View predefined reports**.
2. Click **Storage Systems**, and then click **Volumes**.
3. Click the **Report View** icon for the Compare Performance of Multiple Volumes report.
4. Type a name for the report view, then click **Finish**.
5. Click the **Schedule** icon for the report view.
6. Check all seven days of the week on the **By Week** tab.
7. In the **Options** area, check **Override the default values**, and select **HTML** and **PDF** formats.
8. In the **Prompt values** area, check **Override the default values**, and then click **Set**.
9. Specify the storage systems and the volumes that you want to monitor, and then click **Next**.
10. Select the following performance metrics:
    - Total Overall I/O Rate (ops/s)
    - Overall Response Time (ms/op)
    - Total Data Rate (MiB/s)
    - Total Overall Cache Hit Percentage
11. Specify an interval and a reporting period for the report view, and then click **Finish**.
12. Click **OK** to save the schedule, format, and prompt value settings.

View the output of this report every day to confirm that these volumes are functioning correctly. Click the report view to view the latest output for the report. If the report output is not displayed, make sure that the **Default action** for report view in the properties of the report view is set to **View most recent report**.

Related reference:

"Compare Performance of Multiple Volumes report" on page 1058
Shows up to four performance metrics for multiple volumes over time. Use the report to compare the performance metrics for volumes that are on storage systems.

Investigating a degradation in the performance of a storage pool

You can use the Compare Performance of One Pool over Time Ranges report to analyze a change in the performance of a storage pool. You can also use this scenario to compare the performance metrics of storage systems, switches, and switch ports over two periods of time.

In the previous three days, the performance of a storage pool has significantly degraded. You want to investigate the reasons for the degradation. To do this, use the Compare Performance of One Pool over Time Ranges report to compare the recent poor performance of the pool with a period when the performance was normal.

1. In the Tivoli Storage Productivity Center web-based GUI, select **Reporting** > **View predefined reports**.
2. Click **Storage Systems**, and then click **Pools**.
3. Click **Compare Performance of One Pool over Time Ranges**.
4. Specify the storage system and the pool whose performance has degraded.
5. Specify the performance metric that you want to analyze, for example **Overall Response Time (ms/op)**.
6. Specify an interval.
7. Specify two custom date ranges for the two reporting periods. Specify the previous three days as one reporting period. For the other reporting period, specify a period when the performance of the pool was normal.
8. Click **Finish**.

To investigate other possible reasons for the degradation, run the report again, and select another metric.

Related reference:

"Compare Performance of One Pool over Time Ranges report" on page 1034
Shows one performance metric on one pool over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Identifying the relationships between a storage system and other resources

In this scenario, you use a predefined relationships report to identify the relationships between a storage system and other servers or hypervisors. You can also use this scenario to identify the relationships between storage systems and NAS, pools, and volumes. Use this scenario to identify the relationships between storage virtualizers and servers, hypervisors, NAS, pools, and volumes.

You want to take a storage system offline and migrate the volumes to another storage system. To prepare for the migration, you first identify which servers or hypervisors are affected when the storage system goes offline.

1. In the Tivoli Storage Productivity Center web-based GUI, select **Reporting** > **View predefined reports**.
2. Click **Storage Resource Relationships Summary**.
3. From the **Storage systems** list, select the storage system that you plan to take offline.
4. Click **Finish**.

View the output of the report to see which servers or hypervisors are affected when the storage system goes offline. You can now migrate the volumes.

Related reference:

"Storage Resource Relationships Summary report" on page 931
Shows the end-to-end relationship from a server to back-end storage. The report shows servers, hypervisors, NAS, storage systems, storage virtualizers, pools, and volumes.

Configuring

You can configure Tivoli Storage Productivity Center reports so that you can save report outputs to your local file system, or as attachments in emails. You can also configure Tivoli Storage Productivity Center reports so that you can run reports for multiple Tivoli Storage Productivity Center servers from one Tivoli Common Reporting server.

Setting options for reports

When you schedule or run a report, you can specify the output format and how the report is delivered.

You can select output formats for the report. For example, you can select PDF, HTML, and other formats. If you choose PDF as the format for your report, you can set options such as the orientation and paper size for reports.

You can also specify to print the report, to send the report by email, and other delivery options.

1. Navigate to a report in Tivoli Common Reporting.
2. Depending on whether you want to schedule a report or run a report, do one of the following steps:
   - To schedule a report, and select format and delivery options, click the **Schedule** icon .
   - To run a report, and select format and delivery options, click the **Run with options** icon . To set advanced options, click **advanced options**, and then select **Run in the background**.

Related information:

IBM Cognos Business Intelligence 10.2.0 Information Center

For more information about report options, see the IBM Cognos Business Intelligence 10.2.0 Information Center at pic.dhe.ibm.com/infocenter/cbi/v10r2m0/index.jsp. To schedule a report, search for “Schedule an Entry”. To select formats for reports, search for “Report Formats”. To select delivery options for reports, search for “Distributing Reports”. To set advanced report options, search for “Advanced Report Options”.

Scheduling a report to run with predefined options

You can schedule a report to run on different systems and with different options that you define. To monitor the performance of different sets of volumes, you can schedule a Compare Performance of Multiple Volumes report for each of the sets.

You can specify schedules and output formats for the reports, and different ways to share the reports.

To schedule a report to run with parameters that you specify, you first create a report view from the report. Then, you schedule the report and set the parameters and options in the report view.

If the report from which a report view was created is updated, the report view is also updated.

1. In the Tivoli Storage Productivity Center web-based GUI, select **Reporting** > **View predefined reports**.

2. Click the **Report View** icon for the report.
3. Type a name for the report view, then click **Finish**.
4. Click the **Schedule** icon for the report view.
5. In the schedule area, specify when and with what frequency you want the report to run.
6. In the **Options** area, check **Override the default values**, and select the formats for the report.
7. In the **Prompt values** area, check **Override the default values**, and then click **Set**.
8. Specify parameters for the report. For example, you can specify the resources, the interval, and the reporting period for the report.
9. Click **OK** to save the schedule, format, and prompt value settings.

After the report runs for the first time, click the report view to view the latest output for the report. If the report output is not displayed, make sure that the **Default action** for report view in the properties of the report view is set to **View most recent report**.

Related tasks:

"Monitoring daily the performance of volumes" on page 802
In this scenario, you create a report view from a predefined report to run a daily report about the performance of volumes. You can also use this scenario to monitor any other resources and performance metrics.

Related information:

IBM Cognos Business Intelligence 10.2.0 Information Center
For more information about report views, see the IBM Cognos Business Intelligence 10.2.0 Information Center at pic.dhe.ibm.com/infocenter/cbi/v10r2m0/index.jsp. Search for "Create a Report View".

Creating a copy of a predefined report

Create a copy of a predefined report if you want to modify report. You might want to modify the queries that the report uses, or to add or remove a column. Changes to the predefined report do not affect copies of the report.

When you create a copy of a report, the links to related reports in the report output do not work correctly.

If the predefined report is changed during an upgrade, the copy of the report is not affected. If you modify a predefined report instead of creating a copy and then you upgrade your reports software, your modifications might be overwritten.

1. In the Tivoli Storage Productivity Center web-based GUI, select **Reporting** > **View predefined reports**.
2. Navigate to the report that to want to copy, and then click the check box that is next to the title of the report.
3. Click the **Copy** icon .
4. Optional: Navigate to a folder where you want to store the copy of the report. For example, **My Folders**.
5. Click the **Paste** icon .
6. Optional: To change the name of the report, click the **Properties** icon , type the name in the **Name** field, and then click **OK**.

Related information:

IBM Cognos Business Intelligence 10.2.0 Information Center
For more information about copies of reports, see the IBM Cognos Business Intelligence 10.2.0 Information Center at pic.dhe.ibm.com/infocenter/cbi/v10r2m0/index.jsp. Search for “Copy an Entry”.

Starting the IBM Cognos Configuration GUI

You can use the IBM Cognos Configuration GUI to configure IBM Cognos components. For example, you can configure an archive location for the output of reports, and configure a connection between Tivoli Common Reporting and an SMTP server.

To start IBM Cognos Configuration GUI, follow the instructions for your operating system:

Windows
Click **Start** > **Tivoli Common Reporting 3.1.0.1** > **Cognos Configuration**.

Alternatively, you can run the `tcr_cogconfig.bat` file in the *JazzSM_install_directory*`\reporting\cognos\bin64` directory.

AIX or Linux
Log on to a graphical session, and then run the `tcr_cogconfig.sh` script in the *JazzSM_install_directory*`/reporting/cognos/bin64` directory.

Sending reports by email

To send reports as attachments in email messages, configure a connection between Tivoli Common Reporting and an SMTP server.

If you use Red Hat Enterprise Linux 5, you must set the *LD_PRELOAD* environment variable to open the IBM Cognos Configuration GUI. To set the environment variable, issue the following command:

```
setenv LD_PRELOAD /usr/lib64/libfreetype.so
```

You use the IBM Cognos Configuration GUI to configure a connection between Tivoli Common Reporting and an SMTP server. When you configure the connection to the SMTP server, you can send the custom and predefined reports that you create as email attachments.

1. Open the IBM Cognos Configuration GUI.
2. On the Explorer pane, click **Data Access**.
3. Click **Notification**, and then enter values in the following fields:

   SMTP Mail Server
   Enter the host name and port number of the mail server.

   Account and password
   Leave the field blank if authentication is not required to log on to the mail server. If authentication is required to log on to the mail server, click the **Edit** icon. Enter a valid user ID and password, and then click **OK**.

   Default sender
   Enter the email address of the sender.
4. To ensure that the connection to the mail server is correctly configured, right-click **Notification**, and then click **Test**.
5. Click **Save**.

6. On the Explorer pane, right-click **Local Configuration**, and then click **Restart**.

Related tasks:

"Starting the IBM Cognos Configuration GUI" on page 806
You can use the IBM Cognos Configuration GUI to configure IBM Cognos components. For example, you can configure an archive location for the output of reports, and configure a connection between Tivoli Common Reporting and an SMTP server.

Specifying locations for saving reports

You can configure reports to save the report output to your local file system. You specify a root directory in the file system in which the output of reports can be saved in the IBM Cognos Configuration GUI. You specify the locations in which you want to save the output of reports. Then, you can configure individual reports to save the report output to a location in the file system. For example, you can schedule a report and configure the report output to be saved to a location like `E:\tpcreports\sales` or `/usr/tpcreports/sales`.

Specifying a root directory for saving reports

You can specify a root directory in the file system in which the output of reports can be saved. To do this, you configure an archive location in IBM Cognos Configuration.

1. Open the IBM Cognos Configuration GUI.
2. Select **Actions** > **Edit Global Configuration**.
3. Click the **General** tab.
4. In the **Archive Location File System Root** field, enter a file URI. For example, enter `file:///e:/tpcreports` or `file:///usr/tpcreports`.
5. Click **OK**.
6. Select **Actions** > **Restart**.

Related tasks:

"Starting the IBM Cognos Configuration GUI" on page 806
You can use the IBM Cognos Configuration GUI to configure IBM Cognos components. For example, you can configure an archive location for the output of reports, and configure a connection between Tivoli Common Reporting and an SMTP server.

Specifying file system locations for reports

You can specify the locations in which you want to save the output of reports in Tivoli Common Reporting.

You must configure the root directory of the locations that you specify in IBM Cognos Configuration.

1. In Tivoli Storage Productivity Center, select **Reporting** > **View Predefined Reports**.
2. From the **Launch** menu in Tivoli Common Reporting, click **Administration**.
3. Click the **Configuration** tab.
4. Click **Dispatchers and Services**.
5. Click the **Define File System Locations** icon .
6. On the Define file system locations page, click **New**.

7. Specify a name and the directory for the location on the file system, and then click **Finish**. For example, enter `sales`. The directory is created the first time that the output of a report for the directory is created. Repeat this step to configure more file system locations.
8. To stop Tivoli Common Reporting, follow the instructions for your operating system:
   - Windows

     ```
     c:\program files\IBM\JazzSM\profiles\bin\stopServer.bat server_name
      -username JazzSM_username -password JazzSM_password
     ```

   - AIX or Linux

     ```
     /opt/IBM/JazzSM/profile/bin/stopServer.sh server_name
      -username JazzSM_username -password JazzSM_password
     ```

9. To restart Tivoli Common Reporting, follow the instructions for your operating system:
   - Windows

     ```
     c:\program files\IBM\JazzSM\profiles\bin\startServer.bat server_name
     ```

   - AIX or Linux

     ```
     /opt/IBM/JazzSM/profile/bin/startServer.sh server_name
     ```

Specifying a report to save the report output to a file system

You can configure a report to save the report output to a location in a file system. You can also specify the file name for the report output, and whether to append a date and time to the file name.

You must configure the locations in which you want to save the output of reports in Tivoli Common Reporting.

1. In Tivoli Common Reporting, click the **Schedule** icon for a report.
2. Select **Override the default values**.
3. In the **Delivery** area, select **Save to the file system**.
4. Optional: Click **Edit the options** to specify the file name to use, and the location of the report on the file system. You can also specify what to do if a file name exists. For example, you can specify to use the report name as the file name. You can click **Make the file names unique and append a timestamp** to add the date and time to the file name.
5. Click **OK**.

Setting up multiple connections to the TPCDB data source

In Tivoli Common Reporting, you can add multiple connections to the Tivoli Storage Productivity Center data source so that you can run reports against multiple Tivoli Storage Productivity Center databases (TPCDBs). If you create multiple connections, you can run reports for multiple Tivoli Storage Productivity Center servers from one Tivoli Common Reporting server. You can also develop and test reports on multiple TPCDB databases in your storage environment.

After you add multiple connections to the TPCDB data source, when you run a report you select the TPCDB database to connect to.

1. On the Tivoli Common Reporting server, add the remote TPCDB database to the system database catalog, and assign an alias to the database. For example, you can catalog the remote TPCDB database as TPCSRV1.
2. In the Tivoli Storage Productivity Center web-based GUI, select **Reporting** > **View predefined reports**.

3. Click **Launch** > **Administration**, and click the **Configuration** tab.
4. Click the **TPCDB** data source, and click the **New Connection** icon .
5. Enter the alias that you created in step 1 on page 808 in the **Name** field, and then click **Next**.
6. From the **Type** list, select **DB2**, and then click **Next**.
7. Enter the alias that you created in step 1 on page 808 in the **DB2 database name** field.
8. Click **Signons**, select **Password**, and then enter a user ID and password for the remote database.
9. Click **Test the connection**, and then click **Test**. Verify that the status of the connection is `Succeeded`.
10. Click **Close**, and then click **Close** again.
11. Click **Finish**.
12. If you scheduled reports, modify the parameters of all scheduled reports to specify the database against which to run the report.

Related information:

IBM Tivoli Storage Productivity Center Information Center
For information about how to install Tivoli Common Reporting on a remote server, see the IBM Tivoli Storage Productivity Center Information Center at publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp. Search for "Installing Tivoli Common Reporting on a remote server".

IBM Cognos Business Intelligence 10.2.0 Information Center
For more information about report schedules, see the IBM Cognos Business Intelligence 10.2.0 Information Center at pic.dhe.ibm.com/infocenter/cbi/v10r2m0/index.jsp. Search for "Schedule Management".

Customizing the logo and title of reports

You can change the logo that is displayed in all Tivoli Storage Productivity Center reports and the title of individual reports. Before you modify predefined reports, create backup copies of them.

Creating a folder in Tivoli Common Reporting

You can create a folder to store backup copies of the predefined reports that you want to modify or for custom reports that you create.

1. In the Tivoli Common Reporting user interface, if you are not on the Tivoli Common Reporting home page, click the **Home** icon .
2. Click the tab where you want to create the folder. For example, click **Public Folders**.
3. Click the **New Folder** icon .
4. Specify a name for the folder, and then click **Finish**. For example, enter `Backups` as the name of the folder.

Creating a backup copy of a report

Before you modify a predefined report, create a backup copy of the report.

1. In Tivoli Common Reporting, create a folder in which to store the backup copy.
2. Navigate to the report of which you want to create a backup copy.
3. Click the check box that is next to the title of the report.

4. Click the **Copy** icon .
5. Navigate to the folder that you created in step 1 on page 809.
6. Click the **Paste** icon .

Related tasks:

"Creating a folder in Tivoli Common Reporting" on page 809
You can create a folder to store backup copies of the predefined reports that you want to modify or for custom reports that you create.

Changing the logo for all reports

You can change the logo that is displayed in the output of all reports. The default logo is the IBM Tivoli logo. You can change the image to your company logo or to another image.

Make sure that your logo is approximately the same size as the IBM Tivoli logo that is in all predefined reports. The image file for the IBM Tivoli logo is at the following location:

JazzSM_installation_directory/IBM/JazzSM/profile/installedApps/JazzSMNode01Cell/IBM Cognos.ear/p2pd.war/tivoli/tcr_common/images/tivoli.gif

1. To store images for your company, create a directory, for example, /myCompany/images, at the following location:

   JazzSM_installation_directory/IBM/JazzSM/profile/installedApps/JazzSMNode01Cell/IBM Cognos.ear/p2pd.war
2. Copy the logo image to the new directory, for example, /myCompany/images/myLogo.jpg.
3. In Tivoli Common Reporting, click the **Public Folders** tab, and then click **Tivoli Storage Productivity Center Report Layouts**.
4. Create a backup copy of the report_page_template file.
5. Click the **Edit with Report Studio** icon for the report_page_template file.
6. Click the IBM Tivoli image in the work area.
7. Double-click the **URL** property in the **Properties** pane.
8. Modify the image URL to refer to your logo image. For example, you can specify the following image URL:

   ../myCompany/images/myLogo.jpg
9. Select **File** > **Save**.

Related tasks:

"Creating a backup copy of a report" on page 809
Before you modify a predefined report, create a backup copy of the report.

Modifying the title of a report

You can change the title that is displayed in the output of a report.

You modify the title in Report Studio. You can also change the title of the report by setting a property for the report in IBM Cognos Connection. However, if you change this property, then other reports cannot link to the report. That is, if you click a link to the report with the changed property in other reports, the report is not displayed.

1. In Tivoli Common Reporting, navigate to the location of the report whose title you want to change.
2. Create a backup copy of the report.

3. Click the **Edit with Report Studio** icon for the report. A message is displayed about the package.
4. Click the **Close** icon on the message.
5. Click the report header in the work area.
6. Double-click the **Overrides** property in the **Properties** pane.
7. Select the **report_name** check box, and click **OK**.
8. Click the **Toolbox** icon in the content pane.
9. Drag a **Text Item** object from the Toolbox and drop the object on the **Drop item to override component child** text in the work area.
10. Type the new title for the report in the Text dialog box, and click **OK**.
11. Select **File** > **Save**.

To change the title back to the default title, double-click the **Overrides** property, and clear the **report_name** check box.

Related tasks:

"Creating a backup copy of a report" on page 809
Before you modify a predefined report, create a backup copy of the report.

Creating a search index of reports, report views, and folders

You can create a search index of the reports, report views, and folders in Tivoli Common Reporting.

Use the search index to search the reports to identify whether the reports contain a property or performance metric.

Tip: If there is no search index, you can search only the titles and descriptions of the reports, not the contents.

1. In Tivoli Storage Productivity Center, select **Reporting** > **View Predefined Reports**.
2. From the **Launch** menu in Tivoli Common Reporting, click **Administration**.
3. Click the **Index Search** tab.
4. Select the following check boxes in the **Indexable Types** area:
   - **Folder**
   - **Report**
   - **Report view**
5. Click **Save**.
6. Click the **Configuration** tab, and then click **Content Administration**.
7. Click the **New Index Update** icon .
8. Type a name for the index update task, then click **Next**.
9. Select the check box beside each of the items that you want to include in the index update task. For example, select the **Public Folders** check box.
10. Click **Next**.
11. To specify the schedule and scope of the index update task, select one of the following options:
    - **Save and run once**: Specify the scope after you click **Finish**.
    - **Save and schedule**: Specify the schedule and the scope after you click **Finish**.

- **Save only**: Specify the scope when you run the task.

  When you specify the scope, you specify whether to index all items, or only items that changed since the last index was updated.

12. Click **Finish**.
13. To run the index update task immediately, click the **Run with options** icon.

Related tasks:

"Searching the predefined reports for a property or performance metric" on page 800
You can search the contents of the predefined reports. You can check which reports contain a property or performance metric. For example, you can list the reports that contain the Storage Pool Name property or the Port Receive I/O Rate performance metric.

Related information:

IBM Cognos Business Intelligence 10.2.0 Information Center
For more information about creating a search index, see the IBM Cognos Business Intelligence 10.2.0 Information Center at pic.dhe.ibm.com/infocenter/cbi/v10r2m0/index.jsp. Search for "Create an Index Update Task".

Configuring security for Tivoli Common Reporting with the default authentication configuration

Use WebSphere Application Server to specify which user IDs can log on to Tivoli Common Reporting. By default, all user IDs have administrative privileges for Tivoli Common Reporting. In Tivoli Common Reporting, you can specify which user IDs and groups of user IDs can use reports, report tools, and report administration functions.

You can configure security in the following ways:

- Create user IDs in WebSphere Application Server.
- Create groups of user IDs in Tivoli Common Reporting.
- Assign administrative privileges to administrator IDs and remove administrative privileges from other user IDs.
- Specify which user IDs or groups can access individual reports or report folders.
- Specify which user IDs or groups can access individual Tivoli Common Reporting functions.

The default authentication configuration for Tivoli Storage Productivity Center uses the built-in file repository of the WebSphere Application Server.

Creating user IDs in WebSphere Application Server

To manage the security of users in Tivoli Common Reporting, you can create user IDs in WebSphere Application Server.

This task is only required if you use the built-in file repository of the WebSphere Application Server to configure security. This task is not required if you use LDAP or OS authentication methods.

To complete this task, you must open the WebSphere Integrated Solutions Console for the instance of WebSphere Application Server that is running Tivoli Common Reporting.

1. To open the WebSphere Integrated Solutions Console, open a web browser, and open one of the following web addresses:

- `http://hostname:port/ibm/console/logon.jsp`
- `https://hostname:port/ibm/console/logon.jsp`

The host name is the server name or IP address for the server that is running WebSphere Application Server. The port is the port number for the same server.

The port number differs depending on which protocol you used (http or https) and the options that you selected when you installed Tivoli Storage Productivity Center. The default port for http 16315 and the default port for https is 16316. For example, you might open the following web address: `http://storage.example.com:16315/ibm/console/logon.jsp`

2. Log on to the WebSphere Integrated Solutions Console with administrator privileges.
3. To create a new user ID, complete the following steps:
   a. Select **Users and Groups** > **Manage Users**.
   b. Click **Create**.
   c. Enter the details of the user ID, and then click **Create**.

   To create more user IDs, repeat this step.
4. To give administrative privileges to users on WebSphere Application Server, complete the following steps:
   a. Select **Users and Groups** > **Administrative User Roles**.
   b. Click **Add**.
   c. Search for the user IDs that you want, select the user IDs, and then add the user IDs to the **Mapped to role** list.
   d. Select **Admin Security Manager** and **Administrator** in the roles list.
   e. Click **OK**, and then click **Save**.

Example: Creating user IDs

Create administrator IDs, normal user IDs, and restricted user IDs.

1. Log on to the WebSphere Integrated Solutions Console.
2. Create two administrator IDs, tcradmin1 and tcradmin2, two normal user IDs, tcrnormaluser1, tcrnormaluser2, and one restricted user, tcrrestricteduser1.
3. Add the roles **Admin Security Manager** and **Administrator** to the two administrator IDs tcradmin1 and tcradmin2.

Create groups of user IDs in Tivoli Common Reporting.

Related tasks:

"Creating groups of user IDs in Tivoli Common Reporting" on page 814
To make it easier to manage security, you can create groups of user IDs in Tivoli Common Reporting.

Related information:

IBM WebSphere Application Server Version 8.5 Information Center
For more information about administrative roles, see the IBM WebSphere Application Server Version 8.5 Information Center at pic.dhe.ibm.com/infocenter/wasinfo/v8r5/index.jsp. Search for "Administrative roles".

Opening Tivoli Common Reporting

You can use Tivoli Common Reporting to configure security for Tivoli Storage Productivity Center reports.

Open a web browser, and open the following web address:
`https://servername:port/tarf/servlet/dispatch`
The host name is the server name or IP address for the server that is running Tivoli Common Reporting. The port is the port number for the same server.
For example, you might open the following web address:
`https://storage.example.com:16311/tarf/servlet/dispatch`

Tip: Alternatively, in the Tivoli Storage Productivity Center web-based GUI, select **Reporting** > **View predefined reports**.

Creating groups of user IDs in Tivoli Common Reporting

To make it easier to manage security, you can create groups of user IDs in Tivoli Common Reporting.

1. Open Tivoli Common Reporting.
2. Log on to Tivoli Common Reporting with administrative privileges.
3. From the **Launch** menu, click **Administration**.
4. Click the **Security** tab.
5. Click the **Cognos** directory.
6. To create a group of user IDs, complete the following steps:
   a. Click the **New Group** icon .
   b. Enter the details of the group, and then click **Next**.
   c. Click **Add** on the "Select the members" page.
   d. Click the **VMMProvider** directory, and after the page refreshes, click **Search**.
   e. Search for the user ID that you want, and select the check box for the user ID.
   f. Add the user ID to the **Selected entries** list, and then click **OK**.
   g. Click **Finish**.

   To create more groups of user IDs, repeat step 6.

Example: Creating groups of user IDs

Organize the different types of IDs into groups. It is easier to manage security when you create groups of user IDs. This example uses the user IDs that you created in WebSphere Application Server.

1. Log on to Tivoli Common Reporting with the tcradmin1 ID.
2. Create a group for administrators, tcradminGroup, and add the user IDs tcradmin1 and tcradmin2 to the group.

   Tip: If you add the user ID of the administrator that was used to install Jazz for Service Management, the default administrator for Jazz for Service Management becomes an administrator of Tivoli Common Reporting also.
3. Create a group for normal users, tcrnormaluserGroup, and add the user IDs tcrnormaluser1 and tcrnormaluser2 to the group.
4. Create a group for administrators, tcrrestricteduserGroup, and add the user ID tcrrestricteduser1 to the group.

Restrict administrative privileges in Tivoli Common Reporting.

Related tasks:

"Opening Tivoli Common Reporting" on page 813
You can use Tivoli Common Reporting to configure security for Tivoli Storage Productivity Center reports.

"Restricting administrative privileges for users in Tivoli Common Reporting"
By default, all user IDs have administrative privileges. To restrict administrative privileges, you must first add the group of user IDs that you want to have administrative privileges to an administrative role in Tivoli Common Reporting. Then, you can remove other user IDs from the administrative role.

"Creating user IDs in WebSphere Application Server" on page 812
To manage the security of users in Tivoli Common Reporting, you can create user IDs in WebSphere Application Server.

Restricting administrative privileges for users in Tivoli Common Reporting

By default, all user IDs have administrative privileges. To restrict administrative privileges, you must first add the group of user IDs that you want to have administrative privileges to an administrative role in Tivoli Common Reporting. Then, you can remove other user IDs from the administrative role.

Tivoli Common Reporting includes predefined security roles. User IDs and groups must be members of security roles to use individual functions. For example, if a user ID is a member of the **System Administrators** role, then that user ID has administrative privileges.

Tivoli Common Reporting also includes predefined groups. The **Everyone** group includes all user IDs. By default, the **System Administrators** role includes the **Everyone** group. Therefore, by default all user IDs have administrative privileges.

To restrict who has administrative privileges, you must first add the group of user IDs that you want to have administrative privileges to the **System Administrators** role. Then, you must remove the **Everyone** group from the **System Administrators** role.

1. Open Tivoli Common Reporting.
2. Log on to Tivoli Common Reporting with administrative privileges.
3. From the **Launch** menu, click **Administration**.
4. Click the **Security** tab.
5. Click the **Properties** icon for the **System Administrators** role. If the **System Administrators** role is not visible, scroll down.
6. Click the **Members** tab.
7. Click **Add** on the "Set properties" page.
8. Click the **Cognos** directory.
9. Select the check box for the group of user IDs that you want to have administrative privileges, and then add the group to the **Selected entries** list.
10. Click **OK**.
11. Select the check box for the **Everyone** group.
12. Click **Remove** and then click **OK**.

Example: Restricting administrative privileges to the users in a group of user IDs

Assign administrative privileges to one group of user IDs, and remove administrative privileges from other user IDs to control which users can complete administrative tasks.

This example uses the following user IDs and group IDs:

tcradmin1
This user ID is a member of tcradminGroup. This group has administrative privileges for WebSphere Application Server and Tivoli Common Reporting.

tcrnormaluser1
This user ID is a member of tcrnormalGroup. This group does not have administrative privileges for WebSphere Application Server or Tivoli Common Reporting.

1. Log on to Tivoli Common Reporting with the tcradmin1 ID.
2. Add the tcradminGroup group to the **System Administrators** role.
3. Remove the **Everyone** group from the **System Administrators** role.
4. Click **Logoff**.
5. To verify that not all user IDs have administrative privileges, complete the following steps:
   a. Log on to Tivoli Common Reporting with the tcrnormaluser1 ID.
   b. Click the **Launch** menu in Tivoli Common Reporting.
      If the **Administration** item is not displayed on the **Launch** menu, not all users have administrative privileges.

Related information:

IBM Cognos Business Intelligence 10.2.0 Information Center
For more information about the Everyone group, see the IBM Cognos Business Intelligence 10.2.0 Information Center at pic.dhe.ibm.com/infocenter/cbi/v10r2m0/index.jsp. Search for “Built-in Entries”.

Specifying access to reports in Tivoli Common Reporting

You can grant or deny access to reports or to a report folder in Tivoli Common Reporting. For example, you can prevent a user or group of users from viewing hypervisor reports.

1. Open Tivoli Common Reporting.
2. Log on to Tivoli Common Reporting with administrator privileges.
3. Click **IBM Tivoli Productivity Center Predefined Reports**.
4. Click the **Properties** icon for the report or reports folder to which you want to restrict access.
5. Click the **Permissions** tab.
6. Select the **Override the access permissions acquired from the parent entry** check box.
7. After the page refreshes, click **Add**.
8. Click the **Cognos** directory and after the page refreshes, click **Search**.
9. Search for the user ID or group that you want, and select the check box for the user ID or group.
10. Add the user ID or group to the **Selected entries** list, and then click **OK**.

11. On the **Permissions** tab, select the check box for the user ID or group.
12. Select the **Deny** check box or the **Grant** check box for each of the permissions that you require.
13. Click **OK**.

Example: Preventing a group of users from viewing hypervisor reports

This example shows how to prevent the users in the tcrnormaluserGroup group from accessing the reports in the **Hypervisors** folder.

This example uses the following user IDs and group IDs:

tcradmin1
: This user ID is a member of tcradminGroup. This group has administrative privileges for WebSphere Application Server and Tivoli Common Reporting.

tcrnormaluser1
: This user ID is a member of tcrnormalGroup. This group does not have administrative privileges for WebSphere Application Server or Tivoli Common Reporting.

1. Log on to Tivoli Common Reporting with the tcradmin1 ID.
2. Click **IBM Tivoli Productivity Center Predefined Reports**.
3. Click the **Properties** icon for the **Hypervisors** folder.
4. Click the **Permissions** tab.
5. Select the **Override the access permissions acquired from the parent entry** check box.
6. After the page refreshes, click **Add**.
7. Search the **Cognos** directory for the tcrnormaluserGroup group, and then add tcrnormaluserGroup to the **Selected entries** list.
8. On the **Permissions** tab, select the check box for tcrnormaluserGroup.
9. Select the **Deny** check box for all of the permissions for tcrnormaluserGroup, and then click **OK**.
10. To verify that user IDs in tcrnormaluserGroup do not have access to the hypervisor reports, complete the following steps:
    a. Log on to Tivoli Common Reporting with the user ID and password of a member of tcrnormaluserGroup. For example, log on with the tcrnormaluser1 ID.
    b. Click **IBM Tivoli Productivity Center Predefined Reports**.

       If the **Hypervisors** folder is not displayed in the reports list, the tcrnormaluser1 ID does not have access to the hypervisor reports.

Controlling access to Report Studio

You can control access to functions in Tivoli Common Reporting by using predefined security roles. You assign groups of user IDs to security roles. For example, you can control which users can create or edit reports in Report Studio.

Tivoli Common Reporting includes predefined security roles. User IDs and groups must be members of security roles to use individual functions. For example, if a user ID is a member of the **Authors** role, the user can use the Report Studio tool to create or edit reports.

Tivoli Common Reporting also includes predefined groups. The **Everyone** group includes all authenticated users. By default, the **Authors** role includes the **Everyone** group. Therefore, by default all authenticated users can use Report Studio.

To control access to Report Studio and other Tivoli Common Reporting functions, use the predefined security roles with the groups of user IDs that you create.

For example, to control who can use Report Studio, you take the following general steps:

1. Add the groups of user IDs that you want to use the tools to the **Authors** and **Express Authors** roles.
2. Remove all other user IDs from the **Authors** and **Express Authors** roles. That is, remove the **Everyone** group from the roles.
3. Grant permissions to the **HTML Items in Report** capability to the **Everyone** group.

This procedure describes the general steps in more detail.

1. Open Tivoli Common Reporting.
2. Log on to Tivoli Common Reporting with administrator privileges.
3. From the **Launch** menu, click **Administration**.
4. Click the **Security** tab.
5. To assign a group of users to the **Authors** role and remove the **Everyone** group from the role, complete the following steps:
   a. Click the **Cognos** directory, and then click the **Properties** icon for the **Authors** role.
   b. Click the **Members** tab, and then click **Add**.
   c. Click the **Cognos** directory.
   d. Select the check box for the group, and then add it to the **Selected entries** list.
   e. Click **OK**.
   f. Select the check box for the **Everyone** user.
   g. Click **Remove**, and then click **OK**.

   Repeat this step for the **Express Authors** role.
6. To grant permissions to the **HTML Items in Report** function, complete the following steps:
   a. On the **Security** tab, click **Capabilities**.
   b. Click **Report Studio** in the **Capability** list.
   c. Click the **Actions** icon beside the **HTML Items in Report** list, and then select **Set properties**.
   d. Click the **Permissions** tab.
   e. Select the **Override the access permissions acquired from the parent entry** check box.
   f. Click **Add**.
   g. Click the **Cognos** directory.
   h. Select the check box for the **Everyone** group, and then add it to the **Selected entries** list.
   i. Click **OK**.

j. Select the check box for the **Everyone** group, and then elect the **Grant** check box for the **Execute** and **Traverse** permissions.

k. Click **OK**.

Example: Allowing and restricting access to Report Studio

This example shows how to allow access to Report Studio to groups of administrators and other users. The example also shows how to restrict access by some users to these tools.

This example uses the following user IDs and group IDs:

This example uses the following user IDs and group IDs:

tcradmin1
: This user ID is a member of tcradminGroup. This group has administrative privileges for WebSphere Application Server and Tivoli Common Reporting.

tcrnormaluser1
: This user ID is a member of tcrnormalGroup. This group does not have administrative privileges for WebSphere Application Server or Tivoli Common Reporting.

tcrrestricteduser1
: This user ID is a member of tcrrestrictedGroup. This group does not have administrative privileges for WebSphere Application Server or Tivoli Common Reporting.

1. Assign the tcrnormalGroup group to the **Authors** role, and remove the **Everyone** group from the role.
2. Assign the tcrnormalGroup group to the **Express Authors** role, and remove the **Everyone** group from the role.
3. Assign the tcradminGroup group to the **Authors** and **Express Authors** roles.
4. For the **HTML Items in Report** capability, grant **Execute** and **Traverse** permissions to the **Everyone** group.
5. To verify that users who are in tcrrestrictedGroup cannot use Report Studio, complete the following steps:
   a. Log on to Tivoli Common Reporting with the user ID and password of a member of tcrrestrictedGroup. For example, log on with the tcrrestricteduser1 ID.
   b. Click the **Launch** menu in Tivoli Common Reporting.

      If the Report Studio item is not on the **Launch** menu, users who are in tcrrestrictedGroup cannot use the tool.

Related tasks:

"Opening Tivoli Common Reporting" on page 813
You can use Tivoli Common Reporting to configure security for Tivoli Storage Productivity Center reports.

Related information:

IBM Cognos Business Intelligence 10.2.0 Information Center
For more information about predefined groups and roles like Authors, see the IBM Cognos Business Intelligence 10.2.0 Information Center at pic.dhe.ibm.com/infocenter/cbi/v10r2m0/index.jsp. Search for "Predefined Entries".

Storage virtualizer reports that are available when you upgrade

If you upgrade from Tivoli Storage Productivity Center V5.1, V5.1.1, or V5.1.1.1 to V5.2, you might experience problems with reports.

Review the following table of solutions to common problems with reports:

Problem	Solution
In Tivoli Storage Productivity Center V5.2, no predefined reports are designed specifically for storage virtualizers.	Run predefined reports about storage systems, and select **Storage Virtualizer** in the **Storage System Configuration** list.
Tivoli Storage Productivity Center V5.2 contains a reports package, Performance. This package replaces the IBM Tivoli Storage Productivity Center Performance package. When you upgrade from Tivoli Storage Productivity Center V5.1, V5.1.1, or V5.1.1.1 to V5.2, you can continue to use the IBM Tivoli Storage Productivity Center Performance package in V5.2. Reports that are designed specifically for storage virtualizers are only available in this package. However, the following reports from V5.1, V5.1.1, and V5.1.1.1 do not function correctly in V5.2: • Storage Systems - Most Active Ports Report • Storage System - Port Details Report • Storage Systems - Most Active Arrays Report • Storage System - Array Details Report	View one of the following predefined reports about storage systems in V5.2: • Most Active Ports • Performance of One Port • Most Active RAID Arrays • Performance of One RAID Array
If you used the IBM Tivoli Storage Productivity Center Performance package in V5.1, V5.1.1, or V5.1.1.1 to create custom reports, you might need to re-create those reports.	If your custom reports use the following data, you must re-create the reports in the Performance package: • Performance metrics for ports on DS8000 and DS6000 systems • Performance metrics for RAID arrays
If you used the IBM Tivoli Storage Productivity Center Resources package in V5.1, V5.1.1, or V5.1.1.1 to create custom reports, you might need to re-create those reports.	If your custom reports use data about VMware file systems, you must re-create the reports in the Capacity and Relationships package.
In Tivoli Storage Productivity Center V5.2, the predefined reports about storage tiering are not in the same folder as V5.1, V5.1.1, or V5.1.1.1.	In Tivoli Storage Productivity Center V5.2, go to the **Public Folders** > **IBM Tivoli Productivity Center Packages** > **Storage Tiering** folder.

Storage virtualizer resource reports

Storage virtualizer resource reports show the space statistics for storage virtualizers. The reports also show the space statistics for local disks, managed disks, pools, and volumes on storage virtualizers.

Running the Storage Virtualizers - Space Report

To see information about the capacity of pools, managed disks (MDisks), and volumes on storage virtualizers, run the Storage Virtualizers - Space Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers in the list are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Resources**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizers - Space Report**.
4. Optional: Select one or more storage virtualizers.
5. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizers - Space Report"
The Storage Virtualizers - Space Report shows the space statistics for pools, managed disks (MDisks), disks, volumes, and virtual volumes on storage virtualizers.

Storage Virtualizers - Space Report

The Storage Virtualizers - Space Report shows the space statistics for pools, managed disks (MDisks), disks, volumes, and virtual volumes on storage virtualizers.

Charts

The bar chart shows the used space and available space for pools on storage virtualizers.

Report output

For each storage virtualizer, the following information is provided:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Location
The physical location of a storage virtualizer. The location is defined when a storage virtualizer is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage virtualizer in the **Properties** pane for the storage system in the web-based GUI.

Storage Virtualizer Total Pool Capacity (GiB)
The total amount of storage space in pools on the storage virtualizer.

Storage Virtualizer Pool Used Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can pre-allocate thin-provisioned volume space when the volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. For pools with compressed volumes on SAN Volume Controller and Storwize V7000, the Storage Virtualizer Pool Used Space reflects the size of compressed data written to disk. As the data changes the Storage Virtualizer Pool Used Space may at times be less than the Storage Virtualizer Pool Allocated Space.

For pools with volumes that are not thin provisioned or compressed in SAN Volume Controller, Storwize V7000, and other storage systems, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a pool. Because data collection is run on a set schedule and the used space on volumes can change rapidly, the value in this column might not be 100% accurate for the current state of volumes.

Storage Virtualizer Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Storage Virtualizer Physical Allocation Percentage
The percentage of physical space in storage virtualizer pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

```
(allocated space ÷ pool capacity) × 100
```

For example, the physical allocation percentage is 25% for a total storage virtualizer pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Virtualizer Virtual Allocation Percentage
The percentage of physical space in storage virtualizer pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

```
(total volume capacity ÷ pool capacity) × 100
```

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB.

This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases,

use the Storage Virtualizer Shortfall Percentage property to estimate how critical the shortage of space is for storage virtualizer pools.

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Shortfall Percentage
The percentage of the remaining unallocated volume space in storage virtualizer pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(unallocatable space ÷ (volume space - used volume space)) × 100`

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space.

This Storage Virtualizer Shortfall Percentage represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB ÷ 100 GiB × 100).

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

This property is available only for storage virtualizer pools with thin-provisioned volumes.

Storage Virtualizer Volume Allocated Space (GiB)
The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

Storage Virtualizer MDisk Capacity (GiB)
The amount of storage space on the managed disk.

Storage Virtualizer MDisk Available Space (GiB)
The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

Storage Virtualizer Total Volume Capacity (GiB)
The amount of virtual volume space on all volumes on the storage virtualizer.

Storage Virtualizer Volume Assigned Space (GiB)
The amount of virtual volume space that is assigned to the server or storage virtualizer.

Storage Virtualizer Disk Capacity (GiB)
The amount of local disk storage on the storage virtualizer. The local disks are the internal disk drives.

Storage Virtualizer Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Virtualizer Number of Volumes
The number of volumes on a resource.

Storage Virtualizer Number of MDisks
The number of managed disks on a resource.

Report links

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click the name of the report. You can open the following related reports:

- The Storage Virtualizer - Pools Space Report
- The Storage Virtualizer - Volumes Space Report
- The Managed Disks Capacity
- The Storage Virtualizer - Local Disks Space Report
- The Storage Virtualizers - Historical Space Report

Related tasks:

"Running the Storage Virtualizers - Space Report" on page 821
To see information about the capacity of pools, managed disks (MDisks), and volumes on storage virtualizers, run the Storage Virtualizers - Space Report.

Running the Managed Disks Capacity report

To see information about used space and available space on managed disks (MDisks) on storage virtualizers, run the Managed Disks Capacity report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Managed Disks**.
4. Click **Managed Disks Capacity**.
5. Optional: Select the configuration of the storage systems.
6. Select one or more storage systems.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Managed Disks Capacity report"
Shows a chart of the 20 managed disks with the most available space, and a table with space statistics for all managed disks. Use the report to review space availability in your managed disks.

Managed Disks Capacity report

Shows a chart of the 20 managed disks with the most available space, and a table with space statistics for all managed disks. Use the report to review space availability in your managed disks.

Restriction: This report applies only to SAN Volume Controller, Storwize, and Hitachi Universal Storage Platform V systems.

Charts

The bar chart shows the used space and the available space for managed disk on storage systems.

Report output

For each managed disk on a storage system, the following information is provided:

Storage System MDisk Name
The name that was assigned to the managed disk on a storage system when it was added to the system.

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage System MDisk Capacity (GiB)
The amount of storage space on the managed disk.

Storage System MDisk Used Space (GiB)
The amount of used space on the managed disk.

Storage System MDisk Available Space (GiB)
The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage MDisk RAID Level
The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

Storage System MDisk Type
The type of managed disk on a storage system. For example, the disk on a storage system can be a local managed disk.

Storage System MDisk Mode
The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

Array Extents are provided from local disks.

Managed
Extents are provided from a back-end storage volume.

Unmanaged
The managed disk is not used in the system.

Storage System MDisk Strip Size (KB)
The RAID strip size on a managed disk on a storage system.

Storage System MDisk Is Balanced
Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is `Yes`, the LUNs are balanced.

Storage System MDisk Fast Write State
Shows whether the cache for a volume on a disk that is managed by a storage system is empty, contains data, or is corrupted.

Storage System MDisk Write Verify
Shows whether all write operations on a managed disk on a storage system are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is `Yes`, all write operations are verified by a follow-up read operation.

Related tasks:

"Running the Managed Disks Capacity report" on page 824
To see information about used space and available space on managed disks (MDisks) on storage virtualizers, run the Managed Disks Capacity report.

Running the Storage Virtualizer - Pools Space Report

To see information about the used space and available space in pools on storage virtualizers, run the Storage Virtualizer - Pools Space Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers in the list are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Resources**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizer - Pools Space Report**.
4. Optional: Select one or more storage virtualizers.
5. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizer - Pools Space Report" on page 827
The Storage Virtualizer - Pools Space Report shows the space statistics for pools

on storage virtualizers.

Storage Virtualizer - Pools Space Report

The Storage Virtualizer - Pools Space Report shows the space statistics for pools on storage virtualizers.

Charts

The bar chart shows the used space and available space for pools on storage virtualizers.

Report output

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Pool Capacity (GiB)
The total amount of storage space in a pool.

Storage Virtualizer Pool Used Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can pre-allocate thin-provisioned volume space when the volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. For pools with compressed volumes on SAN Volume Controller and Storwize V7000, the Storage Virtualizer Pool Used Space reflects the size of compressed data written to disk. As the data changes the Storage Virtualizer Pool Used Space may at times be less than the Storage Virtualizer Pool Allocated Space.

For pools with volumes that are not thin provisioned or compressed in SAN Volume Controller, Storwize V7000, and other storage systems, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a pool. Because data collection is run on a set schedule and the used space on volumes can change rapidly, the value in this column might not be 100% accurate for the current state of volumes.

Storage Virtualizer Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Storage Virtualizer Pool Allocated Space (GiB)
The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes.

The space that is allocated for thin-provisioned volumes is less than their virtual capacity. The Storage Virtualizer Pool Total Volume Capacity (GiB) property shows the virtual capacity of the volumes. If a pool does not contain thin-provisioned volumes, this value is the same as the value of the Storage Virtualizer Pool Total Volume Capacity (GiB) property.

This value is equal to Storage Virtualizer Used Space (GiB) for volumes that are not thin provisioned.

Storage Virtualizer Pool Used Volume Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Virtualizer Pool Total Volume Capacity (GiB)
The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Virtualizer Pool Extent Size (MiB)
The size of the extents that were specified when a pool was created. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Number of Storage Volumes
The number of volumes in the storage virtualizer pool.

Storage Virtualizer Pool Number of MDisks
The number of managed disks on a resource.

Storage Virtualizer Pool Assigned Volume Space (GiB)
The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Virtualizer Pool Unassigned Volume Space (GiB)
The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Virtualizer Pool Physical Allocation Percentage
The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$(\text{allocated space} \div \text{pool capacity}) \times 100$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Virtualizer Pool Virtual Allocation Percentage
The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(total volume capacity ÷ pool capacity) × 100`

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. In this case, the virtual capacity that is committed to the volumes in the pool is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, use the Storage Virtualizer Pool Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

This value is only available for pools with thin-provisioned volumes.

Storage Virtualizer Pool Shortfall Percentage
The percentage of the remaining unallocated volume space in a pool that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(unallocatable space ÷ (total volume capacity − allocated space)) × 100`

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool RAID Level
The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Virtualizer Pool Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Virtualizer Pool Is Solid State
Shows whether there are solid-state drives in the pool. This property can contain the following values:

Mixed The pool contains both hard disk drives and solid-state drives.

Non solid state
The pool contains no solid-state drives.

Solid state
The pool contains a least one solid-state drive.

Storage Virtualizer Pool Warning Level
The percentage of used capacity of the resource at which a warning is generated. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Report links

To open related reports, click the name of the storage pool and then click the name of the report. You can open the following related reports:

- The Storage Virtualizer - Volumes Space Report
- The Managed Disks Capacity
- The Storage Virtualizer - Pools Historical Space Report

Related tasks:

"Running the Storage Virtualizer - Pools Space Report" on page 826
To see information about the used space and available space in pools on storage virtualizers, run the Storage Virtualizer - Pools Space Report.

Running the Storage Virtualizer - Volumes Space Report

To see information about the capacity of volumes on storage virtualizers, run the Storage Virtualizer - Volumes Space Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers in the list are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Resources**.
2. Click `Storage Virtualizer`.
3. Click the **Storage Virtualizer - Volumes Space Report** link.
4. Optional: Select one or more storage virtualizers.
5. Optional: Select the type of volume.

Option	Description
All	Select **All**, the default, to view information about all of the volumes on the selected storage virtualizers.
Primary	Select **Primary** to view information about primary volumes on the selected storage virtualizers.

Option	Description
Secondary	Select **Secondary** to view information about secondary volumes on the selected storage virtualizers.

6. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizer - Volumes Space Report"
The Storage Virtualizer - Volumes Space Report shows the space statistics for volumes on storage virtualizers.

Storage Virtualizer - Volumes Space Report

The Storage Virtualizer - Volumes Space Report shows the space statistics for volumes on storage virtualizers.

Charts

The bar chart shows the used space for volumes on storage virtualizers.

Report output

For each volume on a storage virtualizer, the following information is provided:

Storage Virtualizer Volume Name
The name that was assigned to the storage virtualizer volume when it was added to the system.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Volume Copy ID
The identifier of a volume copy. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Volume Copy Is Primary
Shows whether the volume copy is the primary volume copy. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Volume Used Space (GiB)
The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Virtualizer Volume Used Space (GiB) might be different than the Storage Virtualizer Volume Allocated Space (GiB). For volumes

that are not thin provisioned, the values for Storage Virtualizer Volume Used Space (GiB) and Storage Virtualizer Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Virtualizer Volume Capacity (GiB)
The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

Storage Virtualizer Volume Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Virtualizer Volume Allocated Space (GiB)
The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

Storage Virtualizer Volume Physical Allocation Percentage
The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(allocated space ÷ capacity) × 100`

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore the space that is reserved for volumes is 50 GiB.

Storage Virtualizer Volume Shortfall Percentage
The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume. The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(unallocatable volume space ÷ unallocated volume space) × 100`

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level for the volume. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

`(100 GiB unallocated volume space - 70 GiB remaining pool space) ÷ 100 GiB unallocated volume space × 100`

This value is only available for pools with thin-provisioned volumes.

Storage Virtualizer Volume Used Allocated Space Percentage
The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(used space ÷ allocated space) × 100`

This property is available only for SAN Volume Controller and Storwize V7000 volumes.

Storage Virtualizer Volume Is Assigned
Shows whether the volume is assigned to a server. If this value is `Yes`, the volume is assigned to a server.

Storage Virtualizer Volume Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Virtualizer Volume Warning Level
The percentage of used capacity of the volume at which a warning is generated. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Volume Type
The type of storage virtualizer volume. For example, the storage virtualizer volume can be striped, sequential, or another type of storage virtualizer volume. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Volume Grain Size (KiB)
The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Volume RAID Level
The RAID level of a storage virtualizer volume, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Related links

To open the related report, click the name of the volume, click **Go To**, and then click **Storage Virtualizer - Volumes Historical Space Report**.

Related tasks:

"Running the Storage Virtualizer - Volumes Space Report" on page 830
To see information about the capacity of volumes on storage virtualizers, run the Storage Virtualizer - Volumes Space Report.

Storage virtualizer performance reports

Storage virtualizer performance reports show the most active storage virtualizers, and the most active storage resources on storage virtualizers.

Running the Storage Virtualizers - Most Active IO Groups Report

To see which I/O groups are most active on your storage virtualizers, run the Storage Virtualizers - Most Active IO Groups Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizers - Most Active IO Groups Report**.
4. Optional: Select a storage virtualizer, and then select a sort order.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizers - Most Active IO Groups Report"
The Storage Virtualizers - Most Active IO Groups Report shows the performance statistics for the most active I/O groups on storage virtualizers. Use the report to monitor the performance of I/O groups over a specified period.

Storage Virtualizers - Most Active IO Groups Report

The Storage Virtualizers - Most Active IO Groups Report shows the performance statistics for the most active I/O groups on storage virtualizers. Use the report to monitor the performance of I/O groups over a specified period.

Charts

The chart shows the most active I/O groups on storage virtualizers. By default, I/O groups are sorted by the total overall I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided for I/O groups is as follows:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer I/O Group Name
The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Normal I/O Rate (ops/s)
The average number of nonsequential I/O operations per second. This value includes both read and write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Related reports

To open related reports, click the name of the resource or storage device, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizer - Details Report
- Storage Virtualizer - IO Group Details Report

Related tasks:

"Running the Storage Virtualizers - Most Active IO Groups Report" on page 834
To see which I/O groups are most active on your storage virtualizers, run the Storage Virtualizers - Most Active IO Groups Report.

Running the Storage Virtualizers - Most Active MDisks Report

To see which managed disks (MDisks) are most active on your storage virtualizers, run the Storage Virtualizers - Most Active MDisks Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.

2. Click Storage Virtualizer.
3. Click **Storage Virtualizers - Most Active MDisks Report**.
4. Optional: Select a storage virtualizer, and then select a sort order.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizers - Most Active MDisks Report"
The Storage Virtualizers - Most Active MDisks Report shows the performance statistics for the most active managed disks (MDisks) on storage virtualizers. Use the report to monitor the performance of MDisks over a specified period.

Storage Virtualizers - Most Active MDisks Report

The Storage Virtualizers - Most Active MDisks Report shows the performance statistics for the most active managed disks (MDisks) on storage virtualizers. Use the report to monitor the performance of MDisks over a specified period.

Charts

The chart shows the most active MDisks on storage virtualizers. By default, MDisks are sorted by the total back-end I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided for MDisks is as follows:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Virtualizer MDisk Name
The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Related reports

To open related reports, click the name of the resource or device, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizer - Details Report
- Storage Virtualizer - Pool Details Report
- Storage Virtualizer - MDisk Details Report
- Storage Resource Relationships Summary

Related tasks:

"Running the Storage Virtualizers - Most Active MDisks Report" on page 835
To see which managed disks (MDisks) are most active on your storage virtualizers, run the Storage Virtualizers - Most Active MDisks Report.

Running the Storage Virtualizers - Most Active Nodes Report

To see which nodes are most active on your storage virtualizers, run the Storage Virtualizers - Most Active Nodes Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizers - Most Active Nodes Report**.
4. Optional: Select a storage virtualizer, and then select a sort order.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizers - Most Active Nodes Report"
The Storage Virtualizers - Most Active Nodes Report shows the performance statistics for the most active nodes on storage virtualizers. Use the report to monitor the performance of nodes over a specified period.

Storage Virtualizers - Most Active Nodes Report

The Storage Virtualizers - Most Active Nodes Report shows the performance statistics for the most active nodes on storage virtualizers. Use the report to monitor the performance of nodes over a specified period.

Charts

The chart shows the most active nodes on storage virtualizers. By default, nodes are sorted by the total overall I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided for nodes is as follows:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer I/O Group Name
The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

Storage Virtualizer Node Name
The name that was assigned to the storage virtualizer node when it was added to the system.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Related reports

To open related reports, click the name of the resource or device, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizer - Details Report
- Storage Virtualizer - IO Group Details Report
- Storage Virtualizer - Node Details Report

Related tasks:

"Running the Storage Virtualizers - Most Active Nodes Report" on page 837
To see which nodes are most active on your storage virtualizers, run the Storage Virtualizers - Most Active Nodes Report.

Running the Storage Virtualizers - Most Active Pools Report

To see which pools are most active on your storage virtualizers, run the Storage Virtualizers - Most Active Pools Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizers - Most Active Pools Report**.
4. Optional: Select a storage virtualizer, and then select a sort order.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizers - Most Active Pools Report"
The Storage Virtualizers - Most Active Pools Report shows the performance statistics for the most active pools on storage virtualizers. Use the report to monitor the performance of pools over a specified period.

Storage Virtualizers - Most Active Pools Report

The Storage Virtualizers - Most Active Pools Report shows the performance statistics for the most active pools on storage virtualizers. Use the report to monitor the performance of pools over a specified period.

Charts

The chart shows the most active pools on storage virtualizers. By default, pools are sorted by the total overall I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided for pools is as follows:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Related reports

To open related reports, click the name of the resource or device, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizer - Details Report
- Storage Virtualizer - Pool Details Report
- Storage Virtualizers - Most Active MDisks Report
- Storage Virtualizers - Most Active Volumes Report
- Storage Resource Relationships Summary

Related tasks:

"Running the Storage Virtualizers - Most Active Pools Report" on page 839
To see which pools are most active on your storage virtualizers, run the Storage Virtualizers - Most Active Pools Report.

Running the Storage Virtualizers - Most Active Volumes Report

To see which volumes are most active on storage virtualizers, run the Storage Virtualizers - Most Active Volumes Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizers - Most Active Volumes Report**.
4. Optional: Select a storage virtualizer, and then select a sort order.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.

7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizers - Most Active Volumes Report"
The Storage Virtualizers - Most Active Volumes Report shows the performance statistics for the most active volumes on storage virtualizers. Use the report to monitor the performance of volumes over a specified period.

Storage Virtualizers - Most Active Volumes Report

The Storage Virtualizers - Most Active Volumes Report shows the performance statistics for the most active volumes on storage virtualizers. Use the report to monitor the performance of volumes over a specified period.

Charts

The chart shows the most active volumes on storage virtualizers. By default, volumes are sorted by the total overall I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided for volumes is as follows:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Virtualizer I/O Group Name
The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

Storage Virtualizer Volume Name
The name that was assigned to the storage virtualizer volume when it was added to the system.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Volume Utilization
The average percentage of time that the volume is busy.

Related reports

To open related reports, click the name of the resource or device, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizer - Details Report
- Storage Virtualizer - Pool Details Report
- Storage Virtualizer - IO Group Details Report
- Storage Virtualizer - Volume Details Report
- Storage Resource Relationships Summary

Related tasks:

"Running the Storage Virtualizers - Most Active Volumes Report" on page 840
To see which volumes are most active on storage virtualizers, run the Storage Virtualizers - Most Active Volumes Report.

Running the Storage Virtualizers - Most Active Ports Report

To see which ports are most active on your storage virtualizers, run the Storage Virtualizers - Most Active Ports Report.

To run the report with the default settings, click **Finish**.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizers - Most Active Ports Report**.
4. Optional: Select a storage virtualizer, and then select a sort order.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizers - Most Active Ports Report"
The Storage Virtualizers - Most Active Ports Report shows the performance statistics for the most active ports on storage virtualizers. Use the report to monitor the performance of ports over a specified period.

Storage Virtualizers - Most Active Ports Report

The Storage Virtualizers - Most Active Ports Report shows the performance statistics for the most active ports on storage virtualizers. Use the report to monitor the performance of ports over a specified period.

Charts

The chart shows the most active ports on storage virtualizers. By default, ports are sorted by the total port I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided for ports is as follows:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Port Name
An alphanumeric string that identifies a storage virtualizer port.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Related reports

To open related reports, click the name of the resource or device, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizer - Details Report
- Storage Virtualizer - Port Details Report

Related tasks:

"Running the Storage Virtualizers - Most Active Ports Report" on page 842
To see which ports are most active on your storage virtualizers, run the Storage Virtualizers - Most Active Ports Report.

Storage virtualizer performance details reports

Storage virtualizer reports about performance details show information about the performance of storage virtualizers, and information about the performance of the resources related to storage virtualizers.

Running the Storage Virtualizer - Details Report

To see information about the performance of storage virtualizers, run the Storage Virtualizer - Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizer Performance Details Report**.
4. Select a storage virtualizer.
5. Optional: Select an interval.

6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizer - Details Report"
The Storage Virtualizer - Details Report shows the performance statistics for a storage virtualizer. Use the report to monitor the performance of a storage virtualizer over a specified period.

Storage Virtualizer - Details Report

The Storage Virtualizer - Details Report shows the performance statistics for a storage virtualizer. Use the report to monitor the performance of a storage virtualizer over a specified period.

Charts

The charts show the following performance statistics for the storage virtualizer:

- Total overall rates for I/O operations, and the total overall response times
- The overall I/O rates for read and write operations, and the response times for read and write operations
- The data rates for read and write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for processors

Report output

The information that is provided for a storage virtualizer is as follows:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Day or Hour or Time
The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Related reports

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click the name of the report. You can open the following related reports:

- Storage Virtualizers - Most Active IO Groups Report
- Storage Virtualizers - Most Active MDisks Report
- Storage Virtualizers - Most Active Nodes Report
- Storage Virtualizers - Most Active Pools Report
- Storage Virtualizers - Most Active Ports Report
- Storage Virtualizers - Most Active Volumes Report

Related tasks:

"Running the Storage Virtualizer - Details Report" on page 843
To see information about the performance of storage virtualizers, run the Storage Virtualizer - Details Report.

Running the Storage Virtualizer - IO Group Details Report

To see information about the performance of an I/O group on a storage virtualizer, run the Storage Virtualizer - IO Group Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizer - IO Group Details Report**.
4. Select a storage virtualizer, and then select an I/O group.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizer - IO Group Details Report"
The Storage Virtualizer - IO Group Details Report shows the performance statistics for an I/O group on a storage virtualizer. Use the report to monitor the performance of an I/O group over a specified period.

Storage Virtualizer - IO Group Details Report

The Storage Virtualizer - IO Group Details Report shows the performance statistics for an I/O group on a storage virtualizer. Use the report to monitor the performance of an I/O group over a specified period.

Charts

The charts show the following performance statistics for an I/O group on a storage virtualizer:

- The total overall I/O rates, and the total response times
- The overall I/O rates for read and write operations, and the response times for read and write operations
- The data rates for read and write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for processors

Report output

The information that is provided for an I/O group is as follows:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer I/O Group Name
The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

Day or Hour or Time
The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Related reports

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click **Storage Virtualizers - Most Active IO Groups Report**.

Related tasks:

"Running the Storage Virtualizer - IO Group Details Report" on page 846
To see information about the performance of an I/O group on a storage virtualizer, run the Storage Virtualizer - IO Group Details Report.

Running the Storage Virtualizer - MDisk Details Report

To see information about the performance of a managed disk (MDisk) on a storage virtualizer, run the Storage Virtualizer - MDisk Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizer - MDisk Details Report**.
4. Select a storage virtualizer, and then select an MDisk.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizer - MDisk Details Report"
The Storage Virtualizer - MDisk Details Report shows the performance statistics for a managed disk (MDisk) on a storage virtualizer. Use the report to monitor the performance of an MDisk over a specified period.

Storage Virtualizer - MDisk Details Report

The Storage Virtualizer - MDisk Details Report shows the performance statistics for a managed disk (MDisk) on a storage virtualizer. Use the report to monitor the performance of an MDisk over a specified period.

Charts

The charts show the following performance statistics for an MDisk on a storage virtualizer:

- The total I/O rates for back-end operations, and the total response times for back-end operations
- The I/O rates for read and write back-end operations, and the response times for read and write back-end operations
- The total data rates for back-end operations, and the data rates for read and write back-end operations
- The overall queue times for back-end operations, and the queue times for read and write back-end operations

Report output

The information that is provided for an MDisk is as follows:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Virtualizer MDisk Name
The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

Day or Hour or Time
The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Related reports

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click **Storage Virtualizers - Most Active MDisks Report**.

Related tasks:

"Running the Storage Virtualizer - MDisk Details Report" on page 848
To see information about the performance of a managed disk (MDisk) on a storage virtualizer, run the Storage Virtualizer - MDisk Details Report.

Running the Storage Virtualizer - Node Details Report

To see information about the performance of a node on a storage virtualizer, run the Storage Virtualizer - Node Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizer - Node Details Report**.
4. Select a storage virtualizer, and then select a node.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizer - Node Details Report"
The Storage Virtualizer - Node Details Report shows the performance statistics for a node on a storage virtualizer. Use the report to monitor the performance of a node over a specified period.

Storage Virtualizer - Node Details Report

The Storage Virtualizer - Node Details Report shows the performance statistics for a node on a storage virtualizer. Use the report to monitor the performance of a node over a specified period.

Charts

The charts show the following performance statistics for a node on a storage virtualizer:

- The total overall I/O rates, and the total response times
- The overall I/O rates for read and write operations, the response times for read and write operations
- The data rates for read and write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for processors

Report output

The information that is provided for a node is as follows:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer I/O Group Name
The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

Storage Virtualizer I/O Group Name
To be provided.

Storage Virtualizer Node Name
The name that was assigned to the storage virtualizer node when it was added to the system.

Day or Hour or Time
The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Related reports

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click **Storage Virtualizers - Most Active Nodes Report**.

Related tasks:

"Running the Storage Virtualizer - Node Details Report" on page 850
To see information about the performance of a node on a storage virtualizer, run the Storage Virtualizer - Node Details Report.

Running the Storage Virtualizer - Pool Details Report

To see information about the performance of a pool on a storage virtualizer, run the Storage Virtualizer - Pool Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizer - Pool Details Report**.
4. Select a storage virtualizer, and then select a pool.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizer - Pool Details Report"
The Storage Virtualizer - Pool Details Report shows the performance statistics for a pool on a storage virtualizer. Use the report to monitor the performance of a pool over a specified period.

Storage Virtualizer - Pool Details Report

The Storage Virtualizer - Pool Details Report shows the performance statistics for a pool on a storage virtualizer. Use the report to monitor the performance of a pool over a specified period.

Charts

The charts show the following performance statistics for a pool on a storage virtualizer:

- The total overall rates for I/O operations, and the total response times
- The overall I/O rates for read and write operations, and the response times for read and write operations
- The total I/O rates for back-end operations, and overall response times for back-end operations
- The I/O rates for back-end read operations and back-end write operations, and the response rates for back-end read and back-end write operations

Report output

The information that is provided for a pool is as follows:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Day or Hour or Time
The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Related reports

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click **Storage Virtualizers - Most Active Pools Report**.

Related tasks:

"Running the Storage Virtualizer - Pool Details Report" on page 852
To see information about the performance of a pool on a storage virtualizer, run the Storage Virtualizer - Pool Details Report.

Running the Storage Virtualizer - Volume Details Report

To see information about the performance of a volume on a storage virtualizer, run the Storage Virtualizer - Volume Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizer - Volume Details Report**.
4. Select a storage virtualizer.
5. Select a volume. To select a volume, complete the following steps:
   a. Type a keyword. For example, type `lun`, to find the names of volumes that start with `lun`. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type `a%` to find volume names that begin with 'a' or 'A'. Type `%a` to find volume names that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the volumes on a resource.

b. Click **Search**. If the search is successful, the names of the volumes are displayed in the **Choice** list.
c. From the list, select a volume.

6. Optional: Select an interval.
7. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizer - Volume Details Report"
The Storage Virtualizer - Volume Details Report shows the performance statistics of a volume on a storage virtualizer. Use the report to monitor the performance of a volume over a specified period.

Storage Virtualizer - Volume Details Report

The Storage Virtualizer - Volume Details Report shows the performance statistics of a volume on a storage virtualizer. Use the report to monitor the performance of a volume over a specified period.

Charts

The charts show the following performance statistics for a volume on a storage virtualizer:

- The total overall rates for I/O operations, and the total response times
- The overall rates for I/O read operations, the overall rates for I/O write operations, the response times for read operations, and the response times for write operations
- The data rates for read operations, the data rates for write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for volumes

Report output

The information that is provided for a volume is as follows:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Day or Hour or Time
The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Virtualizer Volume Name
The name that was assigned to the storage virtualizer volume when it was added to the system.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Volume Utilization
The average percentage of time that the volume is busy.

Related reports

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click **Storage Virtualizers - Most Active Volumes Report**.

Related tasks:

"Running the Storage Virtualizer - Volume Details Report" on page 854
To see information about the performance of a volume on a storage virtualizer, run the Storage Virtualizer - Volume Details Report.

Running the Storage Virtualizer - Port Details Report

To see information about the performance of a port on a storage virtualizer, run the Storage Virtualizer - Port Details Report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Performance**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizer - Port Details Report**.
4. Select a storage virtualizer, and then select a port.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Virtualizer - Port Details Report"
The Storage Virtualizer - Port Details Report shows the performance statistics for a port on a storage virtualizer. Use the report to monitor the performance of a port over a specified period.

Storage Virtualizer - Port Details Report

The Storage Virtualizer - Port Details Report shows the performance statistics for a port on a storage virtualizer. Use the report to monitor the performance of a port over a specified period.

Charts

The charts show the following performance statistics for a port on a storage virtualizer:

- The total I/O rates
- The send I/O rates for ports, and the receive I/O rates for ports
- The send data rates for ports, the receive data rates for ports, and the total data rates for ports
- The overall bandwidth percentage for ports

Report output

The information that is provided for ports is as follows:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Port Name
An alphanumeric string that identifies a storage virtualizer port.

Day or Hour or Time
The day, hour, or time at which the information about the specified resource or device is provided. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Related reports

To open related reports, click the name of the storage virtualizer, click **Go To**, and then click **Storage Virtualizers - Most Active Ports Report**.

Related tasks:

"Running the Storage Virtualizer - Port Details Report" on page 857
To see information about the performance of a port on a storage virtualizer, run the Storage Virtualizer - Port Details Report.

Predefined reports about historical information

You can run predefined reports about historical information to monitor how space is used on storage systems and storage virtualizers over a specified period.

You can run predefined reports that show how space is used on storage systems and related resources such as pools and volumes.

Running the Storage Virtualizers - Historical Space Report

To see how space on storage virtualizers is used, run the Storage Virtualizers - Historical Space Report.

To run the report with the default settings, click **Finish**. By default all of the storage virtualizers are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Historical Information**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizers - Historical Space Report**.
4. Optional: Select one or more storage virtualizers.
5. Optional: Select an interval.
6. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Storage Virtualizers - Historical Space Report

The Storage Virtualizers - Historical Space Report shows the usage of space on storage virtualizers over a specified period. Use the report to determine whether additional storage resources are required.

Charts

The chart shows the following space statistics for the storage virtualizers:

- The capacity of the pools on the storage virtualizers
- The used space in the pools on the storage virtualizers
- The available space in the pools on the storage virtualizers
- The capacity of the virtual volumes on the storage virtualizers
- The capacity of the managed disks (MDisks) on the storage virtualizers

Report output

The following information is included in reports for storage virtualizers:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Date The date that the data was collected for the report.

Storage Virtualizer Pool Total Volume Capacity (GiB)
The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Virtualizer Pool Used Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can pre-allocate thin-provisioned volume space when the volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. For pools with compressed volumes on SAN Volume Controller and Storwize V7000, the Storage Virtualizer Pool Used Space reflects the size of compressed data written to disk. As the data changes the Storage Virtualizer Pool Used Space may at times be less than the Storage Virtualizer Pool Allocated Space.

For pools with volumes that are not thin provisioned or compressed in SAN Volume Controller, Storwize V7000, and other storage systems, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a pool. Because data collection is run on a set schedule and the used space on volumes can change rapidly, the value in this column might not be 100% accurate for the current state of volumes.

Storage Virtualizer Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Storage Virtualizer Volume Allocated Space (GiB)
The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

Storage Virtualizer MDisk Capacity (GiB)
The amount of storage space on the managed disk.

Storage Virtualizer MDisk Available Space (GiB)
The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

Storage Virtualizer Total Volume Capacity (GiB)
The amount of virtual volume space on all volumes on the storage virtualizer.

Storage Virtualizer Volume Assigned Space (GiB)
The amount of virtual volume space that is assigned to the server or storage virtualizer.

Storage Virtualizer Disk Capacity (GiB)
The amount of local disk storage on the storage virtualizer. The local disks are the internal disk drives.

Storage Virtualizer Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Virtualizer Number of Volumes
The number of volumes in the storage virtualizer pool.

Storage Virtualizer Number of MDisks
The number of managed disks on a resource.

Running the Storage Virtualizer - Pools Historical Space Report

To see how space in pools on a storage virtualizer is used, run the Storage Virtualizer - Pools Historical Space Report.

You can see how space in a pool on a storage virtualizer, or in pools in a storage resource group, are used.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Historical Information**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizer - Pools Historical Space Report**.
4. Required: Select the type of historical space report that you want to create.
5. Optional: You can filter the type of data that you want to see in a report. You can select either of the following options:
   - **Of a Pool in a Storage Virtualizer**
     a. Click **Next**.
     b. Select a storage virtualizer, and then select a pool.
   - **Of Pools in a Storage Resource Group**
     a. Click **Next**.
     b. Select a storage resource group.
6. Optional: Select an interval.
7. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Storage Virtualizer - Pools Historical Space Report

The Storage Virtualizer - Pools Historical Space Report shows the usage of space in a pool on a storage virtualizer, or of pools in a storage resource group over a specified period. Use the report to determine whether additional storage resources are required.

Charts

The chart shows the following space statistics for the pool on the storage virtualizer:

- The capacity of the pool
- The used space in the pool
- The available space in the pool

- The virtual volume capacity of the pool
- The real volume capacity of the pool

The chart shows the following space statistics for the pools in the storage resource group:

- The capacity of the pools
- The used space in the pools
- The available space in the pools
- The virtual volume capacity of the pools
- The real volume capacity of the pools

Report output

The following information is included in reports for a pool on a storage virtualizer, or for pools in a storage resource group:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Date The date that the data was collected for the report.

Storage Virtualizer Pool Capacity (GiB)
The total amount of storage space in a pool.

Storage Virtualizer Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Storage Virtualizer Pool Used Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can pre-allocate thin-provisioned volume space when the volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. For pools with compressed volumes on SAN Volume Controller and Storwize V7000, the Storage Virtualizer Pool Used Space reflects the size of compressed data written to disk. As the data changes the Storage Virtualizer Pool Used Space may at times be less than the Storage Virtualizer Pool Allocated Space.

For pools with volumes that are not thin provisioned or compressed in SAN Volume Controller, Storwize V7000, and other storage systems, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a pool. Because data collection is run on a set schedule and the used space on volumes can change rapidly, the value in this column might not be 100% accurate for the current state of volumes.

Storage Virtualizer Pool Used Volume Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Virtualizer Pool Allocated Space (GiB)
The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes.

The space that is allocated for thin-provisioned volumes is less than their virtual capacity. The Storage Virtualizer Pool Total Volume Capacity (GiB) property shows the virtual capacity of the volumes. If a pool does not contain thin-provisioned volumes, this value is the same as the value of the Storage Virtualizer Pool Total Volume Capacity (GiB) property.

This value is equal to Storage Virtualizer Used Space (GiB) for volumes that are not thin provisioned.

Storage Virtualizer Pool Total Volume Capacity (GiB)
The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Virtualizer Pool Assigned Volume Space (GiB)
The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Virtualizer Pool Unassigned Volume Space (GiB)
The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Virtualizer Pool Number of Storage Volumes
The number of volumes in the storage virtualizer pool.

Storage Virtualizer Pool Number of MDisks
The number of managed disks on a resource.

Running the Storage Virtualizer - Volumes Historical Space Report

To see how space on volumes on a storage virtualizer is used, run the Storage Virtualizer - Volumes Historical Space Report

You can see how space on a volume on a storage virtualizer, or on volumes in a storage resource group, are used.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Historical Information**.
2. Click `Storage Virtualizer`.
3. Click **Storage Virtualizer - Volumes Historical Space Report**.

4. Required: Select the type of historical space report that you want to create.
5. Optional: You can filter the type of data that you want to see in a report. You can select either of the following options:
   - **Of a Volume in a Storage Virtualizer**
     a. Click **Next**.
     b. Select a storage virtualizer.
     c. Type a keyword and click **Search**. For example, type `lun`, to find the names of volumes that start with `lun`. To change the default search option, click **Options**.

        Tip: Use the percent sign (%) as a wildcard. For example, type `a%` to find volume names that begin with 'a' or 'A'. Type `%a` to find volume names that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the volumes on a resource.
     d. Select a volume.
   - **Of Volumes in a Storage Resource Group**
     a. Click **Next**.
     b. Select a storage resource group.
6. Optional: Select an interval.
7. Optional: To limit the time frame of the report, specify a date filter or a custom date range. To reduce the amount of information that is shown in the report, change the default date filter. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of filters.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Storage Virtualizer - Volumes Historical Space Report

The Storage Virtualizer - Volumes Historical Space Report shows the usage of space on a volume on a storage virtualizer, or on volumes in a storage resource group over a specified period. Use the report to determine whether additional storage resources are required.

Charts

The chart shows the following space statistics for the volume on the storage virtualizer:

- The capacity of the volume
- The used space on the volume
- The real space on the volume

The chart shows the following space statistics for the volumes in the storage resource group:

- The capacity of the volumes
- The used space on the volumes
- The real space on the volumes

If the volume is thin provisioned, the chart shows the total amount of virtual space that is allocated to the volume, and the total amount of real space that is used by the volume.

Report output

The following information is included in reports for a volume on a storage virtualizer, or for volumes in a storage resource group:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Virtualizer Volume Name
The name that was assigned to the storage virtualizer volume when it was added to the system.

Date The date that the data was collected for the report.

Storage Virtualizer Volume Capacity (GiB)
The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

Storage Virtualizer Volume Overhead
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Virtualizer Volume Used Space (GiB)
The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Virtualizer Volume Used Space (GiB) might be different than the Storage Virtualizer Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Virtualizer Volume Used Space (GiB) and Storage Virtualizer Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Virtualizer Volume Allocated Space (GiB)
The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

Running the Storage Systems Historical Capacity report

To see changes in the usage of space on storage systems over a specified period, run the Storage Systems Historical Capacity report.

To run the report with the default settings, click **Finish**. By default all of the storage systems are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Storage Systems Historical Capacity**.
4. Optional: Select the configuration of the storage systems.
5. Select one or more storage systems.
6. Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Systems Historical Capacity report"
Shows the usage of space on storage systems over a specified period. Use the report to analyze changes in your storage system requirements.

Storage Systems Historical Capacity report

Shows the usage of space on storage systems over a specified period. Use the report to analyze changes in your storage system requirements.

Charts

The chart shows the following space statistics for the storage systems:

- The capacity of the pools on the storage systems
- The used space in the pools on the storage systems
- The available space in the pools on the storage systems
- The capacity of the volumes on the storage systems
- The available space on the managed disks on the storage systems
- The capacity of the managed disks on the storage systems

Report output

The following information is included in reports for storage systems:

Storage System Name
A user-defined name of the storage system. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Date The date that the data was collected for the report.

Storage System Pool Capacity (GiB)
The amount of storage space in pools that are on the storage system. For an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage System Used Pool Space (GiB)
The amount of space that is in use in all pools on a storage system.

Storage System Available Pool Space (GiB)
The amount of unused space in pools that are on the storage system.

Storage System Volume Capacity (GiB)
The amount of space on all volumes on the storage system.

Storage System Number of Volumes
The number of volumes on a resource.

Storage System Number of Disks
The number of physical disks on the storage system. For a Storwize V7000 storage system that is configured as a back-end device, the value is the number of managed disks on the storage system.

Storage System Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage System Unallocated Disk Space (GiB)
The amount of disk space that can be added to a pool.

Storage System Volume Capacity Assigned to MDisks (GiB)
The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

Storage System Disk Capacity (GiB)
The amount of storage space on physical disks in a storage system, excluding spare disks. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

Storage System Total Disk Capacity (GiB)
The amount of space on physical disks on a storage system, including spare disks.

Storage System Assigned Volume Space (GiB)
The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage System Volume Capacity for z/OS (GiB)
The amount of space on all volumes on the storage system that the z/OS operating system can use.

Storage System Real Available Pool Space (GiB)
The amount of unused space in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System Real Configured Pool Space (GiB)
The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System MDisk Capacity (GiB)
The amount of storage space on the managed disk.

Storage System MDisk Available Space (GiB)
The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Related tasks:

"Running the Storage Systems Historical Capacity report" on page 866
To see changes in the usage of space on storage systems over a specified period, run the Storage Systems Historical Capacity report.

Running the Pools Historical Capacity report

To see how space in pools on storage systems is used, run the Pools Historical Capacity report.

You can see how space in a pool on a storage system, or in pools in a storage resource group, are used.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Pools**.
4. Click **Pools Historical Capacity**.
5. Select the type of historical space report that you want to create. You can select either of the following options:
   - **Of a Pool in a Storage System**
     a. Click **Next**.
     b. Optional: Select the configuration of the storage system.
     c. Select a storage system, and then select a pool.
   - **Of Pools in a Storage Resource Group**
     a. Click **Next**.
     b. Optional: Select the configuration of the storage system.
     c. Select a storage resource group.
6. Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Pools Historical Capacity report"
Shows the usage of space on storage pools over a specified period. Use the report to analyze changes in your storage pool requirements.

Pools Historical Capacity report

Shows the usage of space on storage pools over a specified period. Use the report to analyze changes in your storage pool requirements.

Charts

The chart shows the following space statistics for the pool on a storage system:

- The capacity of the pool
- The used space in the pool
- The available space in the pool
- The virtual volume capacity of the pool
- The used volume space in the pool
- The allocated space in the pool

The chart shows the following space statistics for the pool in the storage resource group:

- The capacity of the pools
- The used space in the pools
- The available space in the pools
- The virtual volume capacity of the pools
- The used volume space in the pools
- The allocated space in the pool

Report output

The following information is included in reports for a pool on a storage system or for pools in a storage resource group:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Date The date that the data was collected for the report.

Storage System Pool Capacity (GiB)
The amount of storage space in pools that are on the storage system. For

an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
pool capacity - allocated space
```

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

Storage Pool Allocated Space (GiB)
The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

Storage Pool Used Space (GiB)
The amount of storage space that is used in the pool.

Storage Pool Total Volume Capacity (GiB)
The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Pool Used Volume Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Pool Assigned Volume Space (GiB)
The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Pool Unassigned Volume Space (GiB)
The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Pool Real Configured Space (GiB)
The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System MDisk Available Space (GiB)
The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage Pool Number of MDisks
The number of managed disks in a storage pool.

Storage Pool Number of Storage Volumes
The number of volumes in the storage pool.

Related tasks:

"Running the Pools Historical Capacity report" on page 868
To see how space in pools on storage systems is used, run the Pools Historical Capacity report.

Running the Volumes Historical Capacity report

To see how space on volumes on storage systems is used, run the Volumes Historical Capacity report.

You can see how space on a volume on a storage system, or on volumes in a storage resource group, are used.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Volumes**.
4. Click **Volumes Historical Capacity**.
5. Select the type of historical capacity report that you want to create. You can select either of the following options:
   - **Of a Volume in a Storage System**
     a. Click **Next**.
     b. Optional: Select the configuration of the storage system.
     c. Select a storage system.
     d. Type a keyword and click **Search**. For example, type `lun`, to find the names of volumes that start with `lun`. To change the default search option, click **Options**.

        Tip: Use the percent sign (%) as a wildcard. For example, type `a%` to find volume names that begin with 'a' or 'A'. Type `%a` to find volume names that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the volumes on a resource.
     e. Select a volume.
   - **Of Volumes in a Storage Resource Group**
     a. Click **Next**.
     b. Select a storage resource group.
6. Select an interval.

7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Volumes Historical Capacity report"
Shows the usage of space on volumes over a specified period. Use the report to analyze changes in your volume requirements, particularly the increasing usage of space on thin-provisioned volumes.

Volumes Historical Capacity report

Shows the usage of space on volumes over a specified period. Use the report to analyze changes in your volume requirements, particularly the increasing usage of space on thin-provisioned volumes.

Charts

The chart shows the following space statistics for the volume on a storage system:

- The capacity of the volume
- The used space on the volume
- The space on the volume that is used for system management
- The allocated space on the volume

The chart shows the following space statistics for the volumes in a storage resource group:

- The capacity of the volumes
- The used space on the volumes
- The space on the volumes that is used for system management
- The allocated space on the volume

If the volume is thin provisioned, the chart shows the total amount of virtual space that is allocated to the volume, and the total amount of real space that is used by the volume.

Report output

The following information is included in reports for a volume on a storage system, or for volumes in a storage resource group:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`,

the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Date The date that the data was collected for the report.

Storage Volume Capacity (GiB)
The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

For XIV systems, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage Volume Used Space (GiB)
The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Volume Allocated Space (GiB)
The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

Storage Volume Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Related tasks:

"Running the Volumes Historical Capacity report" on page 871
To see how space on volumes on storage systems is used, run the Volumes Historical Capacity report.

Storage virtualizer custom reports about performance

You can use IBM Tivoli Common Reporting to create custom performance reports about storage virtualizers. Performance reports can contain detailed information about the performance of monitored storage virtualizers, and some of the properties of those storage virtualizers.

You can create performance reports for storage virtualizers and their components. You can add performance metrics for volumes, back-end arrays, caches, ports, port errors, and other metrics to performance reports. For example, you can add key metrics to a report to create and generate a performance chart for a specific volume of a storage virtualizer. All the performance metrics from the stand-alone GUI are available in the new reporting interface.

Data for storage virtualizers in performance reports

You can include general information, capacity data, properties, and other information about storage virtualizers in performance reports.

Information about storage virtualizers

You can create performance reports that include the following information:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Capacity and usage data

You can create performance reports that include the following information:

Storage Virtualizer Total Pool Capacity (GiB)
The total amount of storage space in pools on the storage virtualizer.

Storage Virtualizer Pool Used Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can pre-allocate thin-provisioned volume space when the volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. For pools with compressed volumes on SAN Volume Controller and Storwize V7000, the Storage Virtualizer Pool Used Space reflects the size of compressed data written to disk. As the data changes the Storage Virtualizer Pool Used Space may at times be less than the Storage Virtualizer Pool Allocated Space.

For pools with volumes that are not thin provisioned or compressed in SAN Volume Controller, Storwize V7000, and other storage systems, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a pool. Because data collection is run on a set schedule and the used space on volumes can change rapidly, the value in this column might not be 100% accurate for the current state of volumes.

Storage Virtualizer Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Storage Virtualizer Total Volume Capacity (GiB)
The amount of virtual volume space on all volumes on the storage virtualizer.

Storage Virtualizer Volume Allocated Space (GiB)
The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

Storage Virtualizer Volume Assigned Space (GiB)
The amount of virtual volume space that is assigned to the server or storage virtualizer.

Storage Virtualizer MDisk Available Space (GiB)
The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

Storage Virtualizer MDisk Capacity (GiB)
The amount of storage space on the managed disk.

Storage Virtualizer Disk Capacity (GiB)
The amount of local disk storage on the storage virtualizer. The local disks are the internal disk drives.

Storage Virtualizer Number of MDisks
The number of managed disks on a resource.

Storage Virtualizer Number of Volumes
The number of volumes on a resource.

Storage Virtualizer Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Virtualizer Physical Allocation Percentage
The percentage of physical space in storage virtualizer pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

`(allocated space ÷ pool capacity) × 100`

For example, the physical allocation percentage is 25% for a total storage virtualizer pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Virtualizer Virtual Allocation Percentage
The percentage of physical space in storage virtualizer pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(total volume capacity ÷ pool capacity) × 100`

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB.

This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage Virtualizer Shortfall Percentage property to estimate how critical the shortage of space is for storage virtualizer pools.

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Shortfall Percentage

The percentage of the remaining unallocated volume space in storage virtualizer pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(unallocatable space ÷ (volume space - used volume space)) × 100`

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space.

This Storage Virtualizer Shortfall Percentage represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB ÷ 100 GiB × 100).

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

This property is available only for storage virtualizer pools with thin-provisioned volumes.

Component properties

You can create performance reports that include the following information:

Storage Virtualizer Access Information
The URL from which you can access the storage system or storage virtualizer.

Storage Virtualizer Alias
Shows an alternative name for a storage virtualizer.

Storage Virtualizer Code Level
The microcode or firmware level of the storage virtualizer.

Storage Virtualizer Time Zone
The time zone in which a resource is located.

Storage Virtualizer Custom Tag 1, 2, and 3
User-defined text that is associated with a storage virtualizer. You can add or edit the custom tags for a storage virtualizer on the **Properties** page of the storage system in the Tivoli Storage Productivity Center web-based GUI.

Storage Virtualizer User Provided Name
The name that was specified for the storage virtualizer in the storage environment.

Storage Virtualizer Is Compression Active
Shows whether the compression feature is enabled on the storage virtualizer. If this value is `Yes`, the compression feature is enabled.

Status information

You can create performance reports that include the following information:

Storage Virtualizer Status
The condition of the resource, for example normal, warning, or error.

Storage Virtualizer Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Storage Virtualizer Last Data Collection
The date and time when storage statistics were last collected from the resource.

Storage Virtualizer Last Data Collection Status
The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

Vendor, model, and device information

You can create performance reports that include the following information:

Storage Virtualizer IP Address
The IP address of the resource.

Storage Virtualizer Serial Number
The serial number of the resource.

Storage Virtualizer Type
The type of storage virtualizer. For example, the storage virtualizer can be an **IBM Storwize V7000** system, an **IBM SAN Volume Controller** system, or another type of storage virtualizer.

Storage Virtualizer Model
The model name or model number of the resource.

Storage Virtualizer Vendor
The vendor who supplied the resource.

Performance metrics for storage virtualizers:

You can include performance data for ports, port errors, volumes, caches, back-end arrays, and other data for storage virtualizers in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port to Host Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to host computers.

Port to Host Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from host computers.

Total Port to Host I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between host computers and the component.

Port to Disk Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to back-end storage resources.

Port to Disk Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from back-end storage resources.

Total Port to Disk I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

Port to Local Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

Total Port to Local Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

Port to Remote Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

Total Port to Remote Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

Port to Host Send Data Rate (MiB/s)
The average rate at which data is sent to host computers. The rate is measured in MiB per second.

Port to Host Receive Data Rate (MiB/s)
The average rate at which data is received from host computers. The rate is measured in MiB per second.

Total Port to Host Data Rate (MiB/s)
The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Disk Send Data Rate (MiB/s)
The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

Port to Disk Receive Data Rate (MiB/s)
The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

Total Port to Disk Data Rate (MiB/s)
The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Local Node Send Data Rate (MiB/s)
The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Receive Data Rate (MiB/s)
The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Local Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Send Data Rate (MiB/s)
The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Receive Data Rate (MiB/s)
The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Remote Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Local Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Overall Port to Local Node Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation with another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Local Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.

Overall Port to Local Node Queue Time (ms/op)
The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Remote Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Overall Port to Remote Node Response Time (ms/op)
The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Remote Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote SAN Volume Controller cluster.

Port to Remote Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote SAN Volume Controller cluster.

Overall Port to Remote Node Queue Time (ms/op)
The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote SAN Volume Controller cluster.

Port error data

You can create performance reports that include the following information:

Link Failure Rate (count/s)
The average number of link errors per second.

Loss of Sync Rate (count/s)
The average number of times per second that synchronization is lost.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)
The average number of times per second that the signal is lost.

CRC Error Rate (count/s)
The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

Primitive Sequence Protocol Error Rate (count/s)
The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)
The average number of bit errors per second that are detected.

Volume data

You can create performance reports that include the following information:

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Read Ahead Percentage of Cache Hits
The percentage of all read-cache hits that occur on pre-staged data.

Dirty Write Percentage of Cache Hits
The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

Write Cache Overflow Percentage
The percentage of write operations that are delayed because there is not enough space in the write cache.

Write Cache Overflow I/O Rate (ops/s)
The average number of write operations per second that are delayed because there is not enough space in the write cache.

Write Cache Flush Through Percentage
The percentage of tracks that is written to disk in flush-through mode.

Write Cache Flush Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in flush-through mode.

Write Cache Write Through Percentage
The percentage of tracks that is written to disk in write-through mode.

Write Cache Write Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in write-through mode.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Peak Back-End Read Response Time (ms)
The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)
The longest time for a back-end storage resource to respond to a write operation.

Peak Back-End Read Queue Time (ms)
The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)
The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write Percentage
The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Global Mirror Overlapping Write I/O Rate (op/s)
The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Peak Read Response Time (ms)
The longest time to complete a read operation.

Peak Write Response Time (ms)
The longest time to complete a write operation.

Global Mirror Secondary Write Lag (ms/op)
The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Overall Host Attributed Response Time Percentage
The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

The value is based on the following conditions:

- For read response time, the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes.
- For write response time, the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Zero Buffer-to-Buffer Credit Timer (microseconds)
The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

Performance metrics for storage systems that are configured to virtualize storage:

A storage environment can include a Hitachi Universal Storage Platform V system that is configured to virtualize storage. You can include performance data for volumes, and cache data for these systems in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Data for managed disks on storage virtualizers in performance reports

You can include general information, capacity data, properties, and other information about managed disks on storage virtualizers in performance reports.

Information about managed disks on storage virtualizers

You can create performance reports that include the following information:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Virtualizer MDisk Name
The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

Capacity and usage data

You can create performance reports that include the following information:

Storage Virtualizer MDisk Capacity (GiB)
The amount of storage space on the managed disk.

Storage Virtualizer MDisk Available Space (GiB)
The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

Storage Virtualizer MDisk Strip Size (KB)
The RAID strip size on a managed disk on a storage virtualizer. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Component properties

You can create performance reports that include the following information:

Storage Virtualizer MDisk RAID Level
The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Spare Goal
The number of spare drives that are required to maintain redundancy. Use spare drives to protect against drive failures in the array on a managed disk on a storage virtualizer.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Type
The type of managed disk on a storage virtualizer. For example, the storage virtualizer can be a local or virtual managed disk.

Storage Virtualizer MDisk Mode
The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

Array Extents are provided from local disks.

Managed
Extents are provided from a back-end storage volume.

Unmanaged
The managed disk is not used in the system.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Is Balanced
Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is `Yes`, the LUNs are balanced.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Fast Write State
Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Write Verify
Shows whether all write operations on a managed disk on a storage virtualizer are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is `Yes`, all write operations are verified by a follow-up read operation.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage Virtualizer MDisk Status
The condition of the resource, for example normal, warning, or error.

Storage Virtualizer MDisk Native Status
Shows the level of access of nodes in the system to a managed disk or volume on a storage virtualizer. The level of access can be online, offline, degraded, or excluded.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Performance metrics for managed disks on storage virtualizers:

You can include performance data for back-end arrays on storage-virtualizer managed disks in performance reports.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Peak Back-End Read Response Time (ms)
The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)
The longest time for a back-end storage resource to respond to a write operation.

Peak Back-End Read Queue Time (ms)
The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Data for storage virtualizer volumes in performance reports

You can include general information, capacity data, properties, and other information about storage virtualizer volumes in performance reports.

Information about storage virtualizer volumes

You can create performance reports that include the following information:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Virtualizer I/O Group Name
The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

Storage Virtualizer Volume Name
The name that was assigned to the storage virtualizer volume when it was added to the system.

Storage Virtualizer Volume Assigned Host Alias
The user-defined name that is assigned to the port on the server. Volumes are associated with the name of the port on the server.

Capacity and usage data

You can create performance reports that include the following information:

Storage Virtualizer Volume Capacity (GiB)
The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

Storage Virtualizer Volume Allocated Space (GiB)
The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

Storage Virtualizer Volume Used Space (GiB)
The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Virtualizer Volume Used Space (GiB) might be different than the Storage Virtualizer Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Virtualizer Volume Used Space (GiB) and Storage Virtualizer Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Virtualizer Volume Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Virtualizer Volume Used Allocated Space Percentage
The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

```
(used space ÷ allocated space) × 100
```

This property is available only for SAN Volume Controller and Storwize V7000 volumes.

Storage Virtualizer Volume Unallocated Space (GiB)
The amount of space in a pool that is not reserved for a volume.

Tivoli Storage Productivity Center uses the following formula to determine the value:

```
capacity - allocated space
```

Storage Virtualizer Volume Physical Allocation Percentage
The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

```
(allocated space ÷ capacity) × 100
```

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore the space that is reserved for volumes is 50 GiB.

Storage Virtualizer Volume Unused Space (GiB)
The amount of space that is allocated to a volume and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
allocated space - used Space
```

This property is available only for SAN Volume Controller and Storwize V7000 pools.

Storage Virtualizer Volume Shortfall Percentage
The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume. The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

```
(unallocatable volume space ÷ unallocated volume space) × 100
```

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level for the volume. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

```
(100 GiB unallocated volume space - 70 GiB remaining pool space) ÷ 100 GiB
unallocated volume space × 100
```

This value is only available for pools with thin-provisioned volumes.

Storage Virtualizer Volume Unallocatable Space (GiB)
The unallocatable storage space in GiB of a thin-provisioned volume. Unallocatable space cannot be supplied by the pool to which the volume belongs.

Storage Virtualizer Volume Uncompressed Capacity (GiB)
The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Component properties

You can create performance reports that include the following information:

Storage Virtualizer Volume Type
The type of storage virtualizer volume. For example, the storage virtualizer volume can be striped, sequential, or another type of storage virtualizer volume. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Volume WWN
The worldwide name of the volume.

Storage Virtualizer Volume Fast Write State
Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted.

Storage Virtualizer Volume Grain Size (KiB)
The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

Storage Virtualizer Volume Is Assigned
Shows whether the volume is assigned to a server. If this value is `Yes`, the volume is assigned to a server.

Storage Virtualizer Volume Is Auto Expand Enabled
Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is `Yes`, the volume or volume copy can allocate new extents from a pool automatically.

Storage Virtualizer Volume Is Encryptable
Shows whether the resource can be encrypted. If this value is `Yes`, the resource can be encrypted.

Storage Virtualizer Volume Is Encrypted
Shows whether the resource is encrypted. If this value is `Yes`, the resource is encrypted.

Storage Virtualizer Volume Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Virtualizer Volume Metro Mirror Name
The name of the Metro Mirror that keeps the synchronous copy of the storage virtualizer volume.

Storage Virtualizer Volume Mirror Count
The number of mirrors that keep a synchronous copy of the resource.

Storage Virtualizer Volume Preferred Node
The number of the storage virtualizer node that is used for I/O operations for the storage virtualizer volume.

Storage Virtualizer Volume Throttle
The maximum number of commands that the storage virtualizer volume can queue.

Storage Virtualizer Volume Warning Level
The percentage of used capacity of the resource at which a warning is generated.

Storage Virtualizer Volume Copy Relationship
Shows whether a volume or a copy of a volume is in a FlashCopy relationship or a point-in-time relationship. This property can contain the following values:

Source Relationship
The volume is the source of the relationship.

Target Relationship
The volume is the target of the relationship.

Both Source and Target Relationship
The volume is both the source and target of the relationship.

Storage Virtualizer Volume Is Compressed
Shows whether the storage virtualizer volume is compressed. If this value is `Yes`, the volume is compressed.

Storage Virtualizer Volume RAID Level
The RAID level of a storage virtualizer volume, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Status information

You can create performance reports that include the following information:

Storage Virtualizer Volume Status
The condition of the resource, for example normal, warning, or error.

Storage Virtualizer Volume Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Storage Virtualizer Volume Native Status
Shows the level of access of nodes in the system to the storage virtualizer managed disk or volume. The level of access can be online, offline, degraded, or excluded.

Performance metrics for storage virtualizer volumes:

You can include performance data for volumes, caches, and other data for storage virtualizer volumes in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Volume Utilization
The average percentage of time that the volume is busy.

Maximum Read I/O Rate (ops/s)
The maximum read I/O rate of the volume in operations per second. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Write I/O Rate (ops/s)
The maximum write I/O rate of the volume in operations per second. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Read Response Time (ms/op)
The maximum read response time of the volume in milliseconds per operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Write Response Time (ms/op)
The maximum write response time of the volume in milliseconds per operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Write I/O Size (KiB)
The maximum size of write I/O operations in kibibytes. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Volume Utilization Percentage
The maximum period in percentage that the volume is busy. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Read Ahead Percentage of Cache Hits
The percentage of all read-cache hits that occur on pre-staged data.

Dirty Write Percentage of Cache Hits
The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

Write Cache Overflow Percentage
The percentage of write operations that are delayed because there is not enough space in the write cache.

Write Cache Overflow I/O Rate (ops/s)
The average number of write operations per second that are delayed because there is not enough space in the write cache.

Write Cache Flush Through Percentage
The percentage of tracks that is written to disk in flush-through mode.

Write Cache Flush Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in flush-through mode.

Write Cache Write Through Percentage
The percentage of tracks that is written to disk in write-through mode.

Write Cache Write Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in write-through mode.

Maximum Read Cache Hit Percentage
The maximum number of cache hits in percentage for sequential and nonsequential read operations for the volume over a specified interval. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Write Cache Delay I/O Rate (ops/s)
The maximum number of delayed I/O operations per second for the volume, where the delay occurred because of space constraints or other conditions.

This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)
The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write Percentage
The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Global Mirror Overlapping Write I/O Rate (op/s)
The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Peak Read Response Time (ms)
The longest time to complete a read operation.

Peak Write Response Time (ms)
The longest time to complete a write operation.

Global Mirror Secondary Write Lag (ms/op)
The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

Overall Host Attributed Response Time Percentage
The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times

and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

The value is based on the following conditions:

- For read response time, the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes.
- For write response time, the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Nonpreferred Node Usage Percentage
The overall percentage of I/O operations that are completed, or data that is transferred over the non-preferred nodes of the virtual disks.

Performance metrics for volumes on storage systems that are configured to virtualize storage:

A storage environment can include a Hitachi Universal Storage Platform V system that is configured to virtualize storage. You can include performance data for volumes and caches for storage system volumes on these systems in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Volume Utilization
The average percentage of time that the volume is busy.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Data for storage virtualizer nodes in performance reports

You can include properties, and other information about storage virtualizer nodes in performance reports.

Information about storage virtualizer nodes

You can create performance reports that include the following information:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Node Name
The name that was assigned to the storage virtualizer node when it was added to the system.

Storage Virtualizer I/O Group Name
The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

Component properties

You can create performance reports that include the following information:

Storage Virtualizer Node IP Address
The IP address of the resource.

Storage Virtualizer Node WWN
The worldwide name of the node.

Storage Virtualizer Node Configuration
The number of the node for the resource. This value can be either `0` or `1`.

Status information

You can create performance reports that include the following information:

Storage Virtualizer Node Status
Shows whether a storage virtualizer node is online.

Performance metrics for storage virtualizer nodes:

You can include port data, front end data, cache data, back-end data, and other data for storage virtualizer nodes in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port to Host Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to host computers.

Port to Host Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from host computers.

Total Port to Host I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between host computers and the component.

Port to Disk Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to back-end storage resources.

Port to Disk Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from back-end storage resources.

Total Port to Disk I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

Port to Local Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

Total Port to Local Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

Port to Remote Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

Total Port to Remote Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

Port to Host Send Data Rate (MiB/s)
The average rate at which data is sent to host computers. The rate is measured in MiB per second.

Port to Host Receive Data Rate (MiB/s)
The average rate at which data is received from host computers. The rate is measured in MiB per second.

Total Port to Host Data Rate (MiB/s)
The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Disk Send Data Rate (MiB/s)
The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

Port to Disk Receive Data Rate (MiB/s)
The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

Total Port to Disk Data Rate (MiB/s)
The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Local Node Send Data Rate (MiB/s)
The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Receive Data Rate (MiB/s)
The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Local Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Send Data Rate (MiB/s)
The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Receive Data Rate (MiB/s)
The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Remote Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Local Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Overall Port to Local Node Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation with another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Local Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.

Overall Port to Local Node Queue Time (ms/op)
The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Remote Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Overall Port to Remote Node Response Time (ms/op)
The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Remote Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote SAN Volume Controller cluster.

Port to Remote Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote SAN Volume Controller cluster.

Overall Port to Remote Node Queue Time (ms/op)
The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote SAN Volume Controller cluster.

Front end data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Read Ahead Percentage of Cache Hits
The percentage of all read-cache hits that occur on pre-staged data.

Dirty Write Percentage of Cache Hits
The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

Write Cache Overflow Percentage
The percentage of write operations that are delayed because there is not enough space in the write cache.

Write Cache Flush Through Percentage
The percentage of tracks that is written to disk in flush-through mode.

Write Cache Write Through Percentage
The percentage of tracks that is written to disk in write-through mode.

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Write Cache Overflow I/O Rate (ops/s)
The average number of write operations per second that are delayed because there is not enough space in the write cache.

Write Cache Flush Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in flush-through mode.

Write Cache Write Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in write-through mode.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Back-end data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Peak Back-End Read Response Time (ms)
The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)
The longest time for a back-end storage resource to respond to a write operation.

Peak Back-End Read Queue Time (ms)
The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)
The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write Percentage
The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Global Mirror Overlapping Write I/O Rate (op/s)
The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Global Mirror Secondary Write Lag (ms/op)
The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

Peak Read Response Time (ms)
The longest time to complete a read operation.

Peak Write Response Time (ms)
The longest time to complete a write operation.

Link Failure Rate (count/s)
The average number of link errors per second.

Loss of Sync Rate (count/s)
The average number of times per second that synchronization is lost.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)
The average number of times per second that the signal is lost.

CRC Error Rate (count/s)
The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

Primitive Sequence Protocol Error Rate (count/s)
The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)
The average number of bit errors per second that are detected.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Overall Host Attributed Response Time Percentage
The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

The value is based on the following conditions:

- For read response time, the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes.
- For write response time, the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Zero Buffer-to-Buffer Credit Timer (microseconds)
The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

Data for storage virtualizer I/O groups in performance reports

You can include general information, capacity data, and properties of storage virtualizer I/O groups in performance reports.

Information about storage virtualizer I/O groups

You can create performance reports that include the following information:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer I/O Group Name
The name that was assigned to the I/O group on a storage virtualizer when it was added to the system.

Capacity and usage data

You can create performance reports that include the following information:

Storage Virtualizer I/O Group Mirroring Total Memory (MiB)
The total amount of memory that is allocated on the nodes in an I/O group to provide volume mirroring functions.

Storage Virtualizer I/O Group Mirroring Free Memory (MiB)
The amount of available memory on the nodes in an I/O group to provide volume mirroring functions.

Storage Virtualizer I/O Group FlashCopy Total Memory (MiB)
The total amount of memory that is allocated on the nodes in an I/O group to provide FlashCopy mirroring functions.

Storage Virtualizer I/O Group FlashCopy Free Memory (MiB)
The amount of available memory on the nodes in an I/O group to provide FlashCopy mirroring functions.

Storage Virtualizer I/O Group Remote Copy Total Memory (MiB)
The total amount of memory that is allocated on the nodes in an I/O group to provide mirroring functions by using a remote copy.

Storage Virtualizer I/O Group Remote Copy Free Memory (MiB)
The amount of available memory on the nodes in an I/O group to provide mirroring functions by using a remote copy.

Storage Virtualizer I/O Group RAID Total Memory (MiB)
The total amount of memory that is allocated on the nodes in an I/O group to provide RAID mirroring functions.

Storage Virtualizer I/O Group RAID Free Memory (MiB)
The amount of available memory on the nodes in an I/O group to provide RAID mirroring functions.

Component properties

You can create performance reports that include the following information:

Storage Virtualizer I/O Group Maintenance
Shows whether the I/O group is in maintenance mode. If this value is `Yes`, the I/O group is in maintenance mode.

Storage Virtualizer I/O Group Is Compression Active
Shows whether the compression feature is enabled on the volumes in the storage virtualizer I/O group. If this value is `Yes`, the compression feature is enabled.

Storage Virtualizer I/O Group Is Compression Supported
Shows whether the compression feature is available for volumes in the storage virtualizer I/O group. If this value is `Yes`, the compression feature is available.

Performance metrics for storage virtualizer I/O groups:

You can include performance data for ports, volumes, caches, back-end arrays, and other data for storage virtualizer I/O groups in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port to Host Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to host computers.

Port to Host Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from host computers.

Total Port to Host I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between host computers and the component.

Port to Disk Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to back-end storage resources.

Port to Disk Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from back-end storage resources.

Total Port to Disk I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

Port to Local Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

Total Port to Local Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

Port to Remote Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

Total Port to Remote Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

Port to Host Send Data Rate (MiB/s)
The average rate at which data is sent to host computers. The rate is measured in MiB per second.

Port to Host Receive Data Rate (MiB/s)
The average rate at which data is received from host computers. The rate is measured in MiB per second.

Total Port to Host Data Rate (MiB/s)
The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Disk Send Data Rate (MiB/s)
The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

Port to Disk Receive Data Rate (MiB/s)
The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

Total Port to Disk Data Rate (MiB/s)
The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Local Node Send Data Rate (MiB/s)
The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Receive Data Rate (MiB/s)
The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Local Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Send Data Rate (MiB/s)
The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Receive Data Rate (MiB/s)
The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Remote Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Local Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Overall Port to Local Node Response Time (ms/op)
The average number of milliseconds to complete a send or receive

operation with another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Local Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.

Overall Port to Local Node Queue Time (ms/op)
The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Remote Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Overall Port to Remote Node Response Time (ms/op)
The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Remote Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote SAN Volume Controller cluster.

Port to Remote Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote SAN Volume Controller cluster.

Overall Port to Remote Node Queue Time (ms/op)
The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote SAN Volume Controller cluster.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Back-end array data

You can create performance reports that include the following information:

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Peak Back-End Read Response Time (ms)
The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)
The longest time for a back-end storage resource to respond to a write operation.

Peak Back-End Read Queue Time (ms)
The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Read Ahead Percentage of Cache Hits
The percentage of all read-cache hits that occur on pre-staged data.

Dirty Write Percentage of Cache Hits
The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

Write Cache Overflow Percentage
The percentage of write operations that are delayed because there is not enough space in the write cache.

Write Cache Overflow I/O Rate (ops/s)
The average number of write operations per second that are delayed because there is not enough space in the write cache.

Write Cache Flush Through Percentage
The percentage of tracks that is written to disk in flush-through mode.

Write Cache Flush Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in flush-through mode.

Write Cache Write Through Percentage
The percentage of tracks that is written to disk in write-through mode.

Write Cache Write Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in write-through mode.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)
The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write Percentage
The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Global Mirror Overlapping Write I/O Rate (op/s)
The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Peak Read Response Time (ms)
The longest time to complete a read operation.

Peak Write Response Time (ms)
The longest time to complete a write operation.

Global Mirror Secondary Write Lag (ms/op)
The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

Link Failure Rate (count/s)
The average number of link errors per second.

Loss of Sync Rate (count/s)
The average number of times per second that synchronization is lost.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)
The average number of times per second that the signal is lost.

CRC Error Rate (count/s)
The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Overall Host Attributed Response Time Percentage
The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

The value is based on the following conditions:

- For read response time, the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes.
- For write response time, the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Nonpreferred Node Usage Percentage
The overall percentage of I/O operations that are completed, or data that is transferred over the non-preferred nodes of the virtual disks.

Primitive Sequence Protocol Error Rate (count/s)
The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)
The average number of bit errors per second that are detected.

Zero Buffer-to-Buffer Credit Timer (microseconds)
The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

Data for storage virtualizer ports in performance reports

You can include properties and other information about storage virtualizer ports in performance reports.

Information about storage virtualizer ports

You can create performance reports that include the following information:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Node Name
The name that was assigned to the storage virtualizer node when it was added to the system.

Storage Virtualizer Port Name
An alphanumeric string that identifies a storage virtualizer port.

Component properties

You can create performance reports that include the following information:

Storage Virtualizer Port Number
The port number on the resource.

Storage Virtualizer Port Type
The type of port on the storage system, storage virtualizer, or switch. For example, the port type can be N_Port, F_Port, or another type of port.

Storage Virtualizer Port Speed (GiB)
The speed of a port, which is measured in GiB per second.

Status information

You can create performance reports that include the following information:

Storage Virtualizer Port Status
The condition of the resource, for example normal, warning, or error.

Storage Virtualizer Port Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Storage Virtualizer Port Enabled State
Shows whether a port is enabled, disabled, or is enabled but offline.

Performance metrics for storage virtualizer ports:

You can include performance data about storage virtualizer ports and port errors in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port to Host Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to host computers.

Port to Host Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from host computers.

Total Port to Host I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between host computers and the component.

Port to Disk Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to back-end storage resources.

Port to Disk Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from back-end storage resources.

Total Port to Disk I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

Port to Local Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

Total Port to Local Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

Port to Remote Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

Total Port to Remote Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

Port to Host Send Data Rate (MiB/s)
The average rate at which data is sent to host computers. The rate is measured in MiB per second.

Port to Host Receive Data Rate (MiB/s)
The average rate at which data is received from host computers. The rate is measured in MiB per second.

Total Port to Host Data Rate (MiB/s)
The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Disk Send Data Rate (MiB/s)
The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

Port to Disk Receive Data Rate (MiB/s)
The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

Total Port to Disk Data Rate (MiB/s)
The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Local Node Receive Data Rate (MiB/s)
The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Send Data Rate (MiB/s)
The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Local Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Receive Data Rate (MiB/s)
The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Send Data Rate (MiB/s)
The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Remote Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port Send Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage
The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port error data

You can create performance reports that include the following information:

Link Failure Rate (count/s)
The average number of link errors per second.

Loss of Sync Rate (count/s)
The average number of times per second that synchronization is lost.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)
The average number of times per second that the signal is lost.

CRC Error Rate (count/s)
The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

Primitive Sequence Protocol Error Rate (count/s)
The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)
The average number of bit errors per second that are detected.

Zero Buffer-to-Buffer Credit Timer (microseconds)
The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

Zero Buffer-to-Buffer Credit Percentage
The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

Performance metrics for ports on storage systems that are configured to virtualize storage:

A storage environment can include a Hitachi Universal Storage Platform V system that is configured to virtualize storage. You can include performance data for ports on these systems in performance reports.

Port data

You can create performance reports that include the following information:

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Data for storage virtualizer pools in performance reports

You can include general information, capacity data, properties, and other information about storage virtualizer pools in performance reports.

Information about storage virtualizer pools

You can create performance reports that include the following information:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Capacity and usage data

You can create performance reports that include the following information:

Storage Virtualizer Pool Capacity (GiB)
The total amount of storage space in a pool.

Storage Virtualizer Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Storage Virtualizer Pool Total Volume Capacity (GiB)
The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Virtualizer Pool Allocated Space (GiB)
The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes.

The space that is allocated for thin-provisioned volumes is less than their virtual capacity. The Storage Virtualizer Pool Total Volume Capacity (GiB) property shows the virtual capacity of the volumes. If a pool does not contain thin-provisioned volumes, this value is the same as the value of the Storage Virtualizer Pool Total Volume Capacity (GiB) property.

This value is equal to Storage Virtualizer Used Space (GiB) for volumes that are not thin provisioned.

Storage Virtualizer Pool Used Volume Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Virtualizer Pool Extent Size (MiB)
The size of the extents that were specified when a pool was created. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Assigned Volume Space (GiB)
The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Virtualizer Pool Unassigned Volume Space (GiB)
The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Virtualizer Pool Number of Storage Volumes
The number of volumes in the storage virtualizer pool.

Storage Virtualizer Pool Number of MDisks
The number of managed disks on a resource.

Storage Virtualizer Pool Physical Allocation Percentage
The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$(\text{allocated space} \div \text{pool capacity}) \times 100$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Virtualizer Pool Virtual Allocation Percentage
The percentage of physical space in a pool that is committed to the total

virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(total volume capacity ÷ pool capacity) × 100`

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. In this case, the virtual capacity that is committed to the volumes in the pool is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, use the Storage Virtualizer Pool Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

This value is only available for pools with thin-provisioned volumes.

Storage Virtualizer Pool Shortfall Percentage
The percentage of the remaining unallocated volume space in a pool that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(unallocatable space ÷ (total volume capacity − allocated space)) × 100`

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Unallocated Volume Space (GiB)
The amount of the Total Volume Capacity in the pool that is not allocated.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

```
total volume capacity – allocated space
```

The space that is allocated for thin-provisioned volumes is typically less than their virtual capacity. Therefore, the unallocated space represents the difference between the virtual capacity and the allocated space for all the volumes in the pool.

Storage Virtualizer Pool Unused Volume Space (GiB)
The amount of space that is allocated to the volumes in a pool and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
allocated space – used pool space
```

This property is available only for SAN Volume Controller and Storwize V7000 pools.

Storage Virtualizer Pool Unallocatable Volume Space (GiB)
The amount of space by which the Total Volume Capacity exceeds the physical capacity of a pool.

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool. If you create volumes with more virtual capacity than can be physically allocated in the pool, you can over commit storage in the pool.

This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.

Storage Virtualizer Pool Used Space Percentage
The percentage of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

Storage Virtualizer Pool Compressed Virtual Capacity (GiB)
The total virtual capacity of all the volumes that are compressed in a storage virtualizer pool.

Storage Virtualizer Pool Compressed Capacity (GiB)
The amount of storage space that is used by compressed volumes in a storage virtualizer pool. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 20.

Storage Virtualizer Pool Uncompressed Capacity (GiB)
The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Component properties

You can create performance reports that include the following information:

Storage Virtualizer Pool Is Encryptable
Shows whether the storage virtualizer pool can be encrypted. If this value is `Yes`, the storage virtualizer pool can be encrypted.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Is Encrypted
Shows whether the resource is encrypted. If this value is `Yes`, the resource is encrypted.

Storage Virtualizer Pool Is Solid State
Shows whether there are solid-state drives in the pool. This property can contain the following values:

Mixed The pool contains both hard disk drives and solid-state drives.

Non solid state
The pool contains no solid-state drives.

Solid state
The pool contains a least one solid-state drive.

Storage Virtualizer Pool Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Virtualizer Pool RAID Level
The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Virtualizer Pool Warning Level
The percentage of used capacity of the resource at which a warning is generated. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Is Compression Active
Shows whether the compression feature is enabled on the storage virtualizer pool. If this value is `Yes`, the compression feature is enabled.

Status information

You can create performance reports that include the following information:

Storage Virtualizer Pool Status
The condition of the resource, for example normal, warning, or error.

Storage Virtualizer Pool Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Performance metrics for storage virtualizer pools:

You can include performance data for volumes and back-end arrays for storage virtualizer pools in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Maximum Read I/O Rate (ops/s)
The maximum read I/O rate of the volume in operations per second. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Peak Back-End Read Response Time (ms)
The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)
The longest time for a back-end storage resource to respond to a write operation.

Peak Back-End Read Queue Time (ms)
The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Maximum Back-End Read I/O Rate (ops/s)
The maximum number of read operations per second that are issued to the back-end storage resources. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Back-End Write I/O Rate (ops/s)
The maximum number of write operations per second that are issued to the back-end storage resources. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Back-End Read Response Time (ms/op)
The maximum number of milliseconds for the back-end storage resources to respond to a read operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Back-End Write Response Time (ms/op)
The maximum number of milliseconds for the back-end storage resources to respond to a write operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Predefined reports about resource relationships

You can run predefined reports about the end-to-end relationship from a server to back-end storage. You can also run reports about the end-to-end relationship from file systems on servers to volumes on storage pools.

Running the Storage Resource Relationships Summary report

To show the relationships between storage resources, run the Storage Resource Relationships Summary report.

To run the report with the default settings, click **Finish**.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Resource Relationships Summary**.
3. Optional: Select one or more of the following:
   a. Servers
   b. Hypervisors
   c. NAS resources
   d. Storage virtualizers
   e. Storage systems
4. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

Restriction: To create a report when you select a single resource, the selected resource must be connected to other resources.

Related reference:

"Storage Resource Relationships Summary report"
Shows the end-to-end relationship from a server to back-end storage. The report shows servers, hypervisors, NAS, storage systems, storage virtualizers, pools, and volumes.

Storage Resource Relationships Summary report

Shows the end-to-end relationship from a server to back-end storage. The report shows servers, hypervisors, NAS, storage systems, storage virtualizers, pools, and volumes.

Report output

Information is provided about related resources and the devices, such as managed disks (MDisks), that are related to resources.

Server Name
The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Hypervisor Name
The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

NAS Name
The fully qualified domain name of the NAS device. For example, the name of a NAS device might be `nas.example.com`.

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Volume Name
The name that was assigned to the storage virtualizer volume when it was added to the system.

Storage Virtualizer Volume Assigned Host Connection
The host connection to which the storage virtualizer volume is assigned. The host connection is a definition in the storage virtualizer that contains the WWPN for the server. The storage virtualizer uses the WWPN to assign volumes to servers.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Virtualizer MDisk Name
The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Volume Assigned Host Connection
The host connection to which the storage volume is assigned. The host connection is a definition in the storage system that contains the WWPN for the server. The storage system uses the WWPN to assign volumes to servers.

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Volume	Performance of One Volume
Managed disk	Performance of One Managed Disk
Storage pool	Most Active Pools, Performance of One Pool

Related tasks:

"Running the Storage Resource Relationships Summary report" on page 930
To show the relationships between storage resources, run the Storage Resource Relationships Summary report.

Running the Storage Resource Relationships Summary (Configurable) report

To show the relationships between resources, run the Storage Resource Relationships Summary (Configurable) report. You can show or hide resources, and you can hide resources that are not related to other resources or storage devices.

To run the report with the default settings, click **Finish**. Servers, storage systems, storage virtualizers, and their related resources are shown in the report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

Tip: To create a report for a single resource, the selected resource must be connected to other storage devices or resources. To create a report about two or more selected resources, the resources that you select must be related.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Resource Relationships Summary (Configurable)**.
3. Optional: Select the storage resources that you want to show in the report.
4. Optional: Specify whether you want to show or hide a resource that is related to other storage devices such as a file system, but not to other resources. For example, a server has a related file system. You can show the server and its related storage devices in the report, or you can hide the resource and its related storage devices in the report.
5. Optional: Click **Finish** to run the report.
6. Optional: Click **Next**. In the Select storage resources page, select one or more of the following:

a. Servers
b. Hypervisors
c. NAS resources
d. Storage virtualizers
e. Storage systems

7. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

Related reference:

"Storage Resource Relationships Summary (Configurable) report"
Shows the end-to-end relationship from a server to back-end storage. You can configure which resources to show in the report.

Storage Resource Relationships Summary (Configurable) report

Shows the end-to-end relationship from a server to back-end storage. You can configure which resources to show in the report.

Report output

Information is provided about related resources and the devices, such as managed disks (MDisks), that are related to resources.

Server Name
The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Server Disk Path
The identifier that specifies a path to a server disk, for example `/dev/hdisk0`.

Logical Volume Path
The path to a logical volume on a resource, for example `/dev/hd1`.

File System Mount Point
The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be `c:\` or `d:\`. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be `/opt` or `/export/home`.

Hypervisor Name
The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

Hypervisor Disk Path
The identifier that specifies a path to a hypervisor server disk, for example `vmhba1:0:0`.

NAS Name
The fully qualified domain name of the NAS device. For example, the name of a NAS device might be `nas.example.com`.

NAS Export Name
The name of the exported file system.

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Volume Assigned Host Connection
The host connection to which the storage virtualizer volume is assigned. The host connection is a definition in the storage virtualizer that contains the WWPN for the server. The storage virtualizer uses the WWPN to assign volumes to servers.

Storage Virtualizer Volume Name
The name that was assigned to the storage virtualizer volume when it was added to the system.

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Virtualizer MDisk Name
The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage Volume Assigned Host Connection
The host connection to which the storage volume is assigned. The host connection is a definition in the storage system that contains the WWPN for the server. The storage system uses the WWPN to assign volumes to servers.

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Volume	Performance of One Volume
Managed disk	Performance of One Managed Disk
Storage pool	Most Active Pools, Performance of One Pool

Related tasks:

"Running the Storage Resource Relationships Summary (Configurable) report" on page 932
To show the relationships between resources, run the Storage Resource Relationships Summary (Configurable) report. You can show or hide resources, and you can hide resources that are not related to other resources or storage devices.

Running the File System to Volume Relationships report

To show the end-to-end relationship from file systems on servers to volumes on storage pools, run the File System to Volume Relationships.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **File System to Volume Relationships**.
3. Select one or more servers.
4. Select the storage virtualizers and storage systems that the storage pools are in. You can select both storage virtualizers and storage systems, storage virtualizers only, or storage systems only.
5. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

Restriction: To create a report when you select a single resource, the selected resource must be connected to other resources.

Related reference:

"File System to Volume Relationships report"
Shows the end-to-end relationship from file systems on servers to volumes on storage pools. The report shows the usage of space on file systems over a specified period. Use the report to analyze changes in your file system requirements.

File System to Volume Relationships report

Shows the end-to-end relationship from file systems on servers to volumes on storage pools. The report shows the usage of space on file systems over a specified period. Use the report to analyze changes in your file system requirements.

Report output

The information that is provided about the servers, storage virtualizers, storage systems, pools, and volumes is as follows:

Server Name
The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

File System Mount Point
The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be `c:\` or `d:\`. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be `/opt` or `/export/home`.

Server Disk Path
The identifier that specifies a path to a server disk, for example `/dev/hdisk0`.

A server disk might have one or more disk paths. To identify unique server disks, add the Server Disk Serial Number property to the report, and group the server disks by the serial number.

File System Capacity (GiB)
The amount of storage space on the file system of the resource.

File System Available Space (GiB)
The amount of unused storage space in the file system of the resource.

File System Available Space Percentage
The percentage of space in the file system that is unused.

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Volume Name
The name that was assigned to the storage virtualizer volume when it was added to the system.

Storage Virtualizer Volume Is Thin Provisioned
Shows whether a volume or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Virtualizer Volume Allocated Space Percentage
The percentage of space on a volume that is reserved.

This property applies only to volumes on SAN Volume Controller and Storwize V7000 systems.

Storage Virtualizer Volume Used Space Percentage
The percentage of reserved space that is being used on a volume.

This property applies only to volumes on SAN Volume Controller and Storwize V7000 systems.

Storage Virtualizer Pool Name
The name that was assigned to the pool when the pool was added to the system.

Storage Virtualizer Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Storage Volume Is Thin Provisioned
Shows whether a volume or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Volume Allocated Space Percentage
The percentage of space on a volume that is reserved.

This property applies only to volumes on SAN Volume Controller and Storwize V7000 systems.

Storage Volume Used Space Percentage
The percentage of reserved space that is being used on a volume.

This property applies only to volumes on SAN Volume Controller and Storwize V7000 systems.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
pool capacity - allocated space
```

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

Related tasks:

"Running the File System to Volume Relationships report" on page 935
To show the end-to-end relationship from file systems on servers to volumes on storage pools, run the File System to Volume Relationships.

Predefined reports about switches and switch ports

You can run predefined reports to analyze and compare the performance of switches and switch ports.

Predefined reports about switches

You can run predefined reports to analyze and compare the performance of switches.

Running the Performance of One Switch report

To see information about the performance of a switch, run the Performance of One Switch report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switches**.
4. Click **Performance of One Switch**.
5. Select a switch.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of One Switch report"
Shows three charts, and a table of performance metrics for a switch over a time period that you specify. Use the report to analyze the performance of a switch.

Performance of One Switch report

Shows three charts, and a table of performance metrics for a switch over a time period that you specify. Use the report to analyze the performance of a switch.

Charts

The charts show the following performance metrics for the ports on the switch:

- The send frame rates, the receive frame rates, and the total frame rates
- The send data rates, the receive data rates, and the total data rates
- The send bandwidth percentages, the receive bandwidth percentages, and the overall bandwidth percentages

Report output

The following information is provided in the report for the period that you specified:

Switch Name
The user-defined name for the switch, or the name that was assigned to the switch by the vendor.

Switch Location
The physical location of a switch. The location is defined when a switch is added to Tivoli Storage Productivity Center. You can add or edit the location of the switch in the **Properties** pane of the switch in the web-based GUI.

Port Send Frame Rate (frames/s)
The average number of frames per second that are sent through the port.

Port Receive Frame Rate (frames/s)
The average number of frames per second that are received by the port.

Total Port Frame Rate (frames/s)
The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Send Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage
The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Related reports

To open the Most Active Switch Ports report, click the name of the switch in the report table.

Related tasks:

"Running the Performance of One Switch report" on page 937
To see information about the performance of a switch, run the Performance of One Switch report.

Running the Compare Performance of Multiple Switches report

Use the Compare Performance of Multiple Switches report to compare up to four performance metrics for multiple switches. For example, you can compare bandwidth percentages, data rates, and frame rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switches**.
4. Click **Compare Performance of Multiple Switches**.
5. Select the switches, and then click **Next**.
6. Select the category of performance metric that you want to include in the report, and then select the performance metric.
7. Optional: To add more metrics to the report, complete the following steps:
   a. Click **Add Metric**.
   b. Select the category of performance metric, and then select the performance metric.
   c. If you want to remove the last metric that you added, click **Remove Metric**.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
10. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference:

"Compare Performance of Multiple Switches report"
Shows up to four performance metrics for multiple switches over time. Use the report to compare the performance metrics for switches.

Compare Performance of Multiple Switches report

Shows up to four performance metrics for multiple switches over time. Use the report to compare the performance metrics for switches.

Charts

The charts show the performance metrics for the switches that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metrics
Shows the data for each performance metric.

Switch names
Shows the names of the switches, and the performance metric data for each switch.

Related tasks:

"Running the Compare Performance of Multiple Switches report" on page 939
Use the Compare Performance of Multiple Switches report to compare up to four performance metrics for multiple switches. For example, you can compare bandwidth percentages, data rates, and frame rates.

Running the Compare Performance of One Switch over Time Ranges report

To compare a performance metric for one switch in two date ranges, run the Compare Performance of One Switch over Time Ranges report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switches**.
4. Click **Compare Performance of One Switch over Time Ranges**.
5. Select the switch.
6. Select the category of performance metric about which you want a report, and then select the performance metric.
7. Select an interval.
8. Specify the two date ranges or reporting periods that you want to compare. For example, you can compare the performance metric for this week with last week.

   Tip: To specify a start date and an end date for a reporting period, select **Custom Date Range** from the list of reporting periods.
9. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Compare Performance of One Switch over Time Ranges report"
Shows one performance metric on one switch over two time periods. For example, you can use the report to compare the I/O rate for this week and last week.

Compare Performance of One Switch over Time Ranges report

Shows one performance metric on one switch over two time periods. For example, you can use the report to compare the I/O rate for this week and last week.

Charts

The chart shows the performance data of the switch on the performance metric you selected, over the time periods that you selected.

Report output

The following information is provided in the report:

Switch Name
The user-defined name for the switch, or the name that was assigned to the switch by the vendor.

Day in Range
The number of the day in the date range at which the performance metric data was collected, relative to the date range that you specified. The number of days that are shown is determined by the reporting period that was selected for the report. If **Last 30 Days** was selected, then entries for days 1 - 30 are shown on the report.

Hour in Range
The number of the hour in the day at which the performance metric data was collected. For example, `0` is between 12:00:00 a.m. and 12:59:59 a.m., 1 is between 01:00:00 a.m. and 01:59:59 a.m.

Minute in Range
The interval of minutes during which the performance metric data was collected. For example, `0` shows that performance data was collected 0 - 5 minutes after the hour. Similarly, `5` shows that the performance data that was collected 5 - 10 minutes after the hour.

First Date Range
The details of the performance metric data in the first date range. The details are displayed in two columns in the First Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the first date range.

Second Date Range
The details of the performance metric data in the second date range. The details are displayed in two columns in the Second Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the second date range.

Related tasks:

"Running the Compare Performance of One Switch over Time Ranges report" on page 940
To compare a performance metric for one switch in two date ranges, run the Compare Performance of One Switch over Time Ranges report.

Predefined reports about switch ports

You can run predefined reports to analyze and compare the performance of switch ports.

Running the Most Active Switch Ports report

To see which ports are most active on switches, run the Most Active Switch Ports report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switch Ports**.
4. Click **Most Active Switch Ports**.
5. Select the switches, and then select a sort order.
6. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active Switch Ports report"
Shows a chart of the 20 switches that are most active, and details for all switches, for a time period that you specify. Use the report to analyze the performance of switches.

Most Active Switch Ports report

Shows a chart of the 20 switches that are most active, and details for all switches, for a time period that you specify. Use the report to analyze the performance of switches.

Charts

The bar chart shows the most active ports on the switch. By default, ports are sorted by the total frame rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided about the most active ports on switches is as follows:

Switch Name
The user-defined name for the switch, or the name that was assigned to the switch by the vendor.

Switch Blade Slot Number
The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

Switch Port Number
The number of the port on the switch.

Switch Port WWPN
The worldwide port name of the port on the switch.

Total Port Frame Rate (frames/s)
The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Switch	Performance of One Switch
Switch port	Performance of One Switch Port

Related tasks:

"Running the Most Active Switch Ports report" on page 942
To see which ports are most active on switches, run the Most Active Switch Ports report.

Running the Performance of One Switch Port report

To see information about the performance of a switch port, run the Performance of One Switch Port report.

You must select a switch and a switch port. If the port is on a blade, you must select the number of the slot on the switch to which the blade is attached.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switch Ports**.

4. Click **Performance of One Switch Port**.
5. Select a switch.
6. Optional: Select the number of the blade slot. If the port is on a blade, you must select the number of the blade slot.
7. Select a port.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of One Switch Port report"
Shows four charts, and a table of performance metrics for a switch port. Use this report to view multiple performance metric types on the same chart. For example, you can view I/O rate and response time metrics on the same chart.

Performance of One Switch Port report

Shows four charts, and a table of performance metrics for a switch port. Use this report to view multiple performance metric types on the same chart. For example, you can view I/O rate and response time metrics on the same chart.

Charts

The charts show the following performance statistics for the port:

- The send frame rate, the receive frame rate, and the total frame rate
- The send data rate, the receive data rate, and the total data rate
- The send bandwidth percentage, the receive bandwidth percentage, and the overall bandwidth percentage
- The error rates

Report output

The following information is provided in the report for the period that you specified:

Switch Name
The user-defined name for the switch, or the name that was assigned to the switch by the vendor.

Switch Blade Slot Number
The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

Switch Port Number
The number of the port on the switch.

Switch Port WWPN
The worldwide port name of the port on the switch.

Port Send Frame Rate (frames/s)
The average number of frames per second that are sent through the port.

Port Receive Frame Rate (frames/s)
The average number of frames per second that are received by the port.

Total Port Frame Rate (frames/s)
The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Send Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage
The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Error Frame Rate (count/s)
The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

Discarded Frame Rate (count/s)
The average number of frames per second that are discarded because host buffers are unavailable for the port.

Link Failure Rate (count/s)
The average number of link errors per second.

CRC Error Rate (count/s)
The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

Discarded Class 3 Frame Rate (count/s)
The average number of class 3 frames per second that are discarded.

Invalid Transmission Word Rate (count/s)
The average number of bit errors per second that are detected.

Related reports

To open the Most Active Switch Ports report, click the name of the switch in the report table.

Related tasks:

"Running the Performance of One Switch Port report" on page 943
To see information about the performance of a switch port, run the Performance of One Switch Port report.

Running the Compare Performance of Multiple Switch Ports report

Use the Compare Performance of Multiple Switch Ports report to compare up to four performance metrics for multiple switch ports. For example, you can compare bandwidth percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switch Ports**.
4. Click **Compare Performance of Multiple Switch Ports**.
5. Select the switches.
6. Select switch ports. To select switch ports, complete the following steps:
   a. Type a keyword. For example, type `example1` to find the names of switch ports that start with `example1`. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type `a%` to find switch ports with names begin with 'a' or 'A'. Type `%a` to find switch ports with names that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the switch ports on a resource.
   b. Click **Search**. If the search is successful, the switch ports are displayed in the **Results** list.
   c. Select the switch ports that you require from the list, and then click **Insert**.
7. Click **Next**.
8. Select the category of performance metric that you want to include in the report, and then select the performance metric.
9. Optional: To add more metrics to the report, complete the following steps:
   a. Click **Add Metric**.
   b. Select the category of performance metric, and then select the performance metric.
   c. If you want to remove the last metric that you added, click **Remove Metric**.
10. Optional: Select an interval.
11. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

    Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
12. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference:

"Compare Performance of Multiple Switch Ports report"
Shows up to four performance metrics for multiple switch ports over time. Use the report to compare the performance metrics for switch ports.

Compare Performance of Multiple Switch Ports report

Shows up to four performance metrics for multiple switch ports over time. Use the report to compare the performance metrics for switch ports.

Charts

The charts show the performance metrics for the switch ports that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metrics
Shows the data for each performance metric.

Switch names
Shows the names of the switches, and the performance metric data for each switch.

Slot numbers
Shows the slot numbers, and the performance metric data for each slot.

Switch port numbers
Shows the port numbers, and the performance metric data for each switch port number.

Related tasks:

"Running the Compare Performance of Multiple Switch Ports report" on page 946
Use the Compare Performance of Multiple Switch Ports report to compare up to four performance metrics for multiple switch ports. For example, you can compare bandwidth percentages, data rates, and I/O rates.

Running the Compare Performance of One Switch Port over Time Ranges report

To compare a performance metric for one switch port in two date ranges, run the Compare Performance of One Switch Port over Time Ranges report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Fabrics and Switches**.
3. Click **Switch Ports**.
4. Click **Compare Performance of One Switch Port over Time Ranges**.
5. Select the switch that the port is on.
6. If the switch is on a blade, select the slot number of the blade that the port is on.
7. Select the number of the port.
8. Select the category of performance metric about which you want a report, and then select the performance metric.
9. Select an interval.

10. Specify the two date ranges or reporting periods that you want to compare. For example, you can compare the performance metric for this week with last week.

    Tip: To specify a start date and an end date for a reporting period, select **Custom Date Range** from the list of reporting periods.
11. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Compare Performance of One Switch Port over Time Ranges report"
Shows one performance metric on one switch port over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Compare Performance of One Switch Port over Time Ranges report

Shows one performance metric on one switch port over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Charts

The chart shows the performance data of the switch port on the performance metric you selected, over the time periods that you selected.

Report output

The following information is provided in the report:

Switch Name
The user-defined name for the switch, or the name that was assigned to the switch by the vendor.

Switch Blade Slot Number
The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

Switch Port Number
The number of the port on the switch.

Switch Port WWPN
The worldwide port name of the port on the switch.

Day in Range
The number of the day in the date range at which the performance metric data was collected, relative to the date range that you specified. The number of days that are shown is determined by the reporting period that was selected for the report. If **Last 30 Days** was selected, then entries for days 1 - 30 are shown on the report.

Hour in Range
The number of the hour in the day at which the performance metric data was collected. For example, `0` is between 12:00:00 a.m. and 12:59:59 a.m., 1 is between 01:00:00 a.m. and 01:59:59 a.m.

Minute in Range
The interval of minutes during which the performance metric data was collected. For example, `0` shows that performance data was collected 0 - 5 minutes after the hour. Similarly, `5` shows that the performance data that was collected 5 - 10 minutes after the hour.

First Date Range
The details of the performance metric data in the first date range. The details are displayed in two columns in the First Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the first date range.

Second Date Range
The details of the performance metric data in the second date range. The details are displayed in two columns in the Second Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the second date range.

Related tasks:

"Running the Compare Performance of One Switch Port over Time Ranges report" on page 947
To compare a performance metric for one switch port in two date ranges, run the Compare Performance of One Switch Port over Time Ranges report.

Predefined reports about hypervisors

You can run predefined reports to analyze the capacity and performance of hypervisors.

Running the Hypervisors Capacity report

To see information about the capacity of disks on hypervisors, run the Hypervisors Capacity report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Hypervisors**.
3. Click **Hypervisors Capacity**.
4. Optional: Select a hypervisor cluster.
5. Select one or more hypervisors.
6. Optional: Select a storage resources sort order.
7. Optional: Select a sort order.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Hypervisors Capacity report"
Shows a chart of the 20 hypervisors with the most allocated space, and a table with space statistics for all hypervisors. Use the report to review space allocation in your hypervisors.

Hypervisors Capacity report

Shows a chart of the 20 hypervisors with the most allocated space, and a table with space statistics for all hypervisors. Use the report to review space allocation in your hypervisors.

Charts

The bar chart shows the used space and available space on hypervisors.

Report output

For each hypervisor, the following information is provided:

Hypervisor Name
The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

Hypervisor Location
The physical location of a hypervisor. The location is defined when a hypervisor is added to Tivoli Storage Productivity Center. You can add or edit the location of the hypervisor in the **Properties** pane of the hypervisor in the web-based GUI.

Hypervisor Cluster Name
The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

Hypervisor Total Disk Capacity (GiB)
The amount of storage space that is on a hypervisor.

Hypervisor Total Disk Available Space (GiB)
The amount of unused storage space on all disks on a hypervisor.

Hypervisor Total Disk Used Space (GiB)
The amount of storage space that is used on all disks on a hypervisor.

Hypervisor Total Disk Used Space (%)
The amount of storage space that is used on a hypervisor.

Hypervisor Data Store Capacity (GiB)
The amount of storage space that is on the hypervisor data store.

Hypervisor Data Store Available Space (GiB)
The amount of unused storage space on a hypervisor data store.

Hypervisor Data Store Used Space (GiB)
The amount of storage space that is used on a hypervisor data store.

Hypervisor Data Store Used Space (%)
The percentage of storage space that is used on all disks on a hypervisor.

Related reports

To open related reports, click the name of the hypervisor and then click the name of the report. You can open the following related reports:

- Hypervisor Disks Capacity
- Hypervisor Data Stores Capacity

Related tasks:

"Running the Hypervisors Capacity report" on page 949
To see information about the capacity of disks on hypervisors, run the Hypervisors Capacity report.

Running the Most Active Hypervisors report

To see which hypervisors are the most active, run the Most Active Hypervisors report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Hypervisors**.
3. Click **Most Active Hypervisors**.
4. Select the hypervisors, and then select a sort order.
5. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
6. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active Hypervisors report"
Shows a chart of the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load.

Most Active Hypervisors report

Shows a chart of the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load.

Charts

The bar chart shows the most active hypervisors. By default, hypervisors are sorted by the total I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided about the most active hypervisors is as follows:

Hypervisor Name
The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Volume Utilization
The average percentage of time that the volume is busy.

Related reports

To open the Summarized Performance of Volumes by Hypervisor report, click the name of the hypervisor in the report table.

Related tasks:

"Running the Most Active Hypervisors report" on page 951
To see which hypervisors are the most active, run the Most Active Hypervisors report.

Running the Summarized Performance of Volumes by Hypervisor report

To see the load that a particular hypervisor adds to your storage environment, use the Summarized Performance of Volumes by Hypervisor report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Hypervisors**.
3. Click **Summarized Performance of Volumes by Hypervisor**.
4. Select one or more hypervisors.
5. Select the volumes.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Summarized Performance of Volumes by Hypervisor report" on page 953
Shows summarized performance metrics for volumes on a hypervisor. Use the report to identify the aggregate load and the average response time of a particular hypervisor.

Summarized Performance of Volumes by Hypervisor report

Shows summarized performance metrics for volumes on a hypervisor. Use the report to identify the aggregate load and the average response time of a particular hypervisor.

Charts

The charts show the following performance metrics for the volumes on a hypervisor:

- The total I/O rates and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates, the write date rates, and the total data rates
- The read cache percentages and the write cache percentages

Report output

The following information is provided in the report for the period that you specified:

Hypervisor Name
The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

Day, hour, or time
The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation, across all volumes on the hypervisor. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second, across all volumes on the hypervisor. This value includes both sequential and nonsequential operations. The average value for each volume is added to give a total I/O rate for all volumes on the hypervisor.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations, across all volumes on the hypervisor. The average value for each volume is added to give a total data rate for all volumes on the hypervisor.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation, across all volumes on the hypervisor.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation, across all volumes on the hypervisor.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read

operations, across all volumes on the hypervisor. The average value for each volume is added to give a read data rate for all volumes on the hypervisor.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations, across all volumes on the hypervisor. The average value for each volume is added to give a write data rate for all volumes on the hypervisor.

Overall Read I/O Rate (ops/s)
The average number of read operations per second, which is totaled for all volumes on the hypervisor. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second, which is totaled for all volumes on the hypervisor. This value includes both sequential and nonsequential write operations.

Overall Read Cache Hit Percentage
The average percentage of all read operations that find data in the cache, across all volumes on the hypervisor. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The average percentage of all write operations that are handled in the cache, across all volumes on the hypervisor. This value includes both sequential and nonsequential write operations.

Related reports

To open the Performance of Volumes by Hypervisor report, click the name of the hypervisor in the report table.

Related tasks:

"Running the Summarized Performance of Volumes by Hypervisor report" on page 952

To see the load that a particular hypervisor adds to your storage environment, use the Summarized Performance of Volumes by Hypervisor report.

Running the Performance of Volumes by Hypervisor report

To see performance metrics for the volumes on a hypervisor, use the Performance of Volumes by Hypervisor report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Hypervisors**.
3. Click **Performance of Volumes by Hypervisor**.
4. Select one or more hypervisors.
5. Select the volumes.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of Volumes by Hypervisor report"
Shows the performance metrics for volumes on a hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment.

Performance of Volumes by Hypervisor report

Shows the performance metrics for volumes on a hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment.

Charts

The charts show the following performance metrics for the volumes:

- The total overall I/O rates. The chart shows the contribution of each volume to the overall load on the hypervisor.
- The total data rates. The chart shows the contribution of each volume to the overall load on the hypervisor.
- The volume utilization percentages. The utilization percentage of each volume is displayed relative to the utilization percentage of the other volumes.
- The total response times. The response time of each volume is displayed relative to the response time of the other volumes.

Report output

The following information is provided in the report for the period that you specified:

Hypervisor Name
The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Day, hour, or time
The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Volume Utilization
The average percentage of time that the volume is busy.

Related tasks:

"Running the Performance of Volumes by Hypervisor report" on page 954
To see performance metrics for the volumes on a hypervisor, use the Performance of Volumes by Hypervisor report.

Running the Hypervisor Data Stores Capacity report

To see information about the capacity of data stores on hypervisors, run the Hypervisor Data Stores Capacity report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Hypervisors**.
3. Click **Data Stores Capacity**.
4. Optional: Select a hypervisor cluster.
5. Select one or more hypervisors.
6. Optional: Select a sort order.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Hypervisor Data Stores Capacity report"
Shows a chart of the 20 hypervisor data stores with the most allocated space. This report also shows a table with space statistics for all hypervisor data stores. Use the report to review space allocation in your hypervisor data stores.

Hypervisor Data Stores Capacity report

Shows a chart of the 20 hypervisor data stores with the most allocated space. This report also shows a table with space statistics for all hypervisor data stores. Use the report to review space allocation in your hypervisor data stores.

Charts

The bar chart shows the used space and available space for data stores on hypervisors.

Report output

For each hypervisor, the following information is provided:

Hypervisor Cluster Name
The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

Hypervisor Data Store
The path to the VMware ESX data store on the hypervisor, for example `datastore_svc3c`.

Hypervisor Data Store Cluster Name
The name of the cluster, if the data store is a member of a data store cluster. If the data store is not a member of a data store cluster, no information is displayed.

Hypervisor Data Store Type
The type of file system that the hypervisor data store uses. For example, the type can be NFS or VMFS.

Hypervisor System Mount Point
The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be `c:\`or `d:\`. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be `/opt`or `/export/home`.

Hypervisor Data Store Capacity (GiB)
The amount of storage space that is on the hypervisor data store.

Hypervisor Data Store Used Space (GiB)
The amount of storage space that is used on a hypervisor data store.

Hypervisor Data Store Available Space (GiB)
The amount of unused storage space on a hypervisor data store.

Hypervisor Data Store Used Space (%)
The percentage of storage space that is used on all disks on a hypervisor.

Hypervisor File System Capacity (GiB)
The amount of storage space on the file system of the hypervisor.

Hypervisor File System Available Space (GiB)
The amount of unused storage space in the file system of the hypervisor.

Related tasks:

"Running the Hypervisor Data Stores Capacity report" on page 956
To see information about the capacity of data stores on hypervisors, run the Hypervisor Data Stores Capacity report.

Running the Hypervisor Disks Capacity report

To see information about the disk capacity on hypervisors, run the Hypervisor Disks Capacity report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Hypervisors**.

3. Click **Hypervisor Disks Capacity**.
4. Optional: Select a hypervisor cluster.
5. Select one or more hypervisors.
6. Optional: Select a sort order.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Hypervisor Disks Capacity report"
Shows a chart of the 20 hypervisor disks with the most allocated space, and a table with space statistics for all hypervisor disks. Use the report to review space allocation in your hypervisor disks.

Hypervisor Disks Capacity report

Shows a chart of the 20 hypervisor disks with the most allocated space, and a table with space statistics for all hypervisor disks. Use the report to review space allocation in your hypervisor disks.

Charts

The bar chart shows the used space and available space on hypervisors.

Report output

For each hypervisor, the following information is provided:

Hypervisor Cluster Name
The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

Hypervisor Disk Path
The identifier that specifies a path to a hypervisor disk, for example `/dev/hdisk0`.

Tip: A hypervisor disk might have one or more disk paths. To identify unique hypervisor disks, add the Hypervisor Disk Serial Number property to the report, and group the server disks by the serial number.

Hypervisor Disk Capacity (GiB)
The amount of storage space that is on a hypervisor disk.

Hypervisor Disk Available Space (GiB)
The amount of unused storage space that is on a hypervisor disk.

Hypervisor Disk Used Space (GiB)
The amount of storage space that is used on a hypervisor disk.

Hypervisor Disk Used Space (%)
The amount of storage space that is used on a hypervisor disk.

Hypervisor Disk Status
The status of a disk on the hypervisor. Use the status to determine the condition of the disk, and if any actions must be taken. For example, if a

disk has an Error status, take immediate action to correct the problem. If the disk has an Operational status, then it is operating normally and no further action is required.

Hypervisor Disk Vendor
The vendor who supplied the hypervisor disk.

Hypervisor Disk Model
The model name or model number of the hypervisor disk.

Hypervisor Disk Serial Number
The serial number of the hypervisor disk.

Hypervisor Disk Firmware Version
The version number of the firmware that is running on the hypervisor disk.

Related tasks:

"Running the Hypervisor Disks Capacity report" on page 957
To see information about the disk capacity on hypervisors, run the Hypervisor Disks Capacity report.

Predefined reports about servers and file systems on servers

You can run predefined reports to analyze the capacity and performance of servers, file systems, and volumes.

Predefined reports about servers

You can run predefined reports to analyze the capacity and performance of servers and volumes.

Running the Servers Capacity report

To see information about the capacity and space of server disks and file systems, run the Servers Capacity report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Servers**.
3. Click **Servers Capacity**.
4. Optional: Select one or more servers.
5. Optional: Select a sort option. You can order the metrics represented in the chart by disk or by file system.

   Tip: When the report is created, you can change the sort option.
6. Optional: Select a sort order.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Servers Capacity report"
Shows a summary of the available and used disk space, and the available and used file systems, for servers. Use the report to identify servers that are running out of space.

Servers Capacity report

Shows a summary of the available and used disk space, and the available and used file systems, for servers. Use the report to identify servers that are running out of space.

Charts

The bar chart shows the capacity, used space, available space, and percentage of used space for disks on servers. Dots are used to represent the total space that is used by the file system, and the total space that is available to the file system on the server. You can position the cursor over the dot to show the amount of used space or the amount of available space for the file system. The position of the dot that is used to represent the total space that is available to the file system is also used to show the total capacity of the file system.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

For each server, the following information is provided:

Server Name
The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Server Location
The physical location of a server. The location is defined when a server is added to Tivoli Storage Productivity Center. You can add or edit the location of the server in the **Properties** pane of the server in the web-based GUI.

Server Total Disk Capacity (GiB)
The amount of storage space that is on a server.

Server Total Disk Used Space (GiB)
The amount of storage space that is used on all disks on a server.

Server Total Disk Available Space (GiB)
The amount of unused storage space on all disks on a server.

Server Total Disk Used Space (%)
The percentage of used storage space on all disks on the server.

Server Total File System Capacity (GiB)
The amount of storage space on the file system of the resource.

Server Total File System Used Space (GiB)
The amount of used storage space in the file system of the resource.

Server Total File System Available Space (GiB)
The amount of unused storage space in the file system of the resource.

Server Total File System Used Space (%)
The amount of used storage space in the file system of the resource.

Related reports

To open related reports, click the name of the server and then click the name of the report. You can open the following related reports:

- Server Disks Capacity
- File Systems Capacity

Related tasks:

"Running the Servers Capacity report" on page 959
To see information about the capacity and space of server disks and file systems, run the Servers Capacity report.

Running the Most Active Servers report

To see which servers are the most active, run the Most Active Servers report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Servers**.
3. Click **Most Active Servers**.
4. Select the servers, and then select a sort order.
5. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
6. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active Servers report"
Shows a chart of the 20 servers that have the most active volumes. The most active volumes have the greatest aggregate load.

Most Active Servers report

Shows a chart of the 20 servers that have the most active volumes. The most active volumes have the greatest aggregate load.

Charts

The bar chart shows the most active servers. By default, servers are sorted by the total I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided about the most active storage systems is as follows:

Server Name
The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Volume Utilization
The average percentage of time that the volume is busy.

Related reports

To open the Summarized Performance of Volumes by Server report, click the name of the server in the report table.

Related tasks:

"Running the Most Active Servers report" on page 961
To see which servers are the most active, run the Most Active Servers report.

Running the Summarized Performance of Volumes by Server report

To see the load that a particular server adds to your storage environment, use the Summarized Performance of Volumes by Server report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Servers**.
3. Click **Summarized Performance of Volumes by Server**.
4. Select one or more servers.
5. Select the volumes.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Summarized Performance of Volumes by Server report" on page 963
Shows summarized performance metrics for volumes on a server. Use the report to identify the aggregate load and the average response time of a particular server.

Summarized Performance of Volumes by Server report

Shows summarized performance metrics for volumes on a server. Use the report to identify the aggregate load and the average response time of a particular server.

Charts

The charts show the following performance metrics for the volumes on a server:

- The total I/O rates and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates, the write date rates, and the total data rates
- The read cache percentages and the write cache percentages

Report output

The following information is provided in the report for the period that you specified:

Server Name
The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Day, hour, or time
The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation, across all volumes on the server. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second, across all volumes on the server. This value includes both sequential and nonsequential operations. The average value for each volume is added to give a total I/O rate for all volumes on the server.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations, across all volumes on the server. The average value for each volume is added to give a total data rate for all volumes on the server.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation, across all volumes on the server.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation, across all volumes on the server.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations, across all volumes on the server. The average value for each volume is added to give a read data rate for all volumes on the server.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write

operations, across all volumes on the server. The average value for each volume is added to give a write data rate for all volumes on the server.

Overall Read I/O Rate (ops/s)
The average number of read operations per second, which is totaled for all volumes on the server. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second, which is totaled for all volumes on the server. This value includes both sequential and nonsequential write operations.

Overall Read Cache Hit Percentage
The average percentage of all read operations that find data in the cache, across all volumes on the server. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.

Related reports

To open the Performance of Volumes by Server report, click the name of the server in the report table.

Related tasks:

"Running the Summarized Performance of Volumes by Server report" on page 962
To see the load that a particular server adds to your storage environment, use the Summarized Performance of Volumes by Server report.

Running the Performance of Volumes by Server report

To see performance metrics for the volumes on a server, use the Performance of Volumes by Server report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Servers**.
3. Click **Performance of Volumes by Server**.
4. Select a server.
5. Select the volumes.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of Volumes by Server report"
Shows the performance metrics for volumes on a server. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

Performance of Volumes by Server report

Shows the performance metrics for volumes on a server. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

Charts

The charts show the following performance metrics for the volumes:

- The total overall I/O rates. The chart shows the contribution of each volume to the overall load on the server.
- The total data rates. The chart shows the contribution of each volume to the overall load on the server.
- The volume utilization percentages. The utilization percentage of each volume is displayed relative to the utilization percentage of the other volumes.
- The total response times. The response time of each volume is displayed relative to the response time of the other volumes.

Report output

The following information is provided in the report for the period that you specified:

Server Name
The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Day, hour, or time
The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Volume Utilization
The average percentage of time that the volume is busy.

Related tasks:

"Running the Performance of Volumes by Server report" on page 964
To see performance metrics for the volumes on a server, use the Performance of Volumes by Server report.

Running the Server Disks Capacity report

To see information about the capacity of disks on servers, run the Server Disks Capacity report.

To run the report with the default settings, click **Finish**.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Servers**.
3. Click **Server Disks Capacity**.
4. Optional: Select one or more servers.
5. Optional: Select a sort order.
6. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Server Disks Capacity report"
Shows the available and used disk space for all disks, which are grouped by server. You can use the report to identify disks that have no allocated space.

Server Disks Capacity report

Shows the available and used disk space for all disks, which are grouped by server. You can use the report to identify disks that have no allocated space.

Charts

The bar chart shows the capacity, used space, available space, and percentage of used space for disks on servers.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

For each server, the following information is provided:

Server Disk Path
The identifier that specifies a path to a server disk, for example `/dev/hdisk0`.

Server Disk Capacity (GiB)
The amount of storage space that is on a server disk.

Server Disk Available Space (GiB)
The amount of unused storage space that is on a server disk.

Server Disk Used Space (GiB)
The amount of storage space that is used on a server disk.

Server Disk Used Space (%)
The percentage of used storage space on the disk.

Server Disk Status
The condition of the resource, for example normal, warning, or error.

Server Disk Firmware Version
The version number of the firmware that is running on the disk.

Server Disk Model
The model name or model number of the resource.

Server Disk Vendor
The vendor who supplied the resource.

Related tasks:

"Running the Server Disks Capacity report" on page 966
To see information about the capacity of disks on servers, run the Server Disks Capacity report.

Predefined reports about file systems on servers

You can run predefined reports to analyze the capacity of file systems that are grouped by server.

Running the File Systems Capacity report

To see information about the capacity of file systems on servers, run the File Systems Capacity report.

To run the report with the default settings, click **Finish**.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Servers**.
3. Click **File Systems**.
4. Click **File Systems Capacity**.
5. Select one or more servers.
6. Optional: Select a sort order.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"File Systems Capacity report"
Shows the available and used capacity of file systems, which are grouped by server. You can use this report to identify file systems that might be running out of space.

File Systems Capacity report

Shows the available and used capacity of file systems, which are grouped by server. You can use this report to identify file systems that might be running out of space.

Charts

The bar chart shows the capacity, used space, available space, and percentage of used space for file systems on servers.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

For each server, the following information is provided:

File System Mount Point
The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be `c:\` or `d:\`. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be `/opt` or `/export/home`.

File System Type
The type of file system that the resource uses.

File System Used Space (GiB)
The amount of used storage space in the file system of the resource.

File System Available Space (GiB)
The amount of unused storage space in the file system of the resource.

File System Capacity (GiB)
The amount of storage space on the file system of the resource.

File System Used Space (%)
The percentage of used space in the file system on the server.

File System Available Inodes
The number of unused inodes in file systems on the operating system.

File System Used Inodes
The number of used inodes in file systems on the operating system.

Tip: For Microsoft Windows systems, this property is blank.

Related tasks:

"Running the File Systems Capacity report" on page 967
To see information about the capacity of file systems on servers, run the File Systems Capacity report.

Predefined reports about storage systems and components

You can run predefined reports to view capacity details and performance metrics for storage systems and their components. The components include controllers, modules, nodes, disks, I/O groups, managed disks, storage pools, ports, RAID arrays, and volumes.

Predefined reports about storage systems

You can run predefined reports to view capacity details and performance metrics for storage systems.

Running the Storage Systems Capacity report

To see information about the capacity of disks, pools, and volumes on storage systems, run the Storage Systems Capacity report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Storage Systems Capacity**.
4. Optional: Select the configuration of the storage systems.
5. Select one or more storage systems.
6. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Systems Capacity report"
Shows a chart of the 20 storage systems with the most used space, and a table with details of all storage systems. Use the report to review the space allocation in your storage systems.

Storage Systems Capacity report

Shows a chart of the 20 storage systems with the most used space, and a table with details of all storage systems. Use the report to review the space allocation in your storage systems.

Charts

The bar chart shows the used space and available space for pools on storage systems.

Report output

For each storage system, the following information is provided:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage System Type
The type of storage system. For example, the storage system can be an **IBM System Storage DS8800** system, an **IBM System Storage DS8700** system, an **IBM System Storage XIV** system, or another type of storage system.

Storage System Location
The physical location of a storage system. The location is defined when a storage system is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage system in the **Properties** pane of the storage system in the web-based GUI.

Storage System Pool Capacity (GiB)
The amount of storage space in pools that are on the storage system. For an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage System Used Pool Space (GiB)
The amount of space that is in use in all pools on a storage system.

Storage System Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

Storage System Physical Allocation Percentage
The percentage of physical space in storage system pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(allocated pool space ÷ pool capacity) × 100`

For example, the physical allocation percentage is 25% for a 200 GiB storage pool. Therefore, the space that is reserved for volumes is 50 GiB.

Storage System Virtual Allocation Percentage
The percentage of physical space in storage system pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(total volume capacity ÷ pool capacity) × 100`

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pools. If the

allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage System Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

Storage System Shortfall Percentage
The percentage of the remaining unallocated volume space in storage system pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(unallocatable space ÷ (volume space - used volume space)) × 100`

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

Storage System Volume Capacity (GiB)
The amount of space on all volumes on the storage system.

Storage System Number of Volumes
The number of volumes on a resource.

Storage System Number of Disks
The number of physical disks on the storage system. For a Storwize V7000 storage system that is configured as a back-end device, the value is the number of managed disks on the storage system.

Storage System Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage System Unallocated Disk Space (GiB)
The amount of disk space that can be added to a pool.

Storage System Volume Capacity Assigned to MDisks (GiB)
The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

Storage System Disk Capacity (GiB)
The total amount of space that is on a physical disk on the storage system.

Storage System Total Disk Capacity (GiB)
The amount of space on physical disks on a storage system, including spare disks.

Storage System Cache (GiB)
The size of the cache on the storage system. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

Storage System Assigned Volume Space (GiB)
The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage System Volume Capacity for z/OS
The amount of space on all volumes on the storage system that the z/OS operating system can use.

Storage System Real Available Pool Space (GiB)
The amount of unused space in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System Real Configured Pool Space (GiB)
The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System MDisk Available Space (GiB)
The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Related reports

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Storage Systems Historical Capacity
- Disks Capacity
- Volumes Capacity
- Pools Capacity

Related tasks:

"Running the Storage Systems Capacity report" on page 969
To see information about the capacity of disks, pools, and volumes on storage systems, run the Storage Systems Capacity report.

Running the Storage Systems Historical Capacity report

To see changes in the usage of space on storage systems over a specified period, run the Storage Systems Historical Capacity report.

To run the report with the default settings, click **Finish**. By default all of the storage systems are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Storage Systems Historical Capacity**.
4. Optional: Select the configuration of the storage systems.
5. Select one or more storage systems.
6. Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Storage Systems Historical Capacity report" on page 866
Shows the usage of space on storage systems over a specified period. Use the report to analyze changes in your storage system requirements.

Storage Systems Historical Capacity report

Shows the usage of space on storage systems over a specified period. Use the report to analyze changes in your storage system requirements.

Charts

The chart shows the following space statistics for the storage systems:

- The capacity of the pools on the storage systems
- The used space in the pools on the storage systems
- The available space in the pools on the storage systems
- The capacity of the volumes on the storage systems
- The available space on the managed disks on the storage systems
- The capacity of the managed disks on the storage systems

Report output

The following information is included in reports for storage systems:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Date The date that the data was collected for the report.

Storage System Pool Capacity (GiB)
The amount of storage space in pools that are on the storage system. For an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage System Used Pool Space (GiB)
The amount of space that is in use in all pools on a storage system.

Storage System Available Pool Space (GiB)
The amount of unused space in pools that are on the storage system.

Storage System Volume Capacity (GiB)
The amount of space on all volumes on the storage system.

Storage System Number of Volumes
The number of volumes on a resource.

Storage System Number of Disks
The number of physical disks on the storage system. For a Storwize V7000 storage system that is configured as a back-end device, the value is the number of managed disks on the storage system.

Storage System Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage System Unallocated Disk Space (GiB)
The amount of disk space that can be added to a pool.

Storage System Volume Capacity Assigned to MDisks (GiB)
The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

Storage System Disk Capacity (GiB)
The amount of storage space on physical disks in a storage system, excluding spare disks. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

Storage System Total Disk Capacity (GiB)
The amount of space on physical disks on a storage system, including spare disks.

Storage System Assigned Volume Space (GiB)
The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage System Volume Capacity for z/OS (GiB)
The amount of space on all volumes on the storage system that the z/OS operating system can use.

Storage System Real Available Pool Space (GiB)
The amount of unused space in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System Real Configured Pool Space (GiB)
The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System MDisk Capacity (GiB)
The amount of storage space on the managed disk.

Storage System MDisk Available Space (GiB)
The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Related tasks:

"Running the Storage Systems Historical Capacity report" on page 866
To see changes in the usage of space on storage systems over a specified period, run the Storage Systems Historical Capacity report.

Running the Most Active Storage Systems report

To see which storage systems are the most active, run the Most Active Storage Systems report.

To run the report with the default settings, click **Finish**. By default all of the storage systems are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Most Active Storage Systems**.
4. Optional: Select the configuration of the storage systems.
5. Select the storage systems, and then select a sort order.
6. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active Storage Systems report"
Shows a chart of the 20 storage systems that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of storage systems.

Most Active Storage Systems report

Shows a chart of the 20 storage systems that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of storage systems.

Charts

The bar chart shows the most active storage systems. By default, storage systems are sorted by the total I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided about the most active storage systems is as follows:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Location
The physical location of a storage system. The location is defined when a storage system is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage system in the **Properties** pane of the storage system in the web-based GUI.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Related reports

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Most Active Host Connections
- Most Active Nodes
- Most Active IO Groups
- Most Active Disks
- Most Active Managed Disks
- Most Active Pools
- Most Active Ports
- Most Active Volumes
- Most Active Controllers or Modules
- Most Active RAID Arrays

- Performance of One Storage System
- Storage Resource Relationships Summary

Related tasks:

"Running the Most Active Storage Systems report" on page 975
To see which storage systems are the most active, run the Most Active Storage Systems report.

Running the Performance of One Storage System report

To see performance metrics for a storage system, run the Performance of One Storage System report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Performance of One Storage System**.
4. Optional: Select the configuration of the storage system.
5. Select a storage system.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of One Storage System report"
Shows four charts, and a table of performance metrics for a storage system over a time period that you specify. Use the report to analyze the performance of a storage system.

Performance of One Storage System report

Shows four charts, and a table of performance metrics for a storage system over a time period that you specify. Use the report to analyze the performance of a storage system.

Charts

The charts show the following performance metrics for a storage system:

- The total I/O rates, and the total response times
- The read I/O rates, write I/O rates, and response times
- The read data rates, and the write data rates
- The read cache hit percentages, and the write cache hit percentages

Report output

The following information is provided in the report for the period that you specified:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Location
The physical location of a storage system. The location is defined when a storage system is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage system in the **Properties** pane of the storage system in the web-based GUI.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Related reports

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Most Active Host Connections
- Most Active Nodes
- Most Active IO Groups
- Most Active Disks
- Most Active Managed Disks
- Most Active Pools
- Most Active Ports
- Most Active Volumes
- Most Active Controllers or Modules
- Most Active RAID Arrays

Related tasks:

"Running the Performance of One Storage System report" on page 977
To see performance metrics for a storage system, run the Performance of One Storage System report.

Running the Compare Performance of Multiple Storage Systems report

Use the Compare Performance of Multiple Storage Systems report to compare up to four performance metrics on multiple storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Compare Performance of Multiple Storage Systems**.
4. Select the storage systems, and then click **Next**.
5. Select the category of performance metric that you want to include in the report, and then select the performance metric.
6. Optional: To add more metrics to the report, complete the following steps:
   a. Click **Add Metric**.
   b. Select the category of performance metric, and then select the performance metric.
   c. If you want to remove the last metric that you added, click **Remove Metric**.
7. Optional: Select an interval.
8. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
9. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference:

"Compare Performance of Multiple Storage Systems report"
Shows up to four performance metrics for multiple storage systems over time. Use the report to compare the performance metrics for storage systems.

Compare Performance of Multiple Storage Systems report

Shows up to four performance metrics for multiple storage systems over time. Use the report to compare the performance metrics for storage systems.

Charts

The charts show the performance metrics for the storage systems that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metrics
Shows the data for each performance metric.

Storage system names
Shows the names of the storage systems, and the performance metric data for each storage system.

Related tasks:

"Running the Compare Performance of Multiple Storage Systems report" on page 979
Use the Compare Performance of Multiple Storage Systems report to compare up to four performance metrics on multiple storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

Running the Compare Performance of One Storage System over Time Ranges report

To compare a performance metric for one storage system in two date ranges, run the Compare Performance of One Storage System over Time Ranges report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Compare Performance of One Storage System over Time Ranges**.
4. Optional: Select the configuration of the storage systems.
5. Select the storage system.
6. Select the category of performance metric about which you want a report, and then select the performance metric.
7. Select an interval.
8. Specify the two date ranges or reporting periods that you want to compare. For example, you can compare the performance metric for this week with last week.

   Tip: To specify a start date and an end date for a reporting period, select **Custom Date Range** from the list of reporting periods.
9. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Compare Performance of One Storage System over Time Ranges report"
Shows one performance metric on one storage system over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Compare Performance of One Storage System over Time Ranges report

Shows one performance metric on one storage system over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Charts

The chart shows the performance data of the storage system on the performance metric you selected, over the time periods that you selected.

Report output

The following information is provided in the report:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Day in Range
The number of the day in the date range at which the performance metric data was collected, relative to the date range that you specified. The number of days that are shown is determined by the reporting period that was selected for the report. If **Last 30 Days** was selected, then entries for days 1 - 30 are shown on the report.

Hour in Range
The number of the hour in the day at which the performance metric data was collected. For example, `0` is between 12:00:00 a.m. and 12:59:59 a.m., 1 is between 01:00:00 a.m. and 01:59:59 a.m.

Minute in Range
The interval of minutes during which the performance metric data was collected. For example, `0` shows that performance data was collected 0 - 5 minutes after the hour. Similarly, `5` shows that the performance data that was collected 5 - 10 minutes after the hour.

First Date Range
The details of the performance metric data in the first date range. The details are displayed in two columns in the First Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the first date range.

Second Date Range
The details of the performance metric data in the second date range. The details are displayed in two columns in the Second Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was

collected. The column with the performance metric name shows the data for the performance metric for the second date range.

Related tasks:

"Running the Compare Performance of One Storage System over Time Ranges report" on page 980
To compare a performance metric for one storage system in two date ranges, run the Compare Performance of One Storage System over Time Ranges report.

Running the Performance Data Export report

To get a spreadsheet of the performance metrics for one or more resources on a storage system, run the Performance Data Export report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Performance Data Export**.
4. Optional: Select the configuration of the storage systems.
5. Select one or more storage systems.
6. Optional: Click **Next**, and then select the devices for which you want performance metrics in the report.
7. Click **Finish**.

You can open or save the spreadsheet from your browser.

Related reference:

"Performance Data Export report"
Exports the performance metrics for one or more resources on a storage system to a spreadsheet file in Microsoft Excel file format. Use the report if you want to analyze performance data in Excel or to send performance data to IBM.

Performance Data Export report

Exports the performance metrics for one or more resources on a storage system to a spreadsheet file in Microsoft Excel file format. Use the report if you want to analyze performance data in Excel or to send performance data to IBM.

Report output

The spreadsheet contains a worksheet for each resource on the storage system. The worksheet for a resource contains the performance metric data for that resource.

Related tasks:

"Running the Performance Data Export report"
To get a spreadsheet of the performance metrics for one or more resources on a storage system, run the Performance Data Export report.

Predefined reports about controllers, modules, and nodes

You can run predefined reports to analyze and compare the performance of controllers, modules, and nodes.

Running the Most Active Controllers or Modules report

To see which controllers or modules are most active on storage systems, run the Most Active Controllers or Modules report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Controllers, Modules, and Nodes**.
4. Click **Most Active Controllers or Modules**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active Controllers or Modules report"
Shows a chart of the 20 controllers or modules that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of controllers or modules.

Most Active Controllers or Modules report

Shows a chart of the 20 controllers or modules that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of controllers or modules.

Restriction: This report applies only to DS8000, DS5000, XIV systems, and vendor-acquired systems.

Charts

The bar chart shows the most active controllers or modules on storage systems. By default, controllers or modules are sorted by the overall I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided about the most active controllers or modules is as follows:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Controller or Module Name
The name that was assigned to the controller or module when it was added to the system.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Controller or module	Performance of One Controller or Module

Related tasks:

"Running the Most Active Controllers or Modules report" on page 982
To see which controllers or modules are most active on storage systems, run the Most Active Controllers or Modules report.

Running the Performance of One Controller or Module report

To see performance metrics for controllers or modules on a storage system, run the Performance of One Controller or Module report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Controllers, Modules, and Nodes**.
4. Click **Performance of One Controller or Module**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a controller or module.
8. Optional: Select an interval.

9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of One Controller or Module report"
Shows four charts, and a table of performance metrics for a controller or module over a time period. Use the report to analyze the performance of a controller or module.

Performance of One Controller or Module report

Shows four charts, and a table of performance metrics for a controller or module over a time period. Use the report to analyze the performance of a controller or module.

Restriction: This report applies only to DS8000, DS5000, XIV systems, and vendor-acquired systems.

Charts

The charts show the following performance metrics for each controller or module on a storage system:

- The total I/O rates, and the total response times
- The read I/O rates, the write I/O rates, and the response times
- The read data rates, and the write data rates
- The overall cache hit percentage for read operations and write operations, and the utilization percentage for processors

Report output

The following information is provided in the report for the period that you specified:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Controller Name or Module Name
The name that was assigned to the controller or module when it was added to the system.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Related reports

To open the Most Active Controllers or Modules report, click the name of the controller or module in the report table.

Related tasks:

"Running the Performance of One Controller or Module report" on page 984
To see performance metrics for controllers or modules on a storage system, run the Performance of One Controller or Module report.

Running the Most Active Nodes report

To see which nodes are most active on your storage systems, run the Most Active Nodes report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Controllers, Modules, and Nodes**.
4. Click **Most Active Nodes**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active Nodes report"
Shows a chart of the 20 nodes that are most active, and details for all nodes, for a time period that you specify. Use the report to analyze the performance of nodes.

Most Active Nodes report

Shows a chart of the 20 nodes that are most active, and details for all nodes, for a time period that you specify. Use the report to analyze the performance of nodes.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Charts

The chart shows the most active nodes on storage systems. By default, nodes are sorted by the total overall I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided for nodes is as follows:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage I/O Group Name
The name that was assigned to the I/O group when it was added to the system.

Storage Node Name
The name that was assigned to the storage system node when it was added to the system.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
I/O group	Performance of One IO Group
Node	Performance of One Node

Related tasks:

"Running the Most Active Nodes report" on page 986
To see which nodes are most active on your storage systems, run the Most Active Nodes report.

Running the Performance of One Node report

To see information about the performance of a node on a storage system, run the Performance of One Node report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Controllers, Modules, and Nodes**.
4. Click **Performance of One Node**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a node.
8. Optional: Select an interval.

9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of One Node report"
Shows four charts, and a table of performance metrics for a node over a time period. Use the report to analyze the performance of a node.

Performance of One Node report

Shows four charts, and a table of performance metrics for a node over a time period. Use the report to analyze the performance of a node.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Charts

The charts show the following performance statistics for a node on a storage system:

- The total overall I/O rates, and the total response times
- The overall I/O rates for read and write operations, the response times for read and write operations
- The data rates for read and write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for processors

Report output

The following information is provided in the report for the period that you specified:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage I/O Group Name
The name that was assigned to the I/O group when it was added to the system.

Storage Node Name
The name that was assigned to the storage system node when it was added to the system.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Most Active Nodes
I/O group	Performance of One IO Group

Related tasks:

"Running the Performance of One Node report" on page 988
To see information about the performance of a node on a storage system, run the Performance of One Node report.

Running the Compare Performance of Multiple Controllers and Modules report

Use this report to compare up to four performance metrics for multiple controllers and modules that are on storage systems. For example, you can compare data rates, I/O rates, and response times.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Controllers, Modules, and Nodes**.
4. Click **Compare Performance of Multiple Controllers and Modules**.
5. Select the storage systems, and then click **Refresh**.
6. Select the controllers and modules, and then click **Next**.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
   a. Click **Add Metric**.
   b. Select the category of performance metric, and then select the performance metric.
   c. If you want to remove the last metric that you added, click **Remove Metric**.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

    Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference:

"Compare Performance of Multiple Controllers and Modules report" on page 992
Shows up to four performance metrics for multiple controllers or modules over time. Use the report to compare the performance metrics for controllers or modules that are on storage systems.

Compare Performance of Multiple Controllers and Modules report

Shows up to four performance metrics for multiple controllers or modules over time. Use the report to compare the performance metrics for controllers or modules that are on storage systems.

Restriction: This report applies only to DS8000, DS5000, XIV systems, and vendor-acquired systems.

Charts

The charts show the performance metrics for the controllers and modules that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metric
Shows the time at which the performance metric data was collected.

Storage system names
Shows the names of the storage systems, and the performance metric data for each storage system.

Controller and modules
Shows the names of the controllers and modules, and the performance metric data for each controller and module.

Related tasks:

"Running the Compare Performance of Multiple Controllers and Modules report" on page 991
Use this report to compare up to four performance metrics for multiple controllers and modules that are on storage systems. For example, you can compare data rates, I/O rates, and response times.

Running the Compare Performance of Multiple Nodes report

Use the Compare Performance of Multiple Controllers and Modules report to compare up to four performance metrics on multiple nodes that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Controllers, Modules, and Nodes**.
4. Click **Compare Performance of Multiple Nodes**.
5. Select the storage systems, and then click **Refresh**.
6. Select the nodes, and then click **Next**.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
   a. Click **Add Metric**.
   b. Select the category of performance metric, and then select the performance metric.

c. If you want to remove the last metric that you added, click **Remove Metric**.

9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

    Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference:

"Compare Performance of Multiple Nodes report"
Shows up to four performance metrics for multiple nodes over time. Use the report to compare the performance metrics for nodes that are on storage systems.

Compare Performance of Multiple Nodes report

Shows up to four performance metrics for multiple nodes over time. Use the report to compare the performance metrics for nodes that are on storage systems.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Charts

The charts show the performance metrics for the nodes that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metrics
Shows the data for each performance metric.

Storage system names
Shows the names of the storage systems, and the performance metric data for each storage system.

Nodes Shows the names of the nodes, and the performance metric data for each node.

Related tasks:

"Running the Compare Performance of Multiple Nodes report" on page 992
Use the Compare Performance of Multiple Controllers and Modules report to compare up to four performance metrics on multiple nodes that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

Predefined reports about disks

You can run predefined reports to analyze the capacity of disks and performance of local disks.

Running the Disks Capacity report

Run the Disks Capacity report to see information about the capacity of disks on storage systems.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Disks**.
4. Click **Disks Capacity**.
5. Optional: Select the configuration of the storage systems.
6. Select one or more storage systems.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Disks Capacity report"
Shows a chart of the disk classes with the most allocated space, and a table with space statistics for all disks by class. Use the report to review space allocation in your disks.

Disks Capacity report

Shows a chart of the disk classes with the most allocated space, and a table with space statistics for all disks by class. Use the report to review space allocation in your disks.

Charts

The bar chart shows the capacity of disks, grouped by storage class, on storage systems.

Report output

For each disk on a storage system, the following information is provided:

Storage Disk Class
The storage technology of the disk. This property applies only to DS8000, DS5000, and DS4000 storage systems. This property can contain the following values:

FC The disk is a Fibre Channel disk.

Nearline ATA
The disk is a nearline Advanced Technology Attachment disk.

Nearline SAS
The disk is a nearline serial-attached SCSI disk.

SAS The disk is a serial-attached SCSI disk.

SATA The disk is a Serial Advanced Technology Attachment disk.

SSD The disk is a solid-state drive.

Unavailable
The disk is on a DS5000 or DS4000 series system, and is not a solid-state drive.

Storage Disk Name
The name that was assigned to the storage disk when it was added to the system.

Storage System Type
The type of storage system. For example, the storage system can be an **IBM System Storage DS8800** system, an **IBM System Storage DS8700** system, an **IBM System Storage XIV** system, or another type of storage system.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Disk Tag
A number that identifies the array to which a disk belongs. This property applies only to DS8000, DS6000, and TotalStorage Enterprise Storage Server storage systems.

Storage Disk Vendor
The vendor who supplied the resource.

Storage Disk Model
The model name or model number of the resource.

Storage Disk Serial Number
The serial number of the resource.

Storage Disk Firmware Version
The version number of the firmware that is running on the disk.

Storage Disk Capacity (GiB)
The amount of storage space that is on a physical disk on the storage system.

Storage Disk Speed (RPM)
The speed of the disk.

Related tasks:

"Running the Disks Capacity report" on page 994
Run the Disks Capacity report to see information about the capacity of disks on storage systems.

Running the Most Active Disks report

To see which local disks that are on storage systems are the most active, use the Most Active Disks report. This report shows information only for SAN Volume Controller and Storwize V7000 storage systems.

To run the report with the default settings, click **Finish**. By default all of the storage systems are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Disks**.
4. Click **Most Active Disks**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active Disks report"
Shows a chart of the 20 local disks that are most active, and details for all local disks, for a time period that you specify. Use the report to analyze the performance of local disks.

Most Active Disks report

Shows a chart of the 20 local disks that are most active, and details for all local disks, for a time period that you specify. Use the report to analyze the performance of local disks.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Charts

The bar chart shows the most active local disks that are on storage systems. By default, local disks are sorted by the total I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided about the most active local disks that are on storage systems is as follows:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage MDisk Name
The name that was assigned to the managed disk when it was added to the system.

Storage Local Disk Name
The name that was assigned to the local disk on the storage system when the local disk was created on the system.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage Local Disk Storage Class
The storage technology of the local disk on a storage system. For example, the storage class can be a serial-attached SCSI (SAS) or solid-state drive (SSD).

Storage Local Disk Speed (RPM)
The speed of the local disk on a storage system.

Storage Local Disk Capacity (GiB)
The amount of storage space on the local disk.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Managed disk	Performance of One Managed Disk
Disk	Performance of One Disk

Related tasks:

"Running the Most Active Disks report" on page 995
To see which local disks that are on storage systems are the most active, use the Most Active Disks report. This report shows information only for SAN Volume Controller and Storwize V7000 storage systems.

Running the Performance of One Disk report

To see performance metrics for a local disk on a storage system, use the Performance of One Disk report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Disks**.
4. Click **Performance of One Disk**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a local disk.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of One Disk report"
Shows four charts, and a table of performance metrics for a local disk over a time period. Use the report to analyze the performance of a local disk.

Performance of One Disk report

Shows four charts, and a table of performance metrics for a local disk over a time period. Use the report to analyze the performance of a local disk.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Charts

The charts show the following performance metrics for a local disk on a storage system:

- The total back-end I/O rates and the overall back-end response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates, the write date rates, and the total date rates
- The read queue times, the write queue times, and the overall queue times

Report output

The following information is provided in the report for the period that you specified:

Storage System Name
A user-defined name of the storage system. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage MDisk Name
The name that was assigned to the managed disk when it was added to the system.

Storage Local Disk Name
The name that was assigned to the local disk on the storage system when the local disk was created on the system.

Storage Local Disk Storage Class
The storage technology of the local disk on a storage system. For example, the storage class can be a serial-attached SCSI (SAS) or solid-state drive (SSD).

Storage Local Disk Speed (RPM)
The speed of the local disk on a storage system.

Storage Local Disk Capacity (GiB)
The amount of storage space on the local disk.

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Related reports

To open the Most Active Disks report, click the name of the storage system in the report table.

Related tasks:

"Running the Performance of One Disk report" on page 998
To see performance metrics for a local disk on a storage system, use the Performance of One Disk report.

Running the Compare Performance of Multiple Disks report

Use the Compare Performance of Multiple Disks report to compare up to four performance metrics for multiple local disks that are on storage systems. You can view reports about local disks on SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, and Storwize V3700 systems. For example, you can compare cache percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Disks**.
4. Click **Compare Performance of Multiple Disks**.
5. Select the storage systems, and then click **Refresh**.
6. Select the local disks, and then click **Next**.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
   a. Click **Add Metric**.
   b. Select the category of performance metric, and then select the performance metric.

c. If you want to remove the last metric that you added, click **Remove Metric**.

9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

    Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference:

"Compare Performance of Multiple Disks report"
Shows up to four performance metrics for multiple local disks over time. Use the report to compare the performance metrics for local disks that are on storage systems.

Compare Performance of Multiple Disks report

Shows up to four performance metrics for multiple local disks over time. Use the report to compare the performance metrics for local disks that are on storage systems.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Charts

The charts show the performance metrics for the local disks that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metric
Shows the time at which the performance metric data was collected.

Storage system names
Shows the names of the storage systems, and the performance metric data for each storage system.

Disks Shows the names of the disks, and the performance metric data for each disk.

Related tasks:

"Running the Compare Performance of Multiple Disks report" on page 1000
Use the Compare Performance of Multiple Disks report to compare up to four performance metrics for multiple local disks that are on storage systems. You can view reports about local disks on SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, and Storwize V3700 systems. For example, you can compare cache percentages, data rates, and I/O rates.

Predefined reports about host connections

You can run predefined reports to analyze and compare the performance of volumes on host connections.

Running the Most Active Host Connections report

To see which host connections are the most active, run the Most Active Host Connections report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Host Connections**.
4. Click **Most Active Host Connections**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active Host Connections report"
Shows a chart of the 20 host connections that have the most active volumes. The most active volumes have the greatest aggregate load.

Most Active Host Connections report

Shows a chart of the 20 host connections that have the most active volumes. The most active volumes have the greatest aggregate load.

Chart

The bar chart shows the most active host connections. By default, host connections are sorted by the total I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided about the most active host connections is as follows:

Storage Host Connection Name
The user-defined name that describes the server or storage device that is assigned to a volume.

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

XIV Cluster
The name of the cluster, if any, as defined on the XIV system. If no cluster is defined for the host connection, this field is blank.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Host connection	Summarized Performance of Volumes by Host Connection
Storage system	Performance of One Storage System

Related tasks:

"Running the Most Active Host Connections report" on page 1002
To see which host connections are the most active, run the Most Active Host Connections report.

Running the Summarized Performance of Volumes by Host Connection report

To see the load that a particular server adds to your storage environment, use the Summarized Performance of Volumes by Host Connection report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.

2. Click **Storage Systems**.
3. Click **Host Connections**.
4. Click **Summarized Performance of Volumes by Host Connection**.
5. Select a host connection. To select a host connection, complete the following steps:
   a. Type a keyword. For example, type `example1` to find the names of host connections that start with `example1`. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type `a%` to find host connections that begin with 'a' or 'A'. Type `%a` to find host connections that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the host connections on a resource.
   b. Click **Search**. The report searches the user-defined names for the host connections. The report does not search the server names that are defined when a server is probed. If the search is successful, the host connections are displayed in the **Choice** list.
   c. From the list, select a host connection, and then click **Refresh**.
6. Select the volumes.
7. Optional: Select an interval.
8. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
9. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Summarized Performance of Volumes by Host Connection report"
Shows summarized performance metrics for volumes on a host connection. Use the report to identify the aggregate load and the average response time of a particular server.

Summarized Performance of Volumes by Host Connection report

Shows summarized performance metrics for volumes on a host connection. Use the report to identify the aggregate load and the average response time of a particular server.

Charts

The charts show the following performance statistics for the volumes on a host connection:

- The total I/O rates and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates, the write date rates, and the total data rates
- The read cache percentages and the write cache percentages

Report output

The following information is provided in the report for the period that you specified:

Storage Host Connection Name
The user-defined name that describes the server or storage device that is assigned to a volume.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation, across all volumes on the server. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second, across all volumes on the server. This value includes both sequential and nonsequential operations. The average value for each volume is added to give a total I/O rate for all volumes on the server.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations, across all volumes on the server. The average value for each volume is added to give a total data rate for all volumes on the server.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation, across all volumes on the server.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation, across all volumes on the server.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations, across all volumes on the server. The average value for each volume is added to give a read data rate for all volumes on the server.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations, across all volumes on the server. The average value for each volume is added to give a write data rate for all volumes on the server.

Overall Read I/O Rate (ops/s)
The average number of read operations per second, which is totaled for all volumes on the server. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second, which is totaled for all volumes on the server. This value includes both sequential and nonsequential write operations.

Overall Read Cache Hit Percentage
The average percentage of all read operations that find data in the cache, across all volumes on the server. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.

Related reports

To open the Performance of Volumes by Host Connection report, click the name of the host connection in the report table.

Related tasks:

"Running the Summarized Performance of Volumes by Host Connection report" on page 1003
To see the load that a particular server adds to your storage environment, use the Summarized Performance of Volumes by Host Connection report.

Running the Performance of Volumes by Host Connection report

To see performance metrics for the volumes on a server, use the Performance of Volumes by Host Connection report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Host Connections**.
4. Click **Performance of Volumes by Host Connection**.
5. Select a host connection. To select a host connection, complete the following steps:
   a. Type a keyword. For example, type `example1` to find the names of host connections that start with `example1`. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type `a%` to find host connections that begin with 'a' or 'A'. Type `%a` to find host connections that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the host connections on a resource.
   b. Click **Search**. The report searches the user-defined names for the host connections. The report does not search the server names that are defined when a server is probed. If the search is successful, the host connections are displayed in the **Choice** list.
   c. From the list, select a host connection, and then click **Refresh**.
6. Select the volumes.
7. Optional: Select an interval.
8. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
9. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of Volumes by Host Connection report" on page 1007
Shows the performance metrics for volumes on a host connection. Use the report to identify the contribution of individual volumes to the load of a particular server

in your storage environment.

Performance of Volumes by Host Connection report

Shows the performance metrics for volumes on a host connection. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

Charts

The charts show the following performance metrics for the volumes:

- The total overall I/O rates. This is a stacked area chart that shows the contribution of each volume to the overall load on the server.
- The total data rates. This is a stacked area chart that shows the contribution of each volume to the overall load on the server.
- The volume utilization percentages. The utilization percentage of each volume is displayed relative to the utilization percentage of the other volumes.
- The total response times. The response time of each volume is displayed relative to the response time of the other volumes.

Report output

The following information is provided in the report for the period that you specified:

Storage Host Connection Name
The user-defined name that describes the server or storage device that is assigned to a volume.

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Volume Utilization
The average percentage of time that the volume is busy.

Related tasks:

"Running the Performance of Volumes by Host Connection report" on page 1006
To see performance metrics for the volumes on a server, use the Performance of Volumes by Host Connection report.

Predefined reports about I/O groups

You can run predefined reports to analyze and compare the performance of I/O groups.

Running the Most Active IO Groups report

To see which I/O groups are most active on storage systems, run the Most Active IO Groups report.

To run the report with the default settings, click **Finish**. By default all of the storage systems are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **IO Groups**.
4. Click **Most Active IO Groups**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active IO Groups report"
Shows a chart of the 20 I/O groups that are most active, and details for all I/O groups, for a time period that you specify. Use the report to analyze the performance of I/O groups.

Most Active IO Groups report

Shows a chart of the 20 I/O groups that are most active, and details for all I/O groups, for a time period that you specify. Use the report to analyze the performance of I/O groups.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Charts

The chart shows the most active I/O groups on storage systems. By default, I/O groups are sorted by the total overall I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided for I/O groups is as follows:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage I/O Group Name
The name that was assigned to the I/O group when it was added to the system.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
I/O group	Performance of One IO Group

Related tasks:

"Running the Most Active IO Groups report" on page 1008
To see which I/O groups are most active on storage systems, run the Most Active IO Groups report.

Running the Performance of One IO Group report

To see information about the performance of an I/O group on a storage system, run the Performance of One IO Group report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **IO Groups**.
4. Click **Performance of One IO Group**.

5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select an I/O group.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of One IO Group report"
Shows four charts, and a table of performance metrics for an I/O group over a time period. Use the report to analyze the performance of an I/O group.

Performance of One IO Group report

Shows four charts, and a table of performance metrics for an I/O group over a time period. Use the report to analyze the performance of an I/O group.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Charts

The charts show the following performance metrics for an I/O group on a storage system:

- The total overall I/O rates, and the total response times
- The overall I/O rates for read and write operations, and the response times for read and write operations
- The data rates for read and write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for processors

Report output

The following information is provided in the report for the period that you specified:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage I/O Group Name
The name that was assigned to the I/O group when it was added to the system.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Related reports

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Performance of One Storage System
- Most Active IO Groups

Related tasks:

"Running the Performance of One IO Group report" on page 1009
To see information about the performance of an I/O group on a storage system, run the Performance of One IO Group report.

Running the Compare Performance of Multiple IO Groups report

Use the Compare Performance of Multiple IO Groups report to compare up to four performance metrics for multiple I/O groups that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **IO Groups**.
4. Click **Compare Performance of Multiple IO Groups**.
5. Select the storage systems, and then click **Refresh**.
6. Select the I/O groups, and then click **Next**.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
   a. Click **Add Metric**.
   b. Select the category of performance metric, and then select the performance metric.
   c. If you want to remove the last metric that you added, click **Remove Metric**.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

    Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference:

"Compare Performance of Multiple IO Groups report"
Shows up to four performance metrics for multiple I/O groups over time. Use the report to compare the performance metrics for I/O groups that are on storage systems.

Compare Performance of Multiple IO Groups report

Shows up to four performance metrics for multiple I/O groups over time. Use the report to compare the performance metrics for I/O groups that are on storage systems.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Charts

The charts show the performance metrics for the I/O groups that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metric
Shows the time at which the performance metric data was collected.

Storage system names
Shows the names of the storage systems, and the performance metric data for each storage system.

I/O groups
Shows the names of the I/O groups, and the performance metric data for each I/O group.

Related tasks:

"Running the Compare Performance of Multiple IO Groups report" on page 1012
Use the Compare Performance of Multiple IO Groups report to compare up to four performance metrics for multiple I/O groups that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

Predefined reports about managed disks

You can run predefined reports about to analyze and compare the capacity and performance of managed disks.

Running the Managed Disks Capacity report

To see information about used space and available space on managed disks (MDisks) on storage virtualizers, run the Managed Disks Capacity report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Managed Disks**.
4. Click **Managed Disks Capacity**.
5. Optional: Select the configuration of the storage systems.
6. Select one or more storage systems.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Managed Disks Capacity report" on page 825
Shows a chart of the 20 managed disks with the most available space, and a table with space statistics for all managed disks. Use the report to review space availability in your managed disks.

Managed Disks Capacity report

Shows a chart of the 20 managed disks with the most available space, and a table with space statistics for all managed disks. Use the report to review space availability in your managed disks.

Restriction: This report applies only to SAN Volume Controller, Storwize, and Hitachi Universal Storage Platform V systems.

Charts

The bar chart shows the used space and the available space for managed disk on storage systems.

Report output

For each managed disk on a storage system, the following information is provided:

Storage System MDisk Name
The name that was assigned to the managed disk on a storage system when it was added to the system.

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage System MDisk Capacity (GiB)
The amount of storage space on the managed disk.

Storage System MDisk Used Space (GiB)
The amount of used space on the managed disk.

Storage System MDisk Available Space (GiB)
The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage MDisk RAID Level
The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

Storage System MDisk Type
The type of managed disk on a storage system. For example, the disk on a storage system can be a local managed disk.

Storage System MDisk Mode
The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

Array Extents are provided from local disks.

Managed
Extents are provided from a back-end storage volume.

Unmanaged
The managed disk is not used in the system.

Storage System MDisk Strip Size (KB)
The RAID strip size on a managed disk on a storage system.

Storage System MDisk Is Balanced
Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is `Yes`, the LUNs are balanced.

Storage System MDisk Fast Write State
Shows whether the cache for a volume on a disk that is managed by a storage system is empty, contains data, or is corrupted.

Storage System MDisk Write Verify
Shows whether all write operations on a managed disk on a storage system are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is `Yes`, all write operations are verified by a follow-up read operation.

Related tasks:

"Running the Managed Disks Capacity report" on page 824
To see information about used space and available space on managed disks (MDisks) on storage virtualizers, run the Managed Disks Capacity report.

Running the Most Active Managed Disks report

To see which managed disks are most active on storage systems, use the Most Active Managed Disks.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Managed Disks**.
4. Click **Most Active Managed Disks**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active Managed Disks report"
Shows a chart of the 20 managed disks that are most active, and details for all managed disks, for a time period that you specify. Use the report to analyze the performance of managed disks.

Most Active Managed Disks report

Shows a chart of the 20 managed disks that are most active, and details for all managed disks, for a time period that you specify. Use the report to analyze the performance of managed disks.

Restriction: This report applies only to SAN Volume Controller, Storwize, and Hitachi Universal Storage Platform V systems.

Charts

The bar chart shows the most active managed disks that are on storage systems. By default, managed disks are sorted by the I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided about the most active managed disks is as follows:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage MDisk Name
The name that was assigned to the managed disk when it was added to the system.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Storage pool	Performance of One Pool
Managed disk	Performance of One Managed Disk, Storage Resource Relationships Summary

Related tasks:

"Running the Most Active Managed Disks report" on page 1015
To see which managed disks are most active on storage systems, use the Most Active Managed Disks.

Running the Performance of One Managed Disk

To see performance metrics for managed disks on a storage system, use the Performance of One Managed Disk.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Managed Disks**.
4. Click **Performance of One Managed Disk**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a managed disk.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of One Managed Disk report"
Shows four charts, and a table of performance metrics for a managed disk over a time period. Use the report to analyze the performance of a managed disk.

Performance of One Managed Disk report

Shows four charts, and a table of performance metrics for a managed disk over a time period. Use the report to analyze the performance of a managed disk.

Restriction: This report applies only to SAN Volume Controller, Storwize, and Hitachi Universal Storage Platform V systems.

Charts

The charts show the following performance statistics for a managed disk on a storage system:

- The total I/O rates and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates and the write date rates
- The read cache percentages, the write cache percentages, and the volume utilization percentages

Report output

The following information is provided in the report for the period that you specified:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage MDisk Name
The name that was assigned to the managed disk when it was added to the system.

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Most Active Managed Disks
Storage pool	Performance of One Pool

Related tasks:

"Running the Performance of One Managed Disk" on page 1017
To see performance metrics for managed disks on a storage system, use the Performance of One Managed Disk.

Running the Compare Performance of Multiple Managed Disks report

Use the Compare Performance of Multiple Managed Disks report to compare up to four performance metrics for multiple managed disks that are on storage systems. For example, you can compare data rates, I/O rates, and response times.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.

3. Click **Managed Disks**.
4. Click **Compare Performance of Multiple Managed Disks**.
5. Select the storage systems.
6. Select managed disks. To select managed disks, complete the following steps:
   a. Type a keyword. For example, type `mdisk1` to find the names of managed disks that start with `mdisk1`. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type `a%` to find managed disks that begin with 'a' or 'A'. Type `%a` to find managed disks that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the managed disks on a resource.
   b. Click **Search**. If the search is successful, the managed disks are displayed in the **Results** list.
   c. Select the managed disks that you require from the list, and then click **Insert**.
7. Click **Next**.
8. Select the category of performance metric that you want to include in the report, and then select the performance metric.
9. Optional: To add more metrics to the report, complete the following steps:
   a. Click **Add Metric**.
   b. Select the category of performance metric, and then select the performance metric.
   c. If you want to remove the last metric that you added, click **Remove Metric**.
10. Optional: Select an interval.
11. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

    Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
12. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference:

"Compare Performance of Multiple Managed Disks report"
Shows up to four performance metrics for multiple managed disks over time. Use the report to compare the performance metrics for managed disks that are on storage systems.

Compare Performance of Multiple Managed Disks report

Shows up to four performance metrics for multiple managed disks over time. Use the report to compare the performance metrics for managed disks that are on storage systems.

Restriction: This report applies only to SAN Volume Controller, Storwize, and Hitachi Universal Storage Platform V systems.

Charts

The charts show the performance metrics for the managed disks that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metric
Shows the time at which the performance metric data was collected.

Storage system names
Shows the names of the storage systems, and the performance metric data for each storage system.

Managed disks
Shows the names of the managed disks, and the performance metric data for each managed disk.

Related tasks:

"Running the Compare Performance of Multiple Managed Disks report" on page 1019
Use the Compare Performance of Multiple Managed Disks report to compare up to four performance metrics for multiple managed disks that are on storage systems. For example, you can compare data rates, I/O rates, and response times.

Predefined reports about storage pools

You can run predefined reports to analyze and compare the capacity and performance of storage pools.

Running the Pools Capacity report

To see information about the used and available space of pools on storage systems, run the Pools Capacity report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Pools**.
4. Click **Pools Capacity**.
5. Optional: Select the configuration of the storage systems.
6. Optional: Select one or more storage systems.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Pools Capacity report" on page 1022
Shows a chart of the 20 pools with the most allocated space, and a table with space statistics for all pools. Use the report to review space allocation in your pools.

Pools Capacity report

Shows a chart of the 20 pools with the most allocated space, and a table with space statistics for all pools. Use the report to review space allocation in your pools.

Charts

The bar chart shows the used space and available space for pools on storage systems.

Report output

For each pool on a storage system, the following information is provided:

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage Pool Capacity (GiB)
The total amount of storage space in a pool.

Storage Pool Used Space (GiB)
The amount of storage space that is used in the pool.

Storage Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
pool capacity - allocated space
```

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

Storage Pool Physical Allocation Percentage
The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

```
(allocated space ÷ pool capacity) × 100
```

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Pool Virtual Allocation Percentage
The percentage of physical space in a pool that is committed to the total

virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

`(total volume capacity ÷ pool capacity) × 100`

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pool is 30 GiB.

This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, you can use the value for Storage Pool Shortfall Percentage to estimate how critical the shortage of space is for a pool.

This value is only available for pools with thin-provisioned volumes.

Storage Pool Shortfall Percentage
The percentage of the remaining unallocated volume space in a pool that is not available to be allocated.

The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(unallocatable space ÷ (virtual capacity − allocated space)) × 100`

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

Storage Pool Assigned Volume Space (GiB)
The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Pool Unassigned Volume Space (GiB)
The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Pool Total Volume Capacity (GiB)
The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Pool Allocated Space (GiB)
The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

Storage Pool Used Volume Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Pool Real Configured Space (GiB)
The amount of storage space that is in an XIV system pool that has an associated soft size. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage Pool Real Available Space (GiB)
The amount of unused space that is in an XIV system pool that has an associated soft size. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage Pool RAID Level
The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Pool Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Pool Is Solid State
Shows whether there are solid-state drives in the pool. This property can contain the following values:

Mixed The pool contains both hard disk drives and solid-state drives.

Non solid state
The pool contains no solid-state drives.

Solid state
The pool contains a least one solid-state drive.

Storage Pool Number of MDisks
The number of managed disks in a storage pool.

Storage Pool Extent Size (MiB)
The size of the extent that was specified when a pool was created. Smaller extent sizes limit the maximum size of the volumes that can be created in a pool. Smaller extent sizes minimize the amount of potentially wasted space per volume.

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage Pool Warning Level
The percentage of used capacity of the storage pool at which a warning is generated.

Storage Pool Repository Capacity (GiB)
The amount of space on all extents in the repository of a pool. This space can be used to allocate track space-efficient volumes. This attribute applies only to DS8000 storage systems.

Storage Pool Used Repository Space (GiB)
The amount of space on all extents in the repository of a pool that are allocated. This space can be used to allocate track space-efficient volumes. This attribute applies only to the DS8000 storage systems.

Report links

To open related reports, click the name of the storage pool and then click the name of the report. You can open the following related reports:

- Pools Historical Capacity
- Volumes Capacity

Related tasks:

"Running the Pools Capacity report" on page 1021
To see information about the used and available space of pools on storage systems, run the Pools Capacity report.

Running the Pools Historical Capacity report

To see how space in pools on storage systems is used, run the Pools Historical Capacity report.

You can see how space in a pool on a storage system, or in pools in a storage resource group, are used.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Pools**.

4. Click **Pools Historical Capacity**.
5. Select the type of historical space report that you want to create. You can select either of the following options:
   - **Of a Pool in a Storage System**
     a. Click **Next**.
     b. Optional: Select the configuration of the storage system.
     c. Select a storage system, and then select a pool.
   - **Of Pools in a Storage Resource Group**
     a. Click **Next**.
     b. Optional: Select the configuration of the storage system.
     c. Select a storage resource group.
6. Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Pools Historical Capacity report" on page 869
Shows the usage of space on storage pools over a specified period. Use the report to analyze changes in your storage pool requirements.

Pools Historical Capacity report

Shows the usage of space on storage pools over a specified period. Use the report to analyze changes in your storage pool requirements.

Charts

The chart shows the following space statistics for the pool on a storage system:

- The capacity of the pool
- The used space in the pool
- The available space in the pool
- The virtual volume capacity of the pool
- The used volume space in the pool
- The allocated space in the pool

The chart shows the following space statistics for the pool in the storage resource group:

- The capacity of the pools
- The used space in the pools
- The available space in the pools
- The virtual volume capacity of the pools
- The used volume space in the pools
- The allocated space in the pool

Report output

The following information is included in reports for a pool on a storage system or for pools in a storage resource group:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Date The date that the data was collected for the report.

Storage System Pool Capacity (GiB)
The amount of storage space in pools that are on the storage system. For an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
pool capacity - allocated space
```

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

Storage Pool Allocated Space (GiB)
The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

Storage Pool Used Space (GiB)
The amount of storage space that is used in the pool.

Storage Pool Total Volume Capacity (GiB)
The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Pool Used Volume Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Pool Assigned Volume Space (GiB)
The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Pool Unassigned Volume Space (GiB)
The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Pool Real Configured Space (GiB)
The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System MDisk Available Space (GiB)
The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage Pool Number of MDisks
The number of managed disks in a storage pool.

Storage Pool Number of Storage Volumes
The number of volumes in the storage pool.

Related tasks:

"Running the Pools Historical Capacity report" on page 868
To see how space in pools on storage systems is used, run the Pools Historical Capacity report.

Running the Most Active Pools report

To see which pools are most active on storage systems, run the Most Active Pools report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Pools**.
4. Click **Most Active Pools**.
5. Optional: Select the configuration of the storage systems.

6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active Pools report"
Shows a chart of the 20 storage pools that are most active, and details for all pools, for a time period that you specify. Use the report to analyze the performance of pools.

Most Active Pools report

Shows a chart of the 20 storage pools that are most active, and details for all pools, for a time period that you specify. Use the report to analyze the performance of pools.

Charts

The bar chart shows the most active pools on storage systems. By default, pools are sorted by the I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided about the most active pools is as follows:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Storage pool	Performance of One Pool, Most Active Volumes, Most Active RAID Arrays, Storage Resource Relationships Summary

Related tasks:

"Running the Most Active Pools report" on page 1028
To see which pools are most active on storage systems, run the Most Active Pools report.

Running the Performance of One Pool report

To see information about the performance of a pool on a storage system, run the Performance of One Pool report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Pools**.
4. Click **Performance of One Pool**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a pool.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of One Pool report"
Shows four charts, and a table of performance metrics for a storage pool over a time period that you specify. Use the report to analyze the performance of a storage pool.

Performance of One Pool report

Shows four charts, and a table of performance metrics for a storage pool over a time period that you specify. Use the report to analyze the performance of a storage pool.

Charts

The charts show the following performance metrics for a pool on a storage system:

- The total overall rates for I/O operations, and the total response times
- The overall I/O rates for read and write operations, and the response times for read and write operations
- The total I/O rates for back-end operations, and overall response times for back-end operations
- The I/O rates for back-end read operations and back-end write operations, and the response rates for back-end read and back-end write operations

Report output

The following information is provided in the report for the period that you specified:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Related reports

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Performance of One Storage System
- Most Active Pools

Related tasks:

"Running the Performance of One Pool report" on page 1030
To see information about the performance of a pool on a storage system, run the Performance of One Pool report.

Running the Compare Performance of Multiple Pools

Use the Compare Performance of Multiple Pools report to compare up to four performance metrics for multiple pools that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Pools**.
4. Click **Compare Performance of Multiple Pools**.
5. Select the storage systems, and then click **Refresh**.
6. Select the pools, and then click **Next**.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
   a. Click **Add Metric**.
   b. Select the category of performance metric, and then select the performance metric.
   c. If you want to remove the last metric that you added, click **Remove Metric**.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference:

"Compare Performance of Multiple Pools report"
Shows up to four performance metrics for multiple storage pools over time. Use the report to compare the performance metrics for storage pools that are on storage systems.

Compare Performance of Multiple Pools report

Shows up to four performance metrics for multiple storage pools over time. Use the report to compare the performance metrics for storage pools that are on storage systems.

Charts

The charts show the performance metrics for the storage pools that you selected for the period that you specified.

Report output

The following information is provided in the report:

Day, hour, or time
The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Performance metrics
Shows the performance metric data for the resource.

Storage system names
Shows the names of the storage systems, and the performance metric data for each storage system.

Pools Shows the names of the storage pools, and the performance metric data for each storage pool.

Related tasks:

"Running the Compare Performance of Multiple Pools" on page 1032
Use the Compare Performance of Multiple Pools report to compare up to four performance metrics for multiple pools that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

Running the Compare Performance of One Pool over Time Ranges report

To compare a performance metric for one pool in two date ranges, run the Compare Performance of One Pool over Time Ranges report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.

3. Click **Pools**.
4. Click **Compare Performance of One Pool over Time Ranges**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage system.
7. Select the category of performance metric about which you want a report, and then select the performance metric.
8. Select an interval.
9. Specify the two date ranges or reporting periods that you want to compare. For example, you can compare the performance metric for this week with last week.

   Tip: To specify a start date and an end date for a reporting period, select **Custom Date Range** from the list of reporting periods.
10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Compare Performance of One Pool over Time Ranges report"
Shows one performance metric on one pool over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Compare Performance of One Pool over Time Ranges report

Shows one performance metric on one pool over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Charts

The chart shows the performance data of the pool on the performance metric you selected, over the time periods that you selected.

Report output

The following information is provided in the report:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Day in Range
The number of the day in the date range at which the performance metric data was collected, relative to the date range that you specified. The number of days that are shown is determined by the reporting period that was selected for the report. If **Last 30 Days** was selected, then entries for days 1 - 30 are shown on the report.

Hour in Range
The number of the hour in the day at which the performance metric data was collected. For example, 0 is between 12:00:00 a.m. and 12:59:59 a.m., 1 is between 01:00:00 a.m. and 01:59:59 a.m.

Minute in Range
The interval of minutes during which the performance metric data was collected. For example, `0` shows that performance data was collected 0 - 5 minutes after the hour. Similarly, `5` shows that the performance data that was collected 5 - 10 minutes after the hour.

First Date Range
The details of the performance metric data in the first date range. The details are displayed in two columns in the First Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the first date range.

Second Date Range
The details of the performance metric data in the second date range. The details are displayed in two columns in the Second Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the second date range.

Related tasks:

"Running the Compare Performance of One Pool over Time Ranges report" on page 1033
To compare a performance metric for one pool in two date ranges, run the Compare Performance of One Pool over Time Ranges report.

Predefined reports about ports

You can run predefined reports to analyze and compare the performance of ports.

Running the Most Active Ports report

To see which ports are most active on storage systems, run the Most Active Ports report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Ports**.
4. Click **Most Active Ports**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active Ports report"
Shows a chart of the 20 ports that are most active, and details for all ports, for a time period that you specify. Use the report to analyze the performance of ports.

Most Active Ports report

Shows a chart of the 20 ports that are most active, and details for all ports, for a time period that you specify. Use the report to analyze the performance of ports.

Charts

The bar chart shows the most active ports on storage systems. By default, ports are sorted by the I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided about the most active ports is as follows:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Port Name
The name that was assigned to the storage port when the storage system was added to the system.

Storage System Port Location
The physical location of a port on DS8000, DS6000, TotalStorage Enterprise Storage Server storage systems, and on XIV systems. For all other systems, this property shows the WWPN of the port.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Overall Port Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Number of the Fibre Channel (FC) port	Performance of One Port

Related tasks:

"Running the Most Active Ports report" on page 1035
To see which ports are most active on storage systems, run the Most Active Ports report.

Running the Performance of One Port report

To see performance metrics for a port on a storage system, run the Performance of One Port report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Ports**.
4. Click **Performance of One Port**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a port.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of One Port report"
Shows four charts, and a table of performance metrics for a port over a time period that you specify. Use the report to analyze the performance of a port.

Performance of One Port report

Shows four charts, and a table of performance metrics for a port over a time period that you specify. Use the report to analyze the performance of a port.

Charts

The charts show the following performance metrics for a port on a storage system:

- The total I/O rates, and the total response times
- The read I/O rates, the write I/O rates, and the response times
- The read data rates, the write data rates, and the total data rate

- The send bandwidth percentage, the receive bandwidth percentage, and the overall bandwidth percentage

Report output

The following information is provided in the report for the period that you specified:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage Port Name
The name that was assigned to the storage port when the storage system was added to the system.

Storage System Port Location
The physical location of a port on DS8000, DS6000, TotalStorage Enterprise Storage Server storage systems, and on XIV systems. For all other systems, this property shows the WWPN of the port.

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port Send Response Time (ms/op)
The average number of milliseconds to complete a send operation.

Port Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation.

Overall Port Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Send Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage
The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Related reports

To open the Most Active Ports report, click the name of the storage system in the report table.

Related tasks:

"Running the Performance of One Port report" on page 1037
To see performance metrics for a port on a storage system, run the Performance of One Port report.

Running the Compare Performance of Multiple Ports report

Use the Compare Performance of Multiple Ports report to compare up to four performance metrics for multiple ports that are on storage systems. For example, you can compare data rates, I/O rates, and bandwidth percentages.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Ports**.
4. Click **Compare Performance of Multiple Ports**.
5. Select the storage systems.
6. Select ports. To select ports, complete the following steps:
   a. Type a keyword. For example, type `port1` to find the names of ports that start with `port1`. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type `a%` to find ports that begin with 'a' or 'A'. Type `%a` to find ports that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the ports on a resource.
   b. Click **Search**. If the search is successful, the ports are displayed in the **Results** list.
   c. Select the ports that you require from the list, and then click **Insert**.
7. Click **Next**.
8. Select the category of performance metric that you want to include in the report, and then select the performance metric.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

    Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference:

"Compare Performance of Multiple Ports report"
Shows up to four performance metrics for multiple ports over time. Use the report to compare the performance metrics for ports that are on storage systems.

Compare Performance of Multiple Ports report

Shows up to four performance metrics for multiple ports over time. Use the report to compare the performance metrics for ports that are on storage systems.

Charts

The charts show the performance metrics for the ports that you selected for the period that you specified.

Report output

The following information is provided in the report:

Day, hour, or time
The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Performance metrics
Shows the performance metric data for the resource.

Storage system names
Shows the names of the storage systems, and the performance metric data for each storage system.

Ports Shows the names of the storage ports, and the performance metric data for each port.

Related tasks:

"Running the Compare Performance of Multiple Ports report" on page 1039
Use the Compare Performance of Multiple Ports report to compare up to four performance metrics for multiple ports that are on storage systems. For example, you can compare data rates, I/O rates, and bandwidth percentages.

Predefined reports about RAID arrays

You can run predefined reports to analyze and compare the performance of RAID arrays.

Running the Most Active RAID Arrays report

To see which arrays are most active on DS8000, DS6000, and IBM TotalStorage Enterprise Storage Server storage systems, run the Most Active RAID Arrays report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **RAID Arrays**.
4. Click **Most Active RAID Arrays**.
5. Select the storage systems, and then select a sort order.
6. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
7. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active RAID Arrays report"
Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

Most Active RAID Arrays report

Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

Restriction: This report applies only to DS8000, DS6000, and TotalStorage Enterprise Storage Server systems.

Charts

The bar chart shows the most active RAID arrays on DS8000, DS6000, and TotalStorage Enterprise Storage Server storage systems. By default, arrays are sorted by the I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided about the most active RAID arrays is as follows:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Controller Name
The name that was assigned to the controller when it was added to the system.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage RAID Array Name
The name that was assigned to the RAID array when it was added to the system.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Disk Utilization Percentage
The average percentage of time that the disks that are associated with an array are busy.

No value is calculated for this property if there are multiple ranks in the extent pool where the thin-provisioned volumes are allocated. In this case, the value `N/A` is displayed for this property. This limitation applies only to DS8000 storage systems.

If there is only a single rank in the extent pool, the value for this property is calculated regardless of the thin-provisioned volumes.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Controller	Performance of One Controller or Module
RAID array	Performance of One RAID Array

Related tasks:

"Running the Most Active RAID Arrays report" on page 1040
To see which arrays are most active on DS8000, DS6000, and IBM TotalStorage Enterprise Storage Server storage systems, run the Most Active RAID Arrays report.

Running the Performance of One RAID Array report

To see performance metrics for a RAID array on DS8000, DS6000, or IBM TotalStorage Enterprise Storage Server storage systems, run the Performance of One RAID Array report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **RAID Arrays**.
4. Click **Performance of One RAID Array**.
5. Select a storage system.
6. Select a RAID array.
7. Optional: Select an interval.
8. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
9. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of One RAID Array report"
Shows four charts, and a table of performance metrics for a RAID array over a time period. Use the report to analyze the performance of a RAID array.

Performance of One RAID Array report

Shows four charts, and a table of performance metrics for a RAID array over a time period. Use the report to analyze the performance of a RAID array.

Restriction: This report applies only to DS8000, DS6000, and TotalStorage Enterprise Storage Server systems.

Charts

The charts show the following performance statistics for arrays on storage systems:

- The total I/O rates, and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The back-end I/O rates, and the back-end I/O response times
- The back-end read I/O rates, the back-end write I/O rates, the read response times, and the write response times

Report output

The following information is provided in the report for the period that you specified:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Controller Name
The name that was assigned to the controller when it was added to the system.

Storage RAID Array Name
The name that was assigned to the RAID array when it was added to the system.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Disk Utilization Percentage

The average percentage of time that the disks that are associated with an array are busy.

No value is calculated for this property if there are multiple ranks in the extent pool where the thin-provisioned volumes are allocated. In this case, the value `N/A` is displayed for this property. This limitation applies only to DS8000 storage systems.

If there is only a single rank in the extent pool, the value for this property is calculated regardless of the thin-provisioned volumes.

Related reports

To open the Most Active RAID Arrays report, click the name of the storage system in the report table.

Related tasks:

"Running the Performance of One RAID Array report" on page 1043
To see performance metrics for a RAID array on DS8000, DS6000, or IBM TotalStorage Enterprise Storage Server storage systems, run the Performance of One RAID Array report.

Running the Compare Performance of Multiple RAID Arrays report

To compare the performance of multiple arrays that are on storage systems, use the Compare Performance of Multiple RAID Arrays report. For example, you can compare data rates, I/O rates, and response times.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **RAID Arrays**.
4. Click **Compare Performance of Multiple RAID Arrays**.
5. Select the storage systems, and then click **Refresh**.
6. Select the RAID arrays, and then click **Next**.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
   a. Click **Add Metric**.
   b. Select the category of performance metric, and then select the performance metric.
   c. If you want to remove the last metric that you added, click **Remove Metric**.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

    Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
11. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference:

"Compare Performance of Multiple RAID Arrays report"
Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

Compare Performance of Multiple RAID Arrays report

Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

Restriction: This report applies only to DS8000, DS6000, and TotalStorage Enterprise Storage Server systems.

Charts

The charts show the performance metrics for the arrays that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metric
Shows the time at which the performance metric data was collected.

Storage system names
Shows the names of the storage systems, and the performance metric data for each storage system.

RAID arrays
Shows the names of the RAID arrays, and the performance metric data for each RAID array.

Related tasks:

"Running the Compare Performance of Multiple RAID Arrays report" on page 1045
To compare the performance of multiple arrays that are on storage systems, use the Compare Performance of Multiple RAID Arrays report. For example, you can compare data rates, I/O rates, and response times.

Predefined reports about volumes

You can run predefined reports to analyze the capacity and performance of volumes.

Running the Volumes Capacity report

To see information about the capacity of volumes on storage systems, run the Volumes Capacity report.

To run the report with the default settings, click **Finish**. By default all of the storage systems in the list are selected.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.

2. Click **Storage Systems**.
3. Click **Volumes**.
4. Click **Volumes Capacity**.
5. Optional: Select the configuration of the storage systems.
6. Select one or more storage systems.
7. Optional: If you selected SAN Volume Controller or Storwize V7000 storage systems, you can select the type of volume.

Option	Description
All Volumes	Select **All**, the default, to view information about all of the volumes on the selected storage systems.
Only Primary Volumes	Select **Primary** to view information about primary volumes on the selected storage systems.
Only Secondary Volumes	Select **Secondary** to view information about secondary volumes on the selected storage systems.

8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Volumes Capacity report"
Shows a chart of the 20 volumes with the most used space, and a table with details of all volumes. Use the report to review the space allocation in your volumes.

Volumes Capacity report

Shows a chart of the 20 volumes with the most used space, and a table with details of all volumes. Use the report to review the space allocation in your volumes.

Charts

The bar chart shows the used space for volumes on storage systems.

Report output

For each volume on a storage system, the following information is provided:

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`,

the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage Volume WWN
The worldwide name of the volume.

Storage Volume Capacity (GiB)
The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

For XIV systems, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage Volume Allocated Space (GiB)
The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

Storage Volume Used Space (GiB)
The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Volume Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Volume Physical Allocation Percentage
The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this value:

`(allocated space ÷ capacity) × 100`

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Volume Shortfall Percentage
The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume.

The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
(unallocatable volume Space ÷ volume unallocated space) × 100
```

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This Storage Volume Shortfall Percentage represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

```
100 GiB unallocated volume space - 70 GiB remaining
pool space ÷ 100 GiB unallocated volume space × 100
```

Storage Volume Used Allocated Space Percentage
The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
(volume used space ÷ volume allocated space) × 100
```

Storage Volume RAID Level
The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Volume Is Assigned
Shows whether the volume is assigned to a server. If this value is `Yes`, the volume is assigned to a server.

Storage Volume Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Volume Type
The type of storage volume. For example, the storage volume can be striped or sequential. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

Storage Volume Grain Size (KiB)
The grain size with which a thin-provisioned volume was created. This

value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

Storage Volume Warning Level
The percentage of volume capacity that is used at which a warning is generated. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

Storage Volume Logical Subsystem
The logical subsystem (LSS) to which a volume or pool belongs.

Storage Volume Copy ID
The identifier of a volume copy.

Storage Volume Copy Is Primary
Shows whether the volume copy is the primary volume copy.

Related reports

To open the Volumes Historical Capacity report, click the name of the volume in the report table.

Related tasks:

"Running the Volumes Capacity report" on page 1046
To see information about the capacity of volumes on storage systems, run the Volumes Capacity report.

Running the Volumes Historical Capacity report

To see how space on volumes on storage systems is used, run the Volumes Historical Capacity report.

You can see how space on a volume on a storage system, or on volumes in a storage resource group, are used.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Volumes**.
4. Click **Volumes Historical Capacity**.
5. Select the type of historical capacity report that you want to create. You can select either of the following options:
   - **Of a Volume in a Storage System**
     a. Click **Next**.
     b. Optional: Select the configuration of the storage system.
     c. Select a storage system.
     d. Type a keyword and click **Search**. For example, type `lun`, to find the names of volumes that start with `lun`. To change the default search option, click **Options**.

        Tip: Use the percent sign (%) as a wildcard. For example, type `a%` to find volume names that begin with 'a' or 'A'. Type `%a` to find volume names that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the volumes on a resource.

e. Select a volume.

- **Of Volumes in a Storage Resource Group**
  a. Click **Next**.
  b. Select a storage resource group.

6. Select an interval.
7. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Volumes Historical Capacity report" on page 872
Shows the usage of space on volumes over a specified period. Use the report to analyze changes in your volume requirements, particularly the increasing usage of space on thin-provisioned volumes.

Volumes Historical Capacity report

Shows the usage of space on volumes over a specified period. Use the report to analyze changes in your volume requirements, particularly the increasing usage of space on thin-provisioned volumes.

Charts

The chart shows the following space statistics for the volume on a storage system:

- The capacity of the volume
- The used space on the volume
- The space on the volume that is used for system management
- The allocated space on the volume

The chart shows the following space statistics for the volumes in a storage resource group:

- The capacity of the volumes
- The used space on the volumes
- The space on the volumes that is used for system management
- The allocated space on the volume

If the volume is thin provisioned, the chart shows the total amount of virtual space that is allocated to the volume, and the total amount of real space that is used by the volume.

Report output

The following information is included in reports for a volume on a storage system, or for volumes in a storage resource group:

Storage System Name
A user-defined name of the storage system. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Date The date that the data was collected for the report.

Storage Volume Capacity (GiB)
The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

For XIV systems, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage Volume Used Space (GiB)
The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Volume Allocated Space (GiB)
The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

Storage Volume Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Related tasks:

"Running the Volumes Historical Capacity report" on page 871
To see how space on volumes on storage systems is used, run the Volumes Historical Capacity report.

Running the Most Active Volumes report

To see which volumes are most active on storage systems, run the Most Active Volumes report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Volumes**.
4. Click **Most Active Volumes**.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period. If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
8. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Most Active Volumes report"
Shows a chart of the 20 volumes that are most active, and details for all volumes, for a time period that you specify. Use the report to analyze the performance of volumes.

Most Active Volumes report

Shows a chart of the 20 volumes that are most active, and details for all volumes, for a time period that you specify. Use the report to analyze the performance of volumes.

Charts

The bar chart shows the most active volumes on storage systems. By default, volumes are sorted by the I/O rate.

Tip: To change the sort order, select a value from the **Sort Order** list.

Report output

The information that is provided about the most active volumes is as follows:

Storage System Name
A user-defined name of the storage system. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Volume Utilization
The average percentage of time that the volume is busy.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Storage pool	Performance of One Pool
Volume	Performance of One Volume, Storage Resource Relationships Summary

Related tasks:

"Running the Most Active Volumes report" on page 1053
To see which volumes are most active on storage systems, run the Most Active Volumes report.

Running the Performance of One Volume report

To see performance metrics for volumes on a storage system, run the Performance of One Volume report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Volumes**.
4. Click **Performance of One Volume**.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a volume. To select a volume, complete the following steps:
   a. Type a keyword. For example, type `lun`, to find the names of volumes that start with `lun`. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type `a%` to find volume names that begin with 'a' or 'A'. Type `%a` to find volume names that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the volumes on a resource.
   b. Click **Search**. If the search is successful, the names of the volumes are displayed in the **Choice** list.
   c. From the list, select a volume.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
10. Click **Finish**.

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference:

"Performance of One Volume report"
Shows four charts, and a table of performance metrics for a volume over a time period that you specify. Use the report to analyze the performance of a volume.

Performance of One Volume report

Shows four charts, and a table of performance metrics for a volume over a time period that you specify. Use the report to analyze the performance of a volume.

Charts

The charts show the following performance metrics for a volume on a storage system:

- The total I/O rates, and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates, and the write date rates
- The read cache percentages, the write cache percentages, and the volume utilization percentages

Report output

The following information is provided in the report for the period that you specified:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Volume Utilization
The average percentage of time that the volume is busy.

Related reports

To open the Most Active Volumes report, click the name of the storage system in the report table.

Related tasks:

"Running the Performance of One Volume report" on page 1055
To see performance metrics for volumes on a storage system, run the Performance of One Volume report.

Running the Compare Performance of Multiple Volumes report

Use the Compare Performance of Multiple Volumes report to compare up to four performance metrics for multiple volumes that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Predefined Reports**.
2. Click **Storage Systems**.
3. Click **Volumes**.
4. Click **Compare Performance of Multiple Volumes**.
5. Select the storage systems.
6. Select volumes. To select volumes, complete the following steps:
   a. Type a keyword. For example, type `example1` to find the names of volumes that start with `example1`. To change the default search option, click **Options**. Use the percent sign as a wildcard. For example, type `a%` to find volumes that begin with 'a' or 'A'. Type `%a` to find volumes that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the volumes on a resource.
   b. Click **Search**. If the search is successful, the volumes are displayed in the **Results** list.
   c. Select the volumes that you require from the list, and then click **Insert**.
7. Click **Next**.
8. Select the category of performance metric that you want to include in the report, and then select the performance metric.
9. Optional: To add more metrics to the report, complete the following steps:
   a. Click **Add Metric**.
   b. Select the category of performance metric, and then select the performance metric.
   c. If you want to remove the last metric that you added, click **Remove Metric**.
10. Optional: Select an interval.
11. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

12. Click **Finish**.

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference:

"Compare Performance of Multiple Volumes report"
Shows up to four performance metrics for multiple volumes over time. Use the report to compare the performance metrics for volumes that are on storage systems.

Compare Performance of Multiple Volumes report

Shows up to four performance metrics for multiple volumes over time. Use the report to compare the performance metrics for volumes that are on storage systems.

Charts

The charts show the performance metrics for the volumes that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metrics
Shows the data for each performance metric.

Storage system names
Shows the names of the storage systems, and the performance metric data for each storage system.

Storage volume names
Shows the names of the storage volumes, and the performance metric data for each volume.

Related tasks:

"Running the Compare Performance of Multiple Volumes report" on page 1057
Use the Compare Performance of Multiple Volumes report to compare up to four performance metrics for multiple volumes that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

Predefined reports about storage tiering

Use storage tier reports to help you determine the best storage-tier configuration for your environment.

Information collection

IBM Tivoli Storage Productivity Center collects information from:

- SAN Volume Controller
- Storwize V7000

- Storwize V7000 Unified

Tivoli Storage Productivity Center provides Tivoli Common Reporting reports that show performance and usage data for the SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified systems in your environment.

Workload activity reports

Before you run workload activity reports, you must set the back-end properties of the managed disk (MDisk) groups in your environment.

You can use the Tivoli Storage Productivity Center CLI called tcptool to set the back-end properties of MDisk groups. Alternatively, you can set the properties in the MDisk Group Details page in the Tivoli Storage Productivity Center stand-alone GUI.

With the information that is collected from Tivoli Storage Productivity Center, you can run workload activity reports. The reports provide information about the workload activity of these resources:

- Managed disk (MDisk) groups
- Virtual disks (VDisks) in MDisk groups
- VDisks in storage resource groups
- VDisks in storage virtualizers

By analyzing workload activity reports, you can determine:

- Whether the workload activity of an MDisk group is too low or too high
- Whether a VDisk should be migrated in one of the following ways:
  - To an MDisk group with a faster read I/O capability rate
  - To an MDisk group with a slower read I/O capability rate

Explanation: In the predefined reports about performance, resources, and historical information, MDisk groups are called pools, and VDisks are called volumes.

Related information:

IBM Tivoli Storage Productivity Center Information Center
You can use commands on the Tivoli Storage Productivity Center command-line interface (CLI) to provide information about RAID arrays and back-end storage subsystems. For more information about these commands, see the Tivoli Storage Productivity Center Information Center at http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp. Search for “setarray”, “autosetarray”, and “lsarray”.

Running the Storage Resource Group - VDisk Workload Activity Details report

To see which virtual disks (VDisks) are most active in a storage resource group, run the Storage Resource Group - VDisk Workload Activity Details report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **VDisk Workload Activity Reports**.
4. Select **Storage Resource Group - VDisk Workload Activity Details**, then click **Next**.
5. Select a storage virtualizer, and then select a storage resource group.

6. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
7. Click **Finish**.

Related reference:

"Storage Resource Group - VDisk Workload Activity Details report"
Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDisks in the group.

Storage Resource Group - VDisk Workload Activity Details report

Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDisks in the group.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Reporting period

The name of the storage virtualizer, the name of the storage resource group, the date filter such as `Last 7 days`, and the reporting period are displayed in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is `myverylongmanageddiskgroupname123` in a chart, it is truncated and displayed as `myverylongmanageddiskgroupname...`

Two charts are displayed:

Most Active VDisks chart

For each storage resource group, two bars are shown:

- A yellow bar that shows average VDisk utilization

- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

```
average(vdisk_peak_utilization)
```

To calculate maximum VDisk utilization, the following formula is used:

```
maximum(vdisk_peak_utilization)
```

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:

   ```
   86.30 + 50.01 + 79.41 = 215.72
   ```

2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.

   ```
   215.72 ÷ 3 = 71.91
   ```

In the sample, the average peak VDisk utilization is 71.91.

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

Least Active VDisks chart

For each storage resource group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active VDisks chart are also used in the least active VDisks chart.

The number of VDisks that a storage resource group uses determines how many VDisks are displayed in each chart, as shown in the table.

Number of VDisks that the storage resource group uses	Number of VDisks in the Most Active VDisks chart	Number of VDisks in the Least Active VDisks chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3

Number of VDisks that the storage resource group uses	Number of VDisks in the Most Active VDisks chart	Number of VDisks in the Least Active VDisks chart
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

Tip: Click the bar that represents a VDisk in a chart to open the VDisk Details Report. Alternatively, click the name of the VDisk in the report. Click the name of the MDisk group in the report to open the MDisk Group Details Report. To monitor the workload activity of VDisks for all storage resource groups in a storage virtualizer, run the Storage Resource Group - VDisk Workload Activity Report.

Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDisks are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: *{%}*. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used: (*maximum_response_time_value* × *maximum_response_time_threshold value*)/100. The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

The following information is provided about VDisks in the storage resource group:

MDisk Group
The name of the MDisk group.

VDisk The name of the VDisk.

VDisk Report Date
The date that the data was collected for the report.

VDisk Capacity (GB)
The total storage capacity of the VDisk in GB.

VDisk Thin Provisioning Capacity Status
The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is ≥ the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is < the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

VDisk Peak Utilization (%)
The highest value of VDisk utilization over a 24-hour period.

VDisk Read I/O Rate (ops/s)
The read I/O rate of the VDisk in operations per second.

VDisk Max Read I/O Rate (ops/s)
The maximum read I/O rate of the VDisk in operations per second.

VDisk Average Read Response Time (ms/op)
The average read response time in milliseconds per operation.

VDisk Max Read Response Time (ms/op)
The maximum read response time in milliseconds per operation.

VDisk Max Read Response Time Threshold of *{%}* **Exceeded By (hr.:min.)**
The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Write I/O Rate (ops/s)
The write I/O rate in operations per second.

VDisk Average Write Response Time (ms/op)
The average write response time in milliseconds per operation.

VDisk Max Write Response Time (ms/op)
The maximum write response time in milliseconds per operation.

VDisk Max Write Response Time Threshold of *{%}* **Exceeded By (hr.:min.)**
The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Max Read Density Rate (ops/s per GB)
The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)
The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDisks in the MDisk group.

VDisk Capacity Share of Total VDisk Capacity (%)
The real capacity of the VDisk as a percentage of the real capacity of all the VDisks in the MDisk group.

VDisk Average Read Cache Hits (%)
The average read cache hits in percent.

VDisk Max Read Cache Hits (%)
The maximum read cache hits in percent.

VDisk Max Read Cache Hits Time Threshold of *{%}* **Exceeded By (hr.:min.)**
The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

VDisk Max Read I/O Rate Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write I/O Rate Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write Cache Delay I/O Rate (ops/s)
The maximum write cache delay I/O rate of the VDisk in operations per second.

Related tasks:

"Running the Storage Resource Group - VDisk Workload Activity Details report" on page 1059
To see which virtual disks (VDisks) are most active in a storage resource group, run the Storage Resource Group - VDisk Workload Activity Details report.

Running the VDisk Details report

To see information about a virtual disk (VDisk), run the VDisk Details report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **VDisk Details**.
4. Select a storage virtualizer, and then select a VDisk.
5. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
6. Click **Finish**.

Related reference:

"VDisk Details report"
Shows the workload activity of a virtual disk (VDisk) in a managed disk group. Use the report to analyze the performance of a VDisk in a managed disk (MDisk) group.

VDisk Details report

Shows the workload activity of a virtual disk (VDisk) in a managed disk group. Use the report to analyze the performance of a VDisk in a managed disk (MDisk) group.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Reporting period

The name of the storage virtualizer, the name of the VDisk, the date filter such as `Last 7 days`, and the reporting period are displayed in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a

report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDisks are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: *{%}*. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used: (*maximum_response_time_value* × *maximum_response_time_threshold value*)/100. The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

The following information is displayed about the VDisk:

MDisk Group
The name of the MDisk group.

VDisk The name of the VDisk.

VDisk Report Date
The date that the data was collected for the report.

VDisk Capacity (GB)
The total storage capacity of the VDisk in GB.

VDisk Thin Provisioning Capacity Status
The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is ≥ the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is < the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

VDisk Peak Utilization (%)
The highest value of VDisk utilization over a 24-hour period.

VDisk Read I/O Rate (ops/s)
The read I/O rate of the VDisk in operations per second.

VDisk Max Read I/O Rate (ops/s)
The maximum read I/O rate of the VDisk in operations per second.

VDisk Average Read Response Time (ms/op)
The average read response time in milliseconds per operation.

VDisk Max Read Response Time (ms/op)
The maximum read response time in milliseconds per operation.

VDisk Max Read Response Time Threshold of *{%}* **Exceeded By (hr.:min.)**
The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Write I/O Rate (ops/s)
The write I/O rate in operations per second.

VDisk Average Write Response Time (ms/op)
The average write response time in milliseconds per operation.

VDisk Max Write Response Time (ms/op)
The maximum write response time in milliseconds per operation.

VDisk Max Write Response Time Threshold of *{%}* **Exceeded By (hr.:min.)**
The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Max Read Density Rate (ops/s per GB)
The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)
The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDisks in the MDisk group.

VDisk Capacity Share of Total VDisk Capacity (%)
The real capacity of the VDisk as a percentage of the real capacity of all the VDisks in the MDisk group.

VDisk Average Read Cache Hits (%)
The average read cache hits in percent.

VDisk Max Read Cache Hits (%)
The maximum read cache hits in percent.

VDisk Max Read Cache Hits Time Threshold of *{%}* **Exceeded By (hr.:min.)**
The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

VDisk Max Read I/O Rate Threshold of *{%}* **Exceeded By (hr.:min.)**
The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write I/O Rate Threshold of *{%}* **Exceeded By (hr.:min.)**
The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write Cache Delay I/O Rate (ops/s)
The maximum write cache delay I/O rate of the VDisk in operations per second.

Related tasks:

"Running the VDisk Details report" on page 1064
To see information about a virtual disk (VDisk), run the VDisk Details report.

Running the MDisk Groups - VDisk Workload Activity report

To see which virtual disks (VDisks) are most active in managed disk (MDisk) groups, run the MDisk Groups - VDisk Workload Activity report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **VDisk Workload Activity Reports**.
4. Select **MDisk Groups - VDisk Workload Activity**, then click **Next**.
5. Select a storage virtualizer.
6. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
7. Click **Finish**.

Related reference:

"MDisk Groups - VDisk Workload Activity report"
Use this report to monitor the performance of virtual disks (VDisks) in managed disk (MDisk) groups. The workload activity of each MDisk group is determined by calculating the average and the maximum peak utilization of the VDisks in the MDisk group.

MDisk Groups - VDisk Workload Activity report

Use this report to monitor the performance of virtual disks (VDisks) in managed disk (MDisk) groups. The workload activity of each MDisk group is determined by calculating the average and the maximum peak utilization of the VDisks in the MDisk group.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Reporting period

The name of the storage virtualizer, the date filter such as `Last 7 days`, and the reporting period are displayed in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated.

To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is myverylongmanageddiskgroupname123 in a chart, it is truncated and displayed as myverylongmanageddiskgroupname...

Two charts are shown:

Most Active MDisk Groups chart

For each MDisk group two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

```
average(vdisk_peak_utilization)
```

To calculate maximum VDisk utilization, the following formula is used:

```
maximum(vdisk_peak_utilization)
```

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:

   ```
   86.30 + 50.01 + 79.41 = 215.72
   ```

2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.

   ```
   215.72 ÷ 3 = 71.91
   ```

In the sample, the average peak VDisk utilization is 71.91.

For each MDisk group, the maximum VDisk utilization is also shown in the chart. The maximum VDisk utilization is the highest value for a VDisk in the VDisk Utilization (%) column. In the sample, the highest value for VDisk1 is 86.30.

Least Active MDisk Groups chart

For each MDisk group two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active MDisk group chart are also used in the least active MDisk group chart.

The number of MDisk groups in the storage virtualizer determines how many MDisk groups are shown in each chart, as shown in the following table.

Number of MDisk groups in the storage virtualizer	Number of MDisk groups in the Most Active MDisk Groups chart	Number of MDisk groups in the Least Active MDisk Groups chart
≥ 10	5	5

Number of MDisk groups in the storage virtualizer	Number of MDisk groups in the Most Active MDisk Groups chart	Number of MDisk groups in the Least Active MDisk Groups chart
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

Tip: Click the bar that represents an MDisk group in a chart to open the MDisk Group - VDisk Workload Activity Details Report. Alternatively, click the name of the MDisk group in the report. Click the name of the VDisk in the report to open the VDisk Details Report.

Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDisks are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: *{%}*. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used: (*maximum_response_time_value* × *maximum_response_time_threshold value*)/100. The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

For each VDisk in an MDisk group, the following information is provided:

MDisk Group
The name of the MDisk group.

VDisk The name of the VDisk.

VDisk Report Date
The date that the data was collected for the report.

VDisk Capacity (GB)
The total storage capacity of the VDisk in GB.

VDisk Thin Provisioning Capacity Status
The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is ≥ the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is < the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

VDisk Peak Utilization (%)
The highest value of VDisk utilization over a 24-hour period.

VDisk Read I/O Rate (ops/s)
The read I/O rate of the VDisk in operations per second.

VDisk Max Read I/O Rate (ops/s)
The maximum read I/O rate of the VDisk in operations per second.

VDisk Average Read Response Time (ms/op)
The average read response time in milliseconds per operation.

VDisk Max Read Response Time (ms/op)
The maximum read response time in milliseconds per operation.

VDisk Max Read Response Time Threshold of *{%}* **Exceeded By (hr.:min.)**
The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Write I/O Rate (ops/s)
The write I/O rate in operations per second.

VDisk Average Write Response Time (ms/op)
The average write response time in milliseconds per operation.

VDisk Max Write Response Time (ms/op)
The maximum write response time in milliseconds per operation.

VDisk Max Write Response Time Threshold of *{%}* **Exceeded By (hr.:min.)**
The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Max Read Density Rate (ops/s per GB)
The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)
The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDisks in the MDisk group.

VDisk Capacity Share of Total VDisk Capacity (%)
The real capacity of the VDisk as a percentage of the real capacity of all the VDisks in the MDisk group.

VDisk Average Read Cache Hits (%)
The average read cache hits in percent.

VDisk Max Read Cache Hits (%)
The maximum read cache hits in percent.

VDisk Max Read Cache Hits Time Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

VDisk Max Read I/O Rate Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write I/O Rate Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write Cache Delay I/O Rate (ops/s)
The maximum write cache delay I/O rate of the VDisk in operations per second.

Related tasks:

"Running the MDisk Groups - VDisk Workload Activity report" on page 1067
To see which virtual disks (VDisks) are most active in managed disk (MDisk) groups, run the MDisk Groups - VDisk Workload Activity report.

Running the Storage Virtualizer - VDisk Workload Activity report

To see which virtual disks (VDisks) are most active on a storage virtualizer, run the Storage Virtualizer - VDisk Workload Activity report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **VDisk Workload Activity Reports**.
4. Select **Storage Virtualizer - VDisk Workload Activity**, then click **Next**.
5. Select a storage virtualizer.
6. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
7. Click **Finish**.

Related reference:

"Storage Virtualizer - VDisk Workload Activity report"
Shows the workload activity of virtual disks (VDisks) in a storage virtualizer. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

Storage Virtualizer - VDisk Workload Activity report

Shows the workload activity of virtual disks (VDisks) in a storage virtualizer. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Reporting period

The name of the storage virtualizer, the date filter such as `Last 7 days`, and the reporting period are shown in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is `myverylongmanageddiskgroupname123` in a chart, it is truncated and displayed as `myverylongmanageddiskgroupname...`

Two charts are shown:

Most Active VDisks chart

For each storage virtualizer, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

```
average(vdisk_peak_utilization)
```

To calculate maximum VDisk utilization, the following formula is used:

```
maximum(vdisk_peak_utilization)
```

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:

   ```
   86.30 + 50.01 + 79.41 = 215.72
   ```

2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.

   ```
   215.72 ÷ 3 = 71.91
   ```

In the sample, the average peak VDisk utilization is 71.91.

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

Least Active VDisks chart

For each storage virtualizer, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active VDisks chart are also used in the least active VDisks chart.

The number of VDisks in a storage virtualizer determines how many VDisks are shown in each chart, as shown in the following table.

Number of VDisks in the storage virtualizer	Number of VDisks in the Most Active VDisks chart	Number of VDisks in the Least Active VDisks chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

Tip: Click the bar that represents a VDisk in a chart to open the VDisk Details Report. Alternatively, click the name of the VDisk in the report. Click the name of the MDisk group in the report to open the MDisk Group Details Report.

Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDisks are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read

response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: *{%}*. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used: (*maximum_response_time_value* × *maximum_response_time_threshold value*)/100. The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

For each VDisk in a storage virtualizer, the following information is shown:

MDisk Group
The name of the MDisk group.

VDisk The name of the VDisk.

VDisk Report Date
The date that the data was collected for the report.

VDisk Capacity (GB)
The total storage capacity of the VDisk in GB.

VDisk Thin Provisioning Capacity Status
The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is ≥ the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is < the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

VDisk Peak Utilization (%)
The highest value of VDisk utilization over a 24-hour period.

VDisk Read I/O Rate (ops/s)
The read I/O rate of the VDisk in operations per second.

VDisk Max Read I/O Rate (ops/s)
The maximum read I/O rate of the VDisk in operations per second.

VDisk Average Read Response Time (ms/op)
The average read response time in milliseconds per operation.

VDisk Max Read Response Time (ms/op)
The maximum read response time in milliseconds per operation.

VDisk Max Read Response Time Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Write I/O Rate (ops/s)
The write I/O rate in operations per second.

VDisk Average Write Response Time (ms/op)
The average write response time in milliseconds per operation.

VDisk Max Write Response Time (ms/op)
The maximum write response time in milliseconds per operation.

VDisk Max Write Response Time Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Max Read Density Rate (ops/s per GB)
The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)
The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDisks in the MDisk group.

VDisk Capacity Share of Total VDisk Capacity (%)
The real capacity of the VDisk as a percentage of the real capacity of all the VDisks in the MDisk group.

VDisk Average Read Cache Hits (%)
The average read cache hits in percent.

VDisk Max Read Cache Hits (%)
The maximum read cache hits in percent.

VDisk Max Read Cache Hits Time Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

VDisk Max Read I/O Rate Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write I/O Rate Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write Cache Delay I/O Rate (ops/s)
The maximum write cache delay I/O rate of the VDisk in operations per second.

Related tasks:

"Running the Storage Virtualizer - VDisk Workload Activity report" on page 1071
To see which virtual disks (VDisks) are most active on a storage virtualizer, run the Storage Virtualizer - VDisk Workload Activity report.

Running the Storage Resource Group - VDisk Workload Activity report

To see which virtual disks (VDisks) are most active in storage resource groups, run the Storage Resource Group - VDisk Workload Activity report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **VDisk Workload Activity Reports**.
4. Select **Storage Resource Group - VDisk Workload Activity**.
5. Select a storage virtualizer.
6. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.

7. Click **Finish**.

Related reference:

"Storage Resource Group - VDisk Workload Activity report"
Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDisks in the group.

Storage Resource Group - VDisk Workload Activity report

Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDisks in the group.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Reporting period

The name of the storage virtualizer, the date filter such as `Last 7 days`, and the reporting period are shown in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is `myverylongmanageddiskgroupname123` in a chart, it is truncated and displayed as `myverylongmanageddiskgroupname...`

Two charts are shown:

Most Active Storage Resource Groups chart

For each storage resource group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

```
average(vdisk_peak_utilization)
```

To calculate maximum VDisk utilization, the following formula is used:

```
maximum(vdisk_peak_utilization)
```

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:
   ```
   86.30 + 50.01 + 79.41 = 215.72
   ```
2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.
   ```
   215.72 ÷ 3 = 71.91
   ```

In the sample, the average peak VDisk utilization is 71.91.

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

Least Active Storage Resource Groups chart

For each storage resource group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active storage resource groups chart are also used in the least active storage resource groups chart.

The number of storage resource groups in a storage virtualizer determines how many storage resource groups are shown in each chart, as shown in the table.

Number of storage resource groups in the storage virtualizer	Number of storage resource groups in the Most Active Storage Resource Groups chart	Number of storage resource groups in the Least Active Storage Resource Groups chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1

Number of storage resource groups in the storage virtualizer	Number of storage resource groups in the Most Active Storage Resource Groups chart	Number of storage resource groups in the Least Active Storage Resource Groups chart
2	1	1
1	1	0

Tip: Click the bar that represents a storage resource group to open the Storage Resource Group: VDisk Workload Activity Details Report. Alternatively, click the name of the storage resource group in the report. Click the name of the VDisk in the report to open the VDisk Details Report. Click the name of the MDisk group in the report to open the MDisk Group Details Report.

Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDisks are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: *{%}*. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used: (*maximum_response_time_value* × *maximum_response_time_threshold value*)/100. The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

For each storage resource group in a storage virtualizer, the following information is provided about VDisks:

Storage Resource Group
The name of the storage resource group.

MDisk Group
The name of the MDisk group.

VDisk The name of the VDisk.

VDisk Report Date
The date that the data was collected for the report.

VDisk Capacity (GB)
The total storage capacity of the VDisk in GB.

VDisk Thin Provisioning Capacity Status
The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is ≥ the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is < the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

VDisk Peak Utilization (%)
The highest value of VDisk utilization over a 24-hour period.

VDisk Read I/O Rate (ops/s)
The read I/O rate of the VDisk in operations per second.

VDisk Max Read I/O Rate (ops/s)
The maximum read I/O rate of the VDisk in operations per second.

VDisk Average Read Response Time (ms/op)
The average read response time in milliseconds per operation.

VDisk Max Read Response Time (ms/op)
The maximum read response time in milliseconds per operation.

VDisk Max Read Response Time Threshold of *{%}* **Exceeded By (hr.:min.)**
The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Write I/O Rate (ops/s)
The write I/O rate in operations per second.

VDisk Average Write Response Time (ms/op)
The average write response time in milliseconds per operation.

VDisk Max Write Response Time (ms/op)
The maximum write response time in milliseconds per operation.

VDisk Max Write Response Time Threshold of *{%}* **Exceeded By (hr.:min.)**
The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Max Read Density Rate (ops/s per GB)
The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)
The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDisks in the MDisk group.

VDisk Capacity Share of Total VDisk Capacity (%)
The real capacity of the VDisk as a percentage of the real capacity of all the VDisks in the MDisk group.

VDisk Average Read Cache Hits (%)
The average read cache hits in percent.

VDisk Max Read Cache Hits (%)
The maximum read cache hits in percent.

VDisk Max Read Cache Hits Time Threshold of *{%}* **Exceeded By (hr.:min.)**
The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

VDisk Max Read I/O Rate Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write I/O Rate Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write Cache Delay I/O Rate (ops/s)
The maximum write cache delay I/O rate of the VDisk in operations per second.

Related tasks:

"Running the Storage Resource Group - VDisk Workload Activity report" on page 1075
To see which virtual disks (VDisks) are most active in storage resource groups, run the Storage Resource Group - VDisk Workload Activity report.

Running the MDisk Group - VDisk Workload Activity Details report

To see which of the virtual disks in a managed disk (MDisk) group are the most active, run the MDisk Group - VDisk Workload Activity Details report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **MDisk Group - VDisk Workload Activity Details**.
4. Select a storage virtualizer, and then select an MDisk group.
5. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
6. Click **Finish**.

Related reference:

"MDisk Group - VDisk Workload Activity Details report"
Shows the workload activity of virtual disks (VDisks) in a managed disk (MDisk) group. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

MDisk Group - VDisk Workload Activity Details report

Shows the workload activity of virtual disks (VDisks) in a managed disk (MDisk) group. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Reporting period

The name of the storage virtualizer, the name of the MDisk group, the date filter such as `Last 7 days`, and the reporting period are displayed on the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is `myverylongmanageddiskgroupname123` in a chart, it is truncated and displayed as `myverylongmanageddiskgroupname...`

Two charts are shown:

Most Active VDisks chart

For each MDisk group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

```
average(vdisk_peak_utilization)
```

To calculate maximum VDisk utilization, the following formula is used:

```
maximum(vdisk_peak_utilization)
```

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:
   ```
   86.30 + 50.01 + 79.41 = 215.72
   ```
2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.
   ```
   215.72 ÷ 3 = 71.91
   ```

In the sample, the average peak VDisk utilization is 71.91.

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

Least Active VDisks chart

For each MDisk group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active MDisk group chart are also used in the least active MDisk group chart.

The number of VDisks in an MDisk group determines how many VDisks are shown in each chart, as shown in the following table.

Number of VDisks in the MDisk group	Number of VDisks in the Most Active VDisks chart	Number of VDisks in the Least Active VDisks chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

Tip: Click the bar that represents a VDisk in a chart to open the VDisk Details Report. Alternatively, click the name of the VDisk in the report. Click the name of the MDisk group in the report to open the MDisk Group Details Report.

Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDisks are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read

response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: *{%}*. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used: (*maximum_response_time_value* × *maximum_response_time_threshold value*)/100. The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

For each VDisk in an MDisk group, the following information is provided:

MDisk Group
The name of the MDisk group.

VDisk The name of the VDisk.

VDisk Report Date
The date that the data was collected for the report.

VDisk Capacity (GB)
The total storage capacity of the VDisk in GB.

VDisk Thin Provisioning Capacity Status
The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is ≥ the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is < the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

VDisk Peak Utilization (%)
The highest value of VDisk utilization over a 24-hour period.

VDisk Read I/O Rate (ops/s)
The read I/O rate of the VDisk in operations per second.

VDisk Max Read I/O Rate (ops/s)
The maximum read I/O rate of the VDisk in operations per second.

VDisk Average Read Response Time (ms/op)
The average read response time in milliseconds per operation.

VDisk Max Read Response Time (ms/op)
The maximum read response time in milliseconds per operation.

VDisk Max Read Response Time Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Write I/O Rate (ops/s)
The write I/O rate in operations per second.

VDisk Average Write Response Time (ms/op)
The average write response time in milliseconds per operation.

VDisk Max Write Response Time (ms/op)
The maximum write response time in milliseconds per operation.

VDisk Max Write Response Time Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Max Read Density Rate (ops/s per GB)
The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)
The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDisks in the MDisk group.

VDisk Capacity Share of Total VDisk Capacity (%)
The real capacity of the VDisk as a percentage of the real capacity of all the VDisks in the MDisk group.

VDisk Average Read Cache Hits (%)
The average read cache hits in percent.

VDisk Max Read Cache Hits (%)
The maximum read cache hits in percent.

VDisk Max Read Cache Hits Time Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

VDisk Max Read I/O Rate Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write I/O Rate Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write Cache Delay I/O Rate (ops/s)
The maximum write cache delay I/O rate of the VDisk in operations per second.

Related tasks:

"Running the MDisk Group - VDisk Workload Activity Details report" on page 1080

To see which of the virtual disks in a managed disk (MDisk) group are the most active, run the MDisk Group - VDisk Workload Activity Details report.

Running the MDisk Group Details report

To see statistics about the performance of a managed disk (MDisk) group, run the MDisk Group Details report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **MDisk Group Details**.
4. Select a storage virtualizer, and then select an MDisk group.
5. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
6. Click **Finish**.

Related reference:

"MDisk Group Details report"
Shows the workload activity of a managed disk (MDisk) group. Use the report to analyze storage services for an MDisk group.

MDisk Group Details report

Shows the workload activity of a managed disk (MDisk) group. Use the report to analyze storage services for an MDisk group.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Reporting period

The name of the storage virtualizer, the name of the MDisk group, the date filter such as `Last 7 days`, and the reporting period are displayed in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

Read I/O rate capability

If you do not set the back-end pool properties of an MDisk group:

- `Not Set` is displayed in the MDisk Group Read I/O Rate Capability column
- `Not Available` is displayed in the MDisk Group Read I/O Rate Availability column

If the read I/O rate capability of an MDisk group is less than the maximum read I/O rate of the MDisk group, negative values are displayed in the report.

Threshold values

The names of the threshold values in the report that are used to monitor workload activity for MDisk groups are:

- Actual Size Threshold
- Read I/O Rate Threshold
- Max Read Response Time Threshold
- Max Write Response Time Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the value of the actual size threshold is changed during a reporting period from 50% to 60%, the name of the column is changed to MDisk Group Available Size < Actual Size Threshold of {50 - 60%}. The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: *{%}*. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:

$(maximum_response_time_value \times maximum_response_time_threshold_value) \div 100.$

Tip: To monitor the workload activity of all MDisk groups in a storage virtualizer, run the MDisk Groups - Workload Activity report.

Report output

The following information is provided about the MDisk group:

MDisk Group
The name of the MDisk group.

MDisk Group Report Date
The date that the data was collected for the report.

MDisk Group Capacity Status
If the storage capacity that is allocated is less than the actual storage capacity, a status value of `OK` is displayed. If the storage capacity that is allocated is more than the actual storage capacity, a status value of `Warning` is displayed. The actual capacity of the MDisk group is calculated as follows:

$$(\text{mdisk_group_real_capacity_value} \div \text{mdisk_group_capacity_value}) \times 100$$

MDisk Group Capacity (GB)
The total storage capacity in GB of all of the disks in the MDisk group.

MDisk Group Real Capacity (GB)
The total amount in GB of real storage capacity in the MDisk group.

MDisk Group Free Space (GB)
The total amount of remaining capacity in GB that can be allocated to increase the size of an existing VDisk, or allocated to an additional VDisk.

MDisk Group Available Size < Actual Size Threshold of *{%}*
If the amount of storage space that is available for allocation to VDisks is less than the actual size threshold, `Yes` is displayed in the column. If the amount of storage space that is available for allocation to VDisks is more than the actual size threshold, `No` is displayed in the column.

MDisk Group Read I/O Rate (ops/s)
The read I/O rate of the MDisk group in operations per second.

MDisk Group Read I/O Rate Capability (ops/s)
The read I/O capability rate of the MDisk group in operations per second.

Tip: To find out how the read I/O capability rate is calculated, see "Read I/O capability formula" on page 1092

MDisk Group Max Read I/O Rate (ops/s)
The maximum read I/O rate of the MDisk group in operations per second.

MDisk Group Read I/O Rate Availability (ops/s)
The difference between the read I/O capability rate and the maximum read I/O rate in operations per second.

MDisk Group Read I/O Rate Availability < Read I/O Rate Capability Threshold of *{%}*
The availability rate for read I/O operations is compared with the read I/O rate capability threshold. If the availability rate for read I/O operations is less than the read I/O rate capability threshold, `Yes` is displayed in the column. If the availability rate for read I/O operations is more than the read I/O rate capability threshold, `No` is displayed in the column.

MDisk Group Max Read Response Time (ms/op)
The maximum read response time in milliseconds per operation.

MDisk Group Max Read Response Time Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum read response time threshold is exceeded over a 24 hour period.

MDisk Group Average Read Response Time (ms/op)
The average read response time in milliseconds per operation.

MDisk Group Max Write I/O Rate (ops/s)
The maximum write I/O rate in operations per second.

MDisk Group Max Write Response Time (ms/op)
The maximum write response time in milliseconds per operation.

MDisk Group Max Write Response Time Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum write response time threshold is exceeded over a 24 hour period.

MDisk Group Average Write Response Time (ms/op)
The average write response time in milliseconds per operation.

MDisk Group Average Write I/O Size (KB)
The average size of write I/O operations in KB.

Related tasks:

"Running the MDisk Group Details report" on page 1084
To see statistics about the performance of a managed disk (MDisk) group, run the MDisk Group Details report.

Running the MDisk Groups - Workload Activity report

To see which managed disk (MDisk) groups are most active, run the MDisk Groups - Workload Activity report.

1. On the **Public Folders** tab, click **IBM Tivoli Productivity Center Packages**.
2. Click **Storage Tiering**.
3. Click **MDisk Groups - Workload Activity**.
4. Select a storage virtualizer.
5. To set the time frame of the report, specify a reporting period. It is easier to analyze small amounts of information over shorter periods.

   Tip: To specify a start date and an end date, select **Custom date range** from the list of reporting periods.
6. Click **Finish**.

Related reference:

"MDisk Groups - Workload Activity report"
Shows the workload activity of managed disk (MDisk) groups. You can use the information that the report provides to investigate which MDisk groups are being underused or overused.

"Read I/O capability formula" on page 1092
The read I/O capability formula is used by IBM Tivoli Storage Productivity Center to estimate the workload capacity of managed disk (MDisk) groups.

MDisk Groups - Workload Activity report

Shows the workload activity of managed disk (MDisk) groups. You can use the information that the report provides to investigate which MDisk groups are being underused or overused.

Restriction: This report applies only to SAN Volume Controller and Storwize systems.

Reporting periods

The name of the storage virtualizer, the date filter such as `Last 7 days`, and the reporting period are displayed in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

Read I/O rate capability

To estimate the read I/O rate capability of an MDisk group, you must set the back-end pool properties of MDisk groups.

If you do not set the back-end pool properties of an MDisk group:

- `Not Set` is displayed in the MDisk Group Read I/O Rate Capability column
- `Not Available` is displayed in the MDisk Group Read I/O Rate Availability column

If the read I/O rate capability of an MDisk group is less than the maximum read I/O rate of the MDisk group, negative values are displayed in the report.

Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is `myverylongmanageddiskgroupname123` in a chart, it is truncated and displayed as `myverylongmanageddiskgroupname...`

Two charts are displayed:

- Most Active MDisk Groups
- Most Available MDisk Groups

Most Active MDisk Groups

For each MDisk group two bars are used:

- A yellow bar that shows the average activity value
- A red bar that shows the maximum activity value

To calculate the average activity value of an MDisk group, the following formula is used:

```
average(maximum_read_I/O_rate ÷ read_I/O_rate_capability)
```

To calculate the maximum activity value, the following formula is used:

```
maximum(maximum_read_I/O_rate ÷ read_I/O_rate_capability)
```

In this sample excerpt of report output for mdiskgroup1, the values that are required to calculate average and maximum activity values are provided.

MDisk Group Report Date	MDisk Group Read I/O Rate Capability (ops/s)	MDisk Group Max Read I/O Rate (ops/s)
Dec 6 2011	331	216.77
Dec 7 2011	331	256.40
Dec 8 2011	331	270.87

The average activity value of an MDisk group is calculated as follows:

1. For each report in the MDisk group, the value in the MDisk Group Max Read I/O Rate column is divided by the value in the MDisk Group Read I/O Rate Capability column:

   ```
   216.77 ÷ 331 = 0.65
   256.40 ÷ 331 = 0.77
   270.87 ÷ 331 = 0.81
   ```

2. The result of the operation for each report is added:

   ```
   0.65 + 0.77 + 0.81 = 2.23
   ```

3. To get the average value, the total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.

   ```
   2.23 ÷ 3 = 0.74
   ```

The maximum value of an MDisk group is calculated as follows. For each report in an MDisk group, the value in the MDisk Group Max Read I/O Rate column is divided by the value in the MDisk Group Read I/O Rate Capability column:

```
216.77 ÷ 331 = 0.65
256.40 ÷ 331 = 0.77
270.87 ÷ 331 = 0.81
```

In this sample, the maximum value is 0.81.

Most Available MDisk Groups

For each MDisk group two bars are used:

- A blue bar to represent minimum values
- A red bar to represent average values

In this sample excerpt of report output for mdiskgroup2, the values that are required to calculate minimum and average availability values for MDisk Groups are provided.

MDisk Group Report Date	MDisk Group Read I/O Rate Availability (ops/s)
Dec 6 2011	435.41
Dec 7 2011	560.92
Dec 8 2011	480.95

The minimum value of an MDisk group is calculated as follows:

- For all of the reports for an MDisk group, the lowest value in the MDisk Group Read I/O Rate Availability column is used. In this sample, the lowest value for mdiskgroup2 is 435.41.

The average value of an MDisk group is calculated as follows:

1. For each report in an MDisk group, the values in the MDisk Group Read I/O Rate Availability column are added:

   ```
   435.41 + 560.92 + 480.95 = 1477.28
   ```

2. To get the average value, the total is divided by the number of days that are specified in the date range. (For example, if the reporting period spans 3 days, the total is divided by 3.)

   `1477.28 ÷ 3 = 492.42`

   In this sample, the average value is 492.42.

The number of MDisk groups in a storage virtualizer determines how many MDisk groups are displayed in each chart, as shown in the following table.

Number of MDisk groups in the storage virtualizer	Number of MDisk groups in the Most Active MDisk Groups chart	Number of MDisk groups in the Least Active MDisk Groups chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

Tip: Click the bar that represents an MDisk group in a chart to open the MDisk Group - VDisk Workload Activity Details Report. Alternatively, you can click the name of the MDisk group in the report.

Threshold values

The names of the threshold values in the report that are used to monitor workload activity for MDisk groups are:

- Actual Size Threshold
- Read I/O Rate Threshold
- Max Read Response Time Threshold
- Max Write Response Time Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the value of the actual size threshold is changed during a reporting period from 50% to 60%, the name of the column is changed to MDisk Group Available Size < Actual Size Threshold of {50 - 60%}. The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: *{%}*. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:

(maximum_response_time_value × *maximum_response_time_threshold_value)*÷ 100.

The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

For each MDisk group, the following information is provided:

MDisk Group
The name of the MDisk group.

MDisk Group Report Date
The date that the data was collected for the report.

MDisk Group Capacity Status
If the storage capacity that is allocated is less than the actual storage capacity, a status value of `OK` is displayed. If the storage capacity that is allocated is more than the actual storage capacity, a status value of `Warning` is displayed. The actual capacity of the MDisk group is calculated as follows:

```
(mdisk_group_real_capacity_value ÷ mdisk_group_capacity_value) × 100
```

MDisk Group Capacity (GB)
The total storage capacity in GB of all of the disks in the MDisk group.

MDisk Group Real Capacity (GB)
The total amount in GB of real storage capacity in the MDisk group.

MDisk Group Free Space (GB)
The total amount of remaining capacity in GB that can be allocated to increase the size of an existing VDisk, or allocated to an additional VDisk.

MDisk Group Available Size < Actual Size Threshold of *{%}*
If the amount of storage space that is available for allocation to VDisks is less than the actual size threshold, `Yes` is displayed in the column. If the amount of storage space that is available for allocation to VDisks is more than the actual size threshold, `No` is displayed in the column.

MDisk Group Read I/O Rate (ops/s)
The read I/O rate of the MDisk group in operations per second.

MDisk Group Read I/O Rate Capability (ops/s)
The read I/O capability rate of the MDisk group in operations per second.

Tip: To find out how the read I/O capability rate is calculated, see "Read I/O capability formula" on page 1092

MDisk Group Max Read I/O Rate (ops/s)
The maximum read I/O rate of the MDisk group in operations per second.

MDisk Group Read I/O Rate Availability (ops/s)
The difference between the read I/O capability rate and the maximum read I/O rate in operations per second.

MDisk Group Read I/O Rate Availability < Read I/O Rate Capability Threshold of *{%}*
The availability rate for read I/O operations is compared with the read I/O rate capability threshold. If the availability rate for read I/O operations is less than the read I/O rate capability threshold, `Yes` is displayed in the column. If the availability rate for read I/O operations is more than the read I/O rate capability threshold, `No` is displayed in the column.

MDisk Group Max Read Response Time (ms/op)
The maximum read response time in milliseconds per operation.

MDisk Group Max Read Response Time Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum read response time threshold is exceeded over a 24 hour period.

MDisk Group Average Read Response Time (ms/op)
The average read response time in milliseconds per operation.

MDisk Group Max Write I/O Rate (ops/s)
The maximum write I/O rate in operations per second.

MDisk Group Max Write Response Time (ms/op)
The maximum write response time in milliseconds per operation.

MDisk Group Max Write Response Time Threshold of *{%}* Exceeded By (hr.:min.)
The number of hours and minutes that the maximum write response time threshold is exceeded over a 24 hour period.

MDisk Group Average Write Response Time (ms/op)
The average write response time in milliseconds per operation.

MDisk Group Average Write I/O Size (KB)
The average size of write I/O operations in KB.

Related tasks:

"Running the MDisk Groups - Workload Activity report" on page 1087
To see which managed disk (MDisk) groups are most active, run the MDisk Groups - Workload Activity report.

Read I/O capability formula

The read I/O capability formula is used by IBM Tivoli Storage Productivity Center to estimate the workload capacity of managed disk (MDisk) groups.

The estimated workload capacity is compared with the actual workload capacity to assess the most active and least active MDisk groups in a storage subsystem. To monitor the workload activity of MDisk groups, you can run customized reports such as the MDisk Groups - Workload Activity Report. This formula is used for MDisk groups in the following storage systems:

- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

The formula that is used by IBM Tivoli Storage Productivity Center to estimate the read I/O capability rate is:

[Estimated Read I/O Workload * Type of Disk I/O Capability * Number of Disks] / [Estimated Read I/O Workload * (1 – Cache Hit Ratio) + Estimated Write I/O Workload * RAID Type Weighted I/O Ratio]

There are two types of value that are used by the formula:

- Fixed values provided by IBM Tivoli Storage Productivity Center
- Variable values that you must specify for each MDisk group

The following fixed values are used by the formula:

Estimated Read I/O Workload
The estimated read I/O workload of an MDisk group. The estimate is set to a fixed value of 0.7 (70%).

Estimated Write I/O Workload
The estimated write I/O workload of an MDisk group. The estimate is set to a fixed value of 0.3 (30%).

You specify the following variable values for each MDisk group:

Type of Disk I/O Capability
When you specify the type of disk that the back-end pool uses:

- You can accept the default I/O capability per second of the disk
- You can enter a new value for the I/O capability per second of the disk

For example, the default value for the I/O capability rate per second of fiber types of disk with 15,000 revolutions per minute (rpm) is 150 I/O operations per second. To change the I/O capability for a type of disk, you issue the **setbackenddisktype** command and you enter a new value for the **-iops** parameter.

Number of Disks
You specify the number of physical disks that the back-end pool uses.

Cache Hit Ratio
When you specify the type of storage system that the back-end pool uses:

- You can accept the default cache hit ratio for the storage system
- You can enter a new value for the cache hit ratio of the storage system

For example, the default cache hit ratio of the XIV system is 0.7 (70%). To change the cache hit ratio for a type of storage system, you issue the **setbackendtype** command and you enter a new value for the **-cachehit** parameter.

RAID Type Weighted I/O Ratio
When you specify the type of RAID that the back-end pool uses:

- You can accept the default value for the weighted I/O ratio of the type of RAID
- You can enter a new value for the weighted I/O ratio of the type of RAID

For example the weighted I/O ratio for RAID5 types of RAID is 4:1. To change the weighted I/O ratio for a type of RAID, you issue the **setbackendraidtype** command and you enter a new value for the **-weightedio** parameter.

Tip: In the stand-alone GUI, you can set the values for the back-end properties of an MDisk group. You can also use the Tivoli Storage Productivity Center CLI called tcptool to set the values for back-end MDisk groups. See the **setarray** command. To change the variable values that are used by the Read I/O capability formula, you must use the Tivoli Storage Productivity Center CLI.

Example

The following values have been set for a back-end pool. The type of storage system is the XIV system. The type of disk is F15 (Fiber - 15 000 rpm), and the type of RAID is X.

Name of variable	Value
Cache Hit Ratio	0.5 (50%)
Type of Disk I/O Capability	150
RAID Type Weighted I/O Ratio	4
Number of Disks	10

Based on the values that are provided, the read I/O capability per second is:

[0.7 * 150 * 10] / [0.7 * (1 – 0.5) + 0.3 * 4]

The read I/O capability per second rate is 677. (The result of the calculation is rounded to the nearest whole number.)

Related information:

IBM Tivoli Storage Productivity Center Information Center
You can use commands on the Tivoli Storage Productivity Center command-line interface (CLI) to provide information about RAID arrays and back-end storage subsystems. For more information about these commands, see the Tivoli Storage Productivity Center Information Center at http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp. Search for “setarray”, “setbackendtype”, “setbackenddisktype”, and “setbackendraidtype”.

Changing threshold values in storage tier reports

Use the **setdscfg** command to change the threshold values in workload activity reports for managed disk groups and virtual disks.

To change threshold values, you must know the value of the property key that is associated with the name of the threshold in the report.

1. To start the Tivoli Storage Productivity Center CLI, issue the **tpctool** command.
2. Issue the **setdscfg** command, and then set the following values by using the **-property** parameter:
   a. Specify the name of the threshold with the *property_key* variable.
   b. Specify the new threshold value with the *value* variable.

   The following sample code sets the read I/O rate for MDisk groups to 80%:

   ```
   tpctool setdscfg -url localhost:9550 -user ***** -pwd ***** -property
   ManagedDiskGroupAvailIORate 80
   ```
3. Press Enter.

The value of the threshold is changed.

Related reference:

“Threshold values in storage tier reports”
Threshold values are used to assess the workload activity of managed disk (MDisk) groups and virtual disks (VDisks) in storage tier reports. To change a threshold value, you must know the value of the property key that corresponds with the threshold value that is used in the report.

Threshold values in storage tier reports

Threshold values are used to assess the workload activity of managed disk (MDisk) groups and virtual disks (VDisks) in storage tier reports. To change a

threshold value, you must know the value of the property key that corresponds with the threshold value that is used in the report.

The following table contains the name of the threshold value and the corresponding value of the property key for the following reports:

- The MDisk Groups - Workload Activity Report
- The MDisk Group Details Report

Threshold name in report	Property key value
Actual Size Threshold	*ManagedDiskGroupAvailSize*
Read I/O Rate Threshold	*ManagedDiskGroupAvailIORate*
Max Read Response Time Threshold	*ManagedDiskGroupBERResponseExceeds*
Max Write Response Time Threshold	*ManagedDiskGroupBEWResponseExceeds*

The following table contains the name of the threshold value and the corresponding value of the property key for the following reports:

- MDisk Groups - VDisk Workload Activity Report
- MDisk Groups - VDisk Workload Activity Details Report
- VDisk Details Report
- Storage Resource Group - VDisk Workload Activity Report

Threshold name in report	Property key value
VDisk Max Read Response Time Threshold	VDiskFERResponseExceeds
VDisk Max Write Response Time Threshold	VDiskFEWResponseExceeds
Max Read Cache Hits Time Threshold	VDiskFERHitExceeds
Max Read I/O Rate Threshold	VDiskFEIOExceeds
Max Write I/O Rate Threshold	VDiskFEIOWExceeds

Related tasks:

"Changing threshold values in storage tier reports" on page 1094
Use the **setdscfg** command to change the threshold values in workload activity reports for managed disk groups and virtual disks.

Custom reports about performance

You can use Tivoli Common Reporting to create custom performance reports. Performance reports can contain detailed information about the performance of monitored resources, and some of the properties of those resources.

The resources for which you can create performance reports are storage systems and their components, storage virtualizers and their components, and switches and switch ports. You can add statistics for volumes, back-end arrays, caches, ports, port errors, and other statistics to performance reports. For example, you can add key statistics to a report to create and generate a performance chart for a specific volume of a storage system.

All the performance metrics from the stand-alone GUI are available in the new reporting interface.

Creating custom performance reports

You can create reports that show performance metrics. Performance reports can contain detailed information about the performance of monitored resources and some of the properties of those resources. You access custom performance reports from the Tivoli Storage Productivity Center web-based GUI, and you create the reports in Tivoli Common Reporting.

When you create a performance report, include properties and metrics for only one type of device in the report.

The list that you use to create performance reports does not include information about the relationships between storage resources. If you include different types of resources in the report, the report might return incorrect results. For example, if you include storage virtualizers and storage systems in a report, the report might return incorrect results.

Tip: Before you run a report, you can preview the report with no data.

You must specify the time interval for which you want to view performance information. You can specify the time interval in the following ways:

- Select a time unit, for example, day, day of week, or week of month.
- Select performance metrics from sample data, hourly data, or daily data in the resource for which you want performance information. You can also add the time and the interval length to reports. The time is the time at which the performance data was collected. The interval length is the length of the data interval in seconds for sample data, or in minutes for hourly data or daily data.

1. Start the web-based GUI for Tivoli Storage Productivity Center.
2. Depending on the type of report that you want to see, select one of the following items:
   - To create a basic report with limited formatting about the performance of resources, select **Reporting** > **Create performance reports**.
   - To create a report about the performance of resources and apply advanced formatting features to the report, complete the following steps:
     a. Select **Reporting** > **Create performance reports**.
     b. Click the **Return** icon .
     c. Click **Launch** > **Report Studio**.
3. Expand **Tivoli Storage Productivity Center** in the list.
4. Expand **Performance**.
5. Optional: Select a time unit from the **Time** item.
6. Explore the folders in the list to locate the resources, properties, and metrics that you want to include in the report.
7. Depending on your browser, add properties and metrics to the report:
   - In Internet Explorer, drag the items that you want to see information about to the work area.
   - In Firefox, select the items that you want to see information about, and then click **Insert**.
8. Use filters with your performance reports to ensure that your reports do not include too much data. To use filters with your performance reports, do the following steps:
   a. Add a resource or resource type to the report.

b. Select the column heading for the resource or resource type in the work area, and then filter the name of the resource or resource type if necessary.
c. Locate the Sample Data, Hourly Data, or Daily Data folder for the resource or resource type in the list. Add a time stamp from the list to the work area. For example, you can add a day time stamp to the report.
d. Select the column heading for the time stamp in the work area, and then filter the time. If you do not enter a time stamp, the report is run on all the records for this resource or resource type in Tivoli Storage Productivity Center.

9. Click **Save**.

Related information:

IBM Cognos Business Intelligence 10.2.0 Information Center

For more information about how to preview and run reports, see the IBM Cognos Business Intelligence 10.2.0 Information Center at pic.dhe.ibm.com/infocenter/cbi/v10r2m0/index.jsp. Search for "Run a Report".

Data and properties in performance reports

You can include capacity data, properties, and other information about storage systems, storage virtualizers, and switches in performance reports.

Data for storage systems in performance reports

You can include general information, capacity data, properties, and other information about storage systems in performance reports.

Information about storage systems

You can create performance reports that include the following information:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Capacity and usage data

You can create performance reports that include the following information:

Storage System Pool Capacity (GiB)
The amount of storage space in pools that are on the storage system. For an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage System Used Pool Space (GiB)
The amount of space that is in use in all pools on a storage system.

Storage System Available Pool Space (GiB)
The amount of unused space in pools that are on the storage system.

Storage System Volume Capacity (GiB)
The amount of space on all volumes on the storage system.

Storage System Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage System Assigned Volume Space (GiB)
The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage System Unallocated Disk Space (GiB)
The amount of disk space that can be added to a pool.

Storage System Number of Disks
The number of physical disks on the storage system. For a Storwize V7000 storage system that is configured as a back-end device, the value is the number of managed disks on the storage system.

Storage System Number of Volumes
The number of volumes on a resource.

Storage System Cache (GiB)
The size of the cache on the storage system. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

Storage System Real Available Pool Space (GiB)
The amount of unused space in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System Real Configured Pool Space (GiB)
The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System Volume Capacity for z/OS (GiB)
The amount of space on all volumes on the storage system that the z/OS operating system can use.

Storage System Volume Capacity Assigned to MDisks (GiB)
The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

Storage System Disk Capacity (GiB)
The amount of storage space on physical disks in a storage system, excluding spare disks. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

Storage System Total Disk Capacity (GiB)
The amount of space on physical disks on a storage system, including spare disks.

Storage System Physical Allocation Percentage
The percentage of physical space in storage system pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

$(\text{allocated pool space} \div \text{pool capacity}) \times 100$

For example, the physical allocation percentage is 25% for a 200 GiB storage pool. Therefore, the space that is reserved for volumes is 50 GiB.

Storage System Virtual Allocation Percentage
The percentage of physical space in storage system pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(total volume capacity ÷ pool capacity) × 100`

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage System Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

Storage System Shortfall Percentage
The percentage of the remaining unallocated volume space in storage system pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(unallocatable space ÷ (volume space - used volume space)) × 100`

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

Storage System MDisk Available Space (GiB)
The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage System MDisk Capacity (GiB)
The amount of storage space on the managed disk.

Component properties

You can create performance reports that include the following information:

Storage System Access Information
The URL from which you can access the storage system or storage virtualizer.

Storage System Alias
An alternative name for a storage system. The user defines this name.

Storage System Code Level
The Shared Ethernet Adapter level of a DS8000 storage system. For other storage systems, this value is the firmware version.

Storage System Time Zone
The time zone in which a resource is located.

Storage System Custom Tag 1, 2, and 3
User-defined text that is associated with a storage system. You can add or edit the custom tags for a storage system in the **Properties** pane of the storage system in the Tivoli Storage Productivity Center web-based GUI.

Storage System Location
The physical location of a storage system. The location is defined when a storage system is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage system in the **Properties** pane of the storage system in the web-based GUI.

Storage System User Provided Name
The name that was specified for the storage system in the storage environment.

Storage System Is Compression Active
Shows whether the compression feature is enabled on the storage system. If this value is `Yes`, the compression feature is enabled.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage System Status
The condition of the resource, for example normal, warning, or error.

Storage System Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Storage System Last Data Collection
The date and time when storage statistics were last collected from the resource.

Storage System Last Data Collection Status
The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

Vendor, model, and device information

You can create performance reports that include the following information:

Storage System Type
The type of storage system. For example, the storage system can be an **IBM System Storage DS8800** system, an **IBM System Storage DS8700** system, an **IBM System Storage XIV** system, or another type of storage system.

Storage System Vendor
The vendor who supplied the resource.

Storage System Model
The model name or model number of the storage system.

Storage System Machine Type
The machine type of a storage system. For example, the storage system can be **DS8800**, **DS8700**, **XIV - 2812**, or another type of storage system.

Storage System IP Address
The IP address of the resource.

Storage System Serial Number
The serial number of the resource.

Performance metrics for SAN Volume Controller and Storwize systems:

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports, port errors, volumes, caches, back-end arrays, and other data for these systems in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port to Host Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to host computers.

Port to Host Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from host computers.

Total Port to Host I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between host computers and the component.

Port to Disk Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to back-end storage resources.

Port to Disk Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from back-end storage resources.

Total Port to Disk I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

Port to Local Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

Total Port to Local Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

Port to Remote Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

Total Port to Remote Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port to Host Send Data Rate (MiB/s)
The average rate at which data is sent to host computers. The rate is measured in MiB per second.

Port to Host Receive Data Rate (MiB/s)
The average rate at which data is received from host computers. The rate is measured in MiB per second.

Total Port to Host Data Rate (MiB/s)
The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Disk Send Data Rate (MiB/s)
The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

Port to Disk Receive Data Rate (MiB/s)
The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

Total Port to Disk Data Rate (MiB/s)
The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Local Node Send Data Rate (MiB/s)
The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Receive Data Rate (MiB/s)
The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Local Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Send Data Rate (MiB/s)
The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Receive Data Rate (MiB/s)
The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Remote Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Local Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Overall Port to Local Node Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation with another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Local Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.

Overall Port to Local Node Queue Time (ms/op)
The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Remote Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Overall Port to Remote Node Response Time (ms/op)
The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Remote Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote SAN Volume Controller cluster.

Port to Remote Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote SAN Volume Controller cluster.

Overall Port to Remote Node Queue Time (ms/op)
The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote SAN Volume Controller cluster.

Port error data

You can create performance reports that include the following information:

Link Failure Rate (count/s)
The average number of link errors per second.

Loss of Sync Rate (count/s)
The average number of times per second that synchronization is lost.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)
The average number of times per second that the signal is lost.

CRC Error Rate (count/s)
The average number of frames per second that are received in which a

cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

Primitive Sequence Protocol Error Rate (count/s)
The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)
The average number of bit errors per second that are detected.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Write Cache Overflow I/O Rate (ops/s)
The average number of write operations per second that are delayed because there is not enough space in the write cache.

Write Cache Flush Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in flush-through mode.

Write Cache Write Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in write-through mode.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Read Ahead Percentage of Cache Hits
The percentage of all read-cache hits that occur on pre-staged data.

Dirty Write Percentage of Cache Hits
The percentage of all cache write operations that occur on data in the cache

that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

Write Cache Overflow Percentage
The percentage of write operations that are delayed because there is not enough space in the write cache.

Write Cache Flush Through Percentage
The percentage of tracks that is written to disk in flush-through mode.

Write Cache Write Through Percentage
The percentage of tracks that is written to disk in write-through mode.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Peak Back-End Read Response Time (ms)
The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)
The longest time for a back-end storage resource to respond to a write operation.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Read Queue Time (ms)
The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)
The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write I/O Rate (op/s)
The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Peak Read Response Time (ms)
The longest time to complete a read operation.

Peak Write Response Time (ms)
The longest time to complete a write operation.

Global Mirror Secondary Write Lag (ms/op)
The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Global Mirror Overlapping Write Percentage
The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Overall Host Attributed Response Time Percentage
The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

The value is based on the following conditions:

- For read response time, the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes.
- For write response time, the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Zero Buffer-to-Buffer Credit Timer (microseconds)
The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

Performance metrics for DS8000 and DS6000 storage systems:

You can include performance metrics for ports, port errors, volumes, caches, back-end arrays, and other data for DS8000 and DS6000 storage systems in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port PPRC Send I/O Rate (ops/s)
The average number of operations per second that are sent by using the Peer-to-Peer Remote Copy protocol.

Port PPRC Receive I/O Rate (ops/s)
The average number of operations per second that are received by using the Peer-to-Peer Remote Copy protocol.

Total Port PPRC I/O Rate (ops/s)
The average number of send operations and receive operations per second using the Peer-to-Peer Remote Copy protocol.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port PPRC Send Data Rate (MiB/s)
The average number of MiB per second that are sent by using the Peer-to-Peer Remote Copy protocol.

Port PPRC Receive Data Rate (MiB/s)
The average number of MiB per second that are received by using the Peer-to-Peer Remote Copy protocol.

Total Port PPRC Data Rate (MiB/s)
The average number of MiB per second that are transferred by using the Peer-to-Peer Remote Copy (PPRC) protocol. This value includes both send and receive PPRC operations.

Port Send Response Time (ms/op)
The average number of milliseconds to complete a send operation.

Port Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation.

Overall Port Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation.

Port PPRC Send Response Time (ms/op)
The average number of milliseconds to complete a send operation by using the Peer-to-Peer Remote Copy protocol.

Port PPRC Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation by using the Peer-to-Peer Remote Copy protocol.

Overall Port PPRC Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation by using the Peer-to-Peer Remote Copy (PPRC) protocol.

Port Send Transfer Size (KiB/op)
The average number of KiB that are transferred per send operation.

Port Receive Transfer Size (KiB/op)
The average number of KiB that are transferred per receive operation.

Overall Port Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Port error data

You can create performance reports that include the following information:

Link Failure Rate (count/s)
The average number of link errors per second.

Loss of Sync Rate (count/s)
The average number of times per second that synchronization is lost.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)
The average number of times per second that the signal is lost.

CRC Error Rate (count/s)
The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

Primitive Sequence Protocol Error Rate (count/s)
The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)
The average number of bit errors per second that are detected.

Error Frame Rate (count/s)
The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

Link Reset Transmitted Rate (count/s)
The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.

Link Reset Received Rate (count/s)
The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.

Out of Order Data Rate (count/s)
The average number of times per second that an out-of-order frame is detected.

Out of Order ACK Rate (count/s)
The average number of times per second that an out-of-order acknowledge (ACK) frame is detected.

An ACK frame is used for end-to-end flow control and is sent to verify receipt of a frame.

Duplicate Frame Rate (count/s)
The average number of duplicate frames per second that are received. A duplicate frame is a frame that the system previously processed for the port.

Invalid Relative Offset Rate (count/s)
The average number of times per second that frames are received with an invalid relative offset in the frame header.

Sequence Timeout Rate (count/s)
The average number of times per second that the port detects a timeout condition after the port receives a sequence initiative for a Fibre Channel exchange.

Volume data

You can create performance reports that include the following information:

Normal Read I/O Rate (ops/s)
The average number of nonsequential read operations per second.

Sequential Read I/O Rate (ops/s)
The average number of sequential read operations per second.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Normal Write I/O Rate (ops/s)
The average number of nonsequential write operations per second.

Sequential Write I/O Rate (ops/s)
The average number of sequential write operations per second.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Normal I/O Rate (ops/s)
The average number of nonsequential I/O operations per second. This value includes both read and write operations.

Total Sequential I/O Rate (ops/s)
The average number of sequential I/O operations per second. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Record Mode Read I/O Rate (ops/s)
The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Normal Read Cache Hit Percentage
The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Sequential Read Cache Hit Percentage
The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Normal Write Cache Hit Percentage
The percentage of nonsequential write operations that are handled in the cache.

Sequential Write Cache Hit Percentage
The percentage of sequential write operations that are handled in the cache.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Normal Cache Hit Percentage
The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Sequential Cache Hit Percentage
The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Record Mode Read Cache Hit Percentage
The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

Other data

You can create performance reports that include the following information:

HPF Read I/O Rate (ops/s)
The average number of read operations per second that are issued by the High Performance FICON® feature of the storage system.

HPF Write I/O Rate (ops/s)
The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

Total HPF I/O Rate (ops/s)
The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

HPF I/O Percentage
The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

Cache Holding Time (s)
The average number of seconds that I/O data for a storage system controller is held in the cache. A short cache-holding time indicates adverse performance.

PPRC Transfer Rate (ops/s)
The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

Performance metrics for an XIV system:

You can include performance data for volumes, caches, and other data for an XIV system in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Read Data Cache Hit Percentage
The percentage of all read data that is read from the cache.

Write Data Cache Hit Percentage
The percentage of all write data that is written to cache slots that are marked as modified.

Overall Data Cache Hit Percentage
The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

SSD Read Cache Hit Percentage
The percentage of read operations that find data in the cache on a solid-state drive (SSD). The value for this metric is also included in the value for the Overall Read Cache Hit Percentage metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

SSD Read Data Cache Hit Percentage
The percentage of all read data that was read from cache memory on a solid-state drive.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

Read Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a read-cache hit operation.

Write Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a write-cache hit operation.

Overall Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

Read Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a read-cache miss operation.

Write Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a write-cache miss operation.

Overall Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

SSD Read Cache Hit Response Time (ms/op)
The average number of milliseconds that it takes to complete a hit operation on the read cache on a solid-state drive. The value for this metric is also included in the value for the Read Cache Hit Response Time metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

Other data

You can create performance reports that include the following information:

Small Transfers I/O Percentage
The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers I/O Percentage
The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers I/O Percentage
The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers I/O Percentage
The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

Performance metrics for block server systems:

You can include performance metrics for volumes, and cache data for block server storage systems in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Data for storage system volumes in performance reports

You can include general information, capacity data, properties, and other information about storage system volumes in performance reports.

Information about storage system volumes

You can create performance reports that include the following information:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Controller Name
The name that was assigned to the controller when it was added to the system.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage RAID Array Name
The name that was assigned to the RAID array when it was added to the system.

Storage I/O Group Name
The name that was assigned to the I/O group when it was added to the system.

Storage Preferred Node Name
The name of the storage node that is used for I/O operations for a volume.

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Storage Volume Assigned Host Connection
The host connection to which the storage volume is assigned. The host connection is a definition in the storage system that contains the WWPN for the server. The storage system uses the WWPN to assign volumes to servers.

Storage Volume Host Connection OS Type
The operating system type of the server or hypervisor that the volume is assigned to.

Storage System Volume Group ID
The unique identifier for a collection of volumes on DS8000 and DS6000 storage systems. The storage system generates this identifier.

Storage Volume Group Name
The user-defined name for a collection of volumes on DS8000 and DS6000 storage systems.

Server Name
The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Hypervisor Name
The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

Hypervisor Cluster Name
The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

Capacity and usage data

You can create performance reports that include the following information:

Storage Volume Capacity (GiB)
The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

For XIV systems, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage Volume Used Space (GiB)
The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Volume Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Volume Allocated Space (GiB)
The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

Storage Volume Unallocated Space (GiB)
The amount of space in a pool that is not reserved for a volume.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
capacity - allocated space
```

Available only for thin provisioned volumes.

Storage Volume Physical Allocation Percentage
The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
(allocated space ÷ capacity) × 100
```

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Volume Shortfall Percentage
The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume.

The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
(unallocatable volume Space ÷ volume unallocated space) × 100
```

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This Storage Volume Shortfall Percentage represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

```
100 GiB unallocated volume space - 70 GiB remaining
pool space ÷ 100 GiB unallocated volume space × 100
```

Storage Volume Unallocatable Space (GiB)
The unallocatable storage space in GiB of a thin-provisioned volume. Unallocatable space cannot be supplied by the pool to which the volume belongs.

Storage Volume Used Allocated Space Percentage
The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
(volume used space ÷ volume allocated space) × 100
```

Storage Volume Unused Space (GiB)
The amount of space that is allocated to a volume and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
Allocated Space - Used Space
```

This property is available only for SAN Volume Controller and Storwize V7000 pools.

Storage Volume Uncompressed Capacity (GiB)
The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Component properties

You can create performance reports that include the following information:

Storage Volume RAID Level
The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Volume Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Volume Is Assigned
Shows whether the volume is assigned to a server. If this value is `Yes`, the volume is assigned to a server.

Storage Volume Is Encrypted
Shows whether the resource is encrypted. If this value is `Yes`, the resource is encrypted.

Storage Volume Is Encryptable
Shows whether the resource can be encrypted. If this value is `Yes`, the resource can be encrypted.

Storage Volume Is Volume Control Manager Database
Shows whether a volume is used as the Volume Control Manager Database. If this value is `Yes`, the volume is used as the Volume Control Manager Database.

Storage Volume Format
Shows the format of the volume or pool. The format can be a Count Key Data (CKD) format, or a fixed block format.

Storage Volume Serial Number
The serial number of the resource.

Storage Volume Number
The number of the volume that is assigned by the system.

Storage Volume WWN
The worldwide name of the volume.

Storage Volume FlashCopy Relationship
Shows whether a volume on a storage system is in a FlashCopy relationship. This property can contain the following values:

Source Relationship
The volume is the source of the relationship.

Target Relationship
The volume is the target of the relationship.

Both Source and Target Relationship
The volume is both the source and target of the relationship.

Not in a FlashCopy Relationship
The volume is not in a FlashCopy relationship.

Storage Volume Remote Copy Relationship
Shows whether a volume on a storage system is in a remote copy relationship. This property can contain the following values:

Not in a Copy Relationship
The volume is not in a remote copy relationship.

Source Relationship
The volume is the source of the relationship.

Target Relationship
The volume is the target of the relationship.

Both Source and Target Relationship
The volume is both the source and target of the relationship.

Storage Volume Logical Subsystem
The logical subsystem (LSS) to which a volume or pool belongs.

Storage Volume Type
The type of storage volume. For example, the storage volume can be striped or sequential. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

Storage Volume Fast Write State
Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted.

Storage Volume Grain Size (KiB)
The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

Storage Volume Is Auto Expand Enabled
Shows whether a volume can automatically allocate new extents from a pool. Volumes might allocate new extents to expand the real capacity of the volume. If this value is `Yes`, the volume can automatically allocate new extents from a pool.

This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

Storage Volume Metro Mirror Name
The name of the Metro Mirror that keeps the synchronous copy of the storage volume. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

Storage Volume Mirror Count
The number of mirrors that keep a synchronous copy of the resource. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

Storage Volume Throttle
The maximum number of commands that the volume can queue. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

Storage Volume Warning Level
The percentage of volume capacity that is used at which a warning is generated. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

Storage Volume Is Compressed
Shows whether the volume is compressed. If this value is `Yes`, the volume is compressed.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage Volume Status
The condition of the resource, for example normal, warning, or error.

Storage Volume Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Storage Volume Native Status
The condition of the volume. The status can be online, offline, degraded, or excluded.

Performance metrics for volumes on SAN Volume Controller and Storwize systems:

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for volumes, caches, and other data for volumes on these systems in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Maximum Read I/O Rate (ops/s)
The maximum read I/O rate of the volume in operations per second. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Write I/O Rate (ops/s)
The maximum write I/O rate of the volume in operations per second. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Write I/O Size (KiB)
The maximum size of write I/O operations in kibibytes. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Maximum Read Response Time (ms/op)
The maximum read response time of the volume in milliseconds per operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Write Response Time (ms/op)
The maximum write response time of the volume in milliseconds per operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Volume Utilization
The average percentage of time that the volume is busy.

Maximum Volume Utilization Percentage
The maximum period in percentage that the volume is busy. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Cache data

You can create performance reports that include the following information:

Write Cache Overflow I/O Rate (ops/s)
The average number of write operations per second that are delayed because there is not enough space in the write cache.

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Write Cache Flush Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in flush-through mode.

Write Cache Write Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in write-through mode.

Maximum Write Cache Delay I/O Rate (ops/s)
The maximum number of delayed I/O operations per second for the volume, where the delay occurred because of space constraints or other conditions.

This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Read Ahead Percentage of Cache Hits
The percentage of all read-cache hits that occur on pre-staged data.

Dirty Write Percentage of Cache Hits
The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

Write Cache Overflow Percentage
The percentage of write operations that are delayed because there is not enough space in the write cache.

Write Cache Flush Through Percentage
The percentage of tracks that is written to disk in flush-through mode.

Write Cache Write Through Percentage
The percentage of tracks that is written to disk in write-through mode.

Maximum Read Cache Hit Percentage
The maximum number of cache hits in percentage for sequential and nonsequential read operations for the volume over a specified interval. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)
The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write Percentage
The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Global Mirror Overlapping Write I/O Rate (op/s)
The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Peak Read Response Time (ms)
The longest time to complete a read operation.

Peak Write Response Time (ms)
The longest time to complete a write operation.

Global Mirror Secondary Write Lag (ms/op)
The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

Overall Host Attributed Response Time Percentage
The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

The value is based on the following conditions:

- For read response time, the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes.
- For write response time, the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Nonpreferred Node Usage Percentage
The overall percentage of I/O operations that are completed, or data that is transferred over the non-preferred nodes of the virtual disks.

Performance metrics for volumes on DS8000 and DS6000 storage systems:

You can include performance data for volumes, caches, and other data for storage system volumes on DS8000 and DS6000 storage systems in performance reports.

Volume data

You can create performance reports that include the following information:

Normal Read I/O Rate (ops/s)
The average number of nonsequential read operations per second.

Sequential Read I/O Rate (ops/s)
The average number of sequential read operations per second.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Normal Write I/O Rate (ops/s)
The average number of nonsequential write operations per second.

Sequential Write I/O Rate (ops/s)
The average number of sequential write operations per second.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Normal I/O Rate (ops/s)
The average number of nonsequential I/O operations per second. This value includes both read and write operations.

Total Sequential I/O Rate (ops/s)
The average number of sequential I/O operations per second. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Record Mode Read I/O Rate (ops/s)
The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Volume Utilization
The average percentage of time that the volume is busy.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Normal Read Cache Hit Percentage
The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Sequential Read Cache Hit Percentage
The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Normal Write Cache Hit Percentage
The percentage of nonsequential write operations that are handled in the cache.

Sequential Write Cache Hit Percentage
The percentage of sequential write operations that are handled in the cache.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Normal Cache Hit Percentage
The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Sequential Cache Hit Percentage
The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Record Mode Read Cache Hit Percentage
The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

Other data

You can create performance reports that include the following information:

HPF Read I/O Rate (ops/s)
The average number of read operations per second that are issued by the High Performance FICON feature of the storage system.

HPF Write I/O Rate (ops/s)
The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

Total HPF I/O Rate (ops/s)
The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

HPF I/O Percentage
The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

PPRC Transfer Rate (ops/s)
The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

Performance metrics for storage system volumes on an XIV system:

You can include performance data for volumes, caches, and other data for storage system volumes on an XIV system in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Volume Utilization
The average percentage of time that the volume is busy.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Read Data Cache Hit Percentage
The percentage of all read data that is read from the cache.

Write Data Cache Hit Percentage
The percentage of all write data that is written to cache slots that are marked as modified.

Overall Data Cache Hit Percentage
The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

Read Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a read-cache miss operation.

Write Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a write-cache miss operation.

Overall Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

Read Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a read-cache hit operation.

Write Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a write-cache hit operation.

Overall Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

SSD Read Cache Hit Percentage
The percentage of read operations that find data in the cache on a solid-state drive (SSD). The value for this metric is also included in the value for the Overall Read Cache Hit Percentage metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

SSD Read Data Cache Hit Percentage
The percentage of all read data that was read from cache memory on a solid-state drive.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

SSD Read Cache Hit Response Time (ms/op)
The average number of milliseconds that it takes to complete a hit operation on the read cache on a solid-state drive. The value for this metric is also included in the value for the Read Cache Hit Response Time metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

Other data

You can create performance reports that include the following information:

Small Transfers I/O Percentage
The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers I/O Percentage
The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers I/O Percentage
The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers I/O Percentage
The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

Performance metrics for storage system volumes on block servers:

You can include performance data for volumes and caches for storage system volumes on block servers in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Volume Utilization
The average percentage of time that the volume is busy.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Data for storage system ports in performance reports

You can include general information, properties, and other information about storage system ports in performance reports.

Information about storage system ports

You can create performance reports that include the following information:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Module Name
The name that was assigned to the module when it was added to the system.

Storage Port Name
The name that was assigned to the storage port when the storage system was added to the system.

Component properties

You can create performance reports that include the following information:

Storage System Port Number
The port number on the resource.

Storage System Port Type
The type of port on the storage system, storage virtualizer, or switch. For example, the port type can be N_Port, F_Port, or another type of port.

Storage System Port Speed (GiB/s)
The speed of a port, which is measured in GiB per second.

Storage System Port Location
The physical location of a port on DS8000, DS6000, TotalStorage Enterprise Storage Server storage systems, and on XIV systems. For all other systems, this property shows the WWPN of the port.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage System Port Status
The condition of the resource, for example normal, warning, or error.

Storage System Port Enabled State
Shows whether a port is enabled, disabled, or is enabled but offline.

Storage System Port Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Performance metrics for ports on SAN Volume Controller and Storwize systems:

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports and port errors on these systems in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port to Host Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to host computers.

Port to Host Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from host computers.

Total Port to Host I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between host computers and the component.

Port to Disk Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to back-end storage resources.

Port to Disk Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from back-end storage resources.

Total Port to Disk I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

Port to Local Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

Total Port to Local Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

Port to Remote Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

Total Port to Remote Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port to Host Send Data Rate (MiB/s)
The average rate at which data is sent to host computers. The rate is measured in MiB per second.

Port to Host Receive Data Rate (MiB/s)
The average rate at which data is received from host computers. The rate is measured in MiB per second.

Total Port to Host Data Rate (MiB/s)
The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Disk Send Data Rate (MiB/s)
The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

Port to Disk Receive Data Rate (MiB/s)
The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

Total Port to Disk Data Rate (MiB/s)
The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Local Node Send Data Rate (MiB/s)
The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Receive Data Rate (MiB/s)
The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Local Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Send Data Rate (MiB/s)
The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Receive Data Rate (MiB/s)
The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Remote Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port Send Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage
The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port error data

You can create performance reports that include the following information:

Link Failure Rate (count/s)
The average number of link errors per second.

Loss of Sync Rate (count/s)
The average number of times per second that synchronization is lost.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)
The average number of times per second that the signal is lost.

CRC Error Rate (count/s)
The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

Primitive Sequence Protocol Error Rate (count/s)
The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)
The average number of bit errors per second that are detected.

Zero Buffer-to-Buffer Credit Timer (microseconds)
The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

Zero Buffer-to-Buffer Credit Percentage
The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

Performance metrics for ports on DS8000 and DS6000 storage systems:

You can include performance data for storage system port and port errors on DS8000 and DS6000 storage systems in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port FICON Send I/O Rate (ops/s)
The average number of send operations per second for Fibre Channel connection (FICON) usage.

Port FICON Receive I/O Rate (ops/s)
The average number of receive operations per second for FICON usage.

Total Port FICON I/O Rate (ops/s)
The average number of send and receive operations per second for FICON usage.

Port FCP Send I/O Rate (ops/s)
The average number of send operations per second for Fibre Channel Protocol (FCP) usage.

Port FCP Receive I/O Rate (ops/s)
The average number of receive operations per second for FCP usage.

Total Port FCP I/O Rate (ops/s)
The average number of send operations and receive operations per second for FCP usage.

Port PPRC Send I/O Rate (ops/s)
The average number of operations per second that are sent by using the Peer-to-Peer Remote Copy protocol.

Port PPRC Receive I/O Rate (ops/s)
The average number of operations per second that are received by using the Peer-to-Peer Remote Copy protocol.

Total Port PPRC I/O Rate (ops/s)
The average number of send operations and receive operations per second using the Peer-to-Peer Remote Copy protocol.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port FICON Send Data Rate (MiB/s)
The average number of MiB per second that is sent for FICON usage.

Port FICON Receive Data Rate (MiB/s)
The average number of MiB per second that is received for FICON usage.

Total Port FICON Data Rate (MiB/s)
The average number of MiB per second that is transferred for FICON usage. This value includes both send and receive FICON operations.

Port FCP Send Data Rate (MiB/s)
The average number of MiB per second that are sent for FCP usage.

Port FCP Receive Data Rate (MiB/s)
The average number of MiB per second that are received for FCP usage.

Total Port FCP Data Rate (MiB/s)
The average number of MiB per second that are transferred for FCP usage. This value includes both send and receive FCP operations.

Port PPRC Send Data Rate (MiB/s)
The average number of MiB per second that are sent by using the Peer-to-Peer Remote Copy protocol.

Port PPRC Receive Data Rate (MiB/s)
The average number of MiB per second that are received by using the Peer-to-Peer Remote Copy protocol.

Total Port PPRC Data Rate (MiB/s)
The average number of MiB per second that are transferred by using the Peer-to-Peer Remote Copy (PPRC) protocol. This value includes both send and receive PPRC operations.

Port Send Response Time (ms/op)
The average number of milliseconds to complete a send operation.

Port Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation.

Overall Port Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation.

Port FCP Send Response Time (ms/op)
The average number of milliseconds to complete a send operation for FCP usage.

Port FCP Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation for FCP usage.

Overall Port FCP Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation for FCP usage.

Port FICON Send Response Time (ms/op)
The average number of milliseconds to complete a send operation for FICON usage.

Port FICON Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation for FICON usage.

Overall Port FICON Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation for FICON usage. This value includes both send and receive FICON operations.

Port PPRC Send Response Time (ms/op)
The average number of milliseconds to complete a send operation by using the Peer-to-Peer Remote Copy protocol.

Port PPRC Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation by using the Peer-to-Peer Remote Copy protocol.

Overall Port PPRC Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation by using the Peer-to-Peer Remote Copy (PPRC) protocol.

Port Send Transfer Size (KiB/op)
The average number of KiB that are transferred per send operation.

Port Receive Transfer Size (KiB/op)
The average number of KiB that are transferred per receive operation.

Overall Port Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Port Send Utilization Percentage
The average percentage of time that the port is busy sending data.

Port Receive Utilization Percentage
The average percentage of time that the port is busy receiving data.

Overall Port Utilization Percentage
The average percentage of time that the port is busy sending or receiving data.

Port Send Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage
The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port error data

You can create performance reports that include the following information:

Error Frame Rate (count/s)
The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

Link Failure Rate (count/s)
The average number of link errors per second.

Loss of Sync Rate (count/s)
The average number of times per second that synchronization is lost.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)
The average number of times per second that the signal is lost.

CRC Error Rate (count/s)
The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

Primitive Sequence Protocol Error Rate (count/s)
The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)
The average number of bit errors per second that are detected.

Out of Order Data Rate (count/s)
The average number of times per second that an out-of-order frame is detected.

Out of Order ACK Rate (count/s)
The average number of times per second that an out-of-order acknowledge (ACK) frame is detected.

An ACK frame is used for end-to-end flow control and is sent to verify receipt of a frame.

Duplicate Frame Rate (count/s)
The average number of duplicate frames per second that are received. A duplicate frame is a frame that the system previously processed for the port.

Invalid Relative Offset Rate (count/s)
The average number of times per second that frames are received with an invalid relative offset in the frame header.

Sequence Timeout Rate (count/s)
The average number of times per second that the port detects a timeout condition after the port receives a sequence initiative for a Fibre Channel exchange.

Link Reset Transmitted Rate (count/s)
The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.

Link Reset Received Rate (count/s)
The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.

Performance metrics for ports on an XIV system:

You can include performance metrics for storage system ports on an XIV system in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Send Response Time (ms/op)
The average number of milliseconds to complete a send operation.

Port Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation.

Overall Port Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation.

Port Send Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage
The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Performance metrics for storage system ports on block servers:

You can include performance metrics for storage system ports on block servers in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Send Response Time (ms/op)
The average number of milliseconds to complete a send operation.

Port Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation.

Overall Port Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation.

Port Send Transfer Size (KiB/op)
The average number of KiB that are transferred per send operation.

Port Receive Transfer Size (KiB/op)
The average number of KiB that are transferred per receive operation.

Overall Port Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Data for storage system RAID arrays in performance reports

You can include general information, capacity data, and properties of storage system RAID arrays in performance reports.

Information about storage system RAID arrays

You can create performance reports that include the following information:

Storage System Name
A user-defined name of the storage system. If a name was not defined,

Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Controller Name
The name that was assigned to the controller when it was added to the system.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage RAID Array Name
The name that was assigned to the RAID array when it was added to the system.

Capacity and usage data

You can create performance reports that include the following information:

Storage RAID Array DDM Capacity (GiB)
The amount of storage space on the disk drive modules in an array.

Component properties

You can create performance reports that include the following information:

Storage RAID Array RAID Level
The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage RAID Array Width
The number of disks that are in the storage array.

Storage RAID Array DDM Speed (RPM)
The speed of the disk drive module in the array.

Storage RAID Array Description
Shows the array site that the array is associated with.

Storage RAID Array Tag
A number that identifies an array.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Performance metrics for storage system arrays:

You can include front end data, cache data, back-end data, and other data for storage system arrays in performance reports.

Front end data

You can create performance reports that include the following information:

Normal Read I/O Rate (ops/s)
The average number of nonsequential read operations per second.

Sequential Read I/O Rate (ops/s)
The average number of sequential read operations per second.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Normal Write I/O Rate (ops/s)
The average number of nonsequential write operations per second.

Sequential Write I/O Rate (ops/s)
The average number of sequential write operations per second.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Normal I/O Rate (ops/s)
The average number of nonsequential I/O operations per second. This value includes both read and write operations.

Total Sequential I/O Rate (ops/s)
The average number of sequential I/O operations per second. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Record Mode Read I/O Rate (ops/s)
The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Normal Read Cache Hit Percentage
The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Sequential Read Cache Hit Percentage
The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Normal Write Cache Hit Percentage
The percentage of nonsequential write operations that are handled in the cache.

Sequential Write Cache Hit Percentage
The percentage of sequential write operations that are handled in the cache.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Normal Cache Hit Percentage
The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Sequential Cache Hit Percentage
The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Record Mode Read Cache Hit Percentage
The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Back-end data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Sequential Back-End I/O Percentage
The percentage of all I/O operations that are issued to the back-end storage resources. The operations are for the array and are sequential operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Other data

You can create performance reports that include the following information:

HPF Read I/O Rate (ops/s)
The average number of read operations per second that are issued by the High Performance FICON feature of the storage system.

HPF Write I/O Rate (ops/s)
The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

Total HPF I/O Rate (ops/s)
The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

HPF I/O Percentage
The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

PPRC Transfer Rate (ops/s)
The average number of tracks per second that are transferred to the

secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

Disk Utilization Percentage
The average percentage of time that the disks that are associated with an array are busy.

No value is calculated for this property if there are multiple ranks in the extent pool where the thin-provisioned volumes are allocated. In this case, the value `N/A` is displayed for this property. This limitation applies only to DS8000 storage systems.

If there is only a single rank in the extent pool, the value for this property is calculated regardless of the thin-provisioned volumes.

Data for storage system controllers in performance reports

You can include general information about storage system controllers in performance reports.

Information about storage system controllers

You can create performance reports that include the following information:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Controller Name
The name that was assigned to the controller when it was added to the system.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Performance metrics for storage system controllers on DS8000 and DS6000 systems:

You can include front end data, cache data, back-end data, and other data for storage system controllers on DS8000 and DS6000 series systems in performance reports.

Front end data

You can create performance reports that include the following information:

Normal Read I/O Rate (ops/s)
The average number of nonsequential read operations per second.

Sequential Read I/O Rate (ops/s)
The average number of sequential read operations per second.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Normal Write I/O Rate (ops/s)
The average number of nonsequential write operations per second.

Sequential Write I/O Rate (ops/s)
The average number of sequential write operations per second.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Normal I/O Rate (ops/s)
The average number of nonsequential I/O operations per second. This value includes both read and write operations.

Total Sequential I/O Rate (ops/s)
The average number of sequential I/O operations per second. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Record Mode Read I/O Rate (ops/s)
The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Normal Read Cache Hit Percentage
The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Sequential Read Cache Hit Percentage
The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Normal Write Cache Hit Percentage
The percentage of nonsequential write operations that are handled in the cache.

Sequential Write Cache Hit Percentage
The percentage of sequential write operations that are handled in the cache.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Normal Cache Hit Percentage
The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Sequential Cache Hit Percentage
The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Record Mode Read Cache Hit Percentage
The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Back-end data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Other data

You can create performance reports that include the following information:

HPF Read I/O Rate (ops/s)
The average number of read operations per second that are issued by the High Performance FICON feature of the storage system.

HPF Write I/O Rate (ops/s)
The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

Total HPF I/O Rate (ops/s)
The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

HPF I/O Percentage
The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

PPRC Transfer Rate (ops/s)
The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

Cache Holding Time (s)
The average number of seconds that I/O data for a storage system controller is held in the cache. A short cache-holding time indicates adverse performance.

Performance metrics for storage system controllers on block servers:

You can include performance data for volumes and caches for storage system controllers on block servers in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Data for storage system modules in performance reports

You can include properties of storage system modules, and other information about storage system modules in performance reports.

Information about storage system modules

You can create performance reports that include the following information:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Module Name
The name that was assigned to the module when it was added to the system.

Component properties

You can create performance reports that include the following information:

Storage Module IP Address
The IP address of the resource.

Storage Module WWN
The worldwide name of the module.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage Module Status
The condition of the resource, for example normal, warning, or error.

Storage Module Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Performance metrics for storage system modules on an XIV system:

You can include performance metrics for volumes, caches, and other data for storage system modules on an XIV system in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Read Data Cache Hit Percentage
The percentage of all read data that is read from the cache.

Write Data Cache Hit Percentage
The percentage of all write data that is written to cache slots that are marked as modified.

Overall Data Cache Hit Percentage
The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

Read Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a read-cache hit operation.

Write Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a write-cache hit operation.

Overall Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

Read Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a read-cache miss operation.

Write Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a write-cache miss operation.

Overall Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

SSD Read Cache Hit Percentage
The percentage of read operations that find data in the cache on a solid-state drive (SSD). The value for this metric is also included in the value for the Overall Read Cache Hit Percentage metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

SSD Read Data Cache Hit Percentage
The percentage of all read data that was read from cache memory on a solid-state drive.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

SSD Read Cache Hit Response Time (ms/op)
The average number of milliseconds that it takes to complete a hit operation on the read cache on a solid-state drive. The value for this metric is also included in the value for the Read Cache Hit Response Time metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

Other data

You can create performance reports that include the following information:

Small Transfers I/O Percentage
The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers I/O Percentage
The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers I/O Percentage
The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers I/O Percentage
The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has

a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

Data for storage system pools in performance reports

You can include general information, capacity data, properties, and other information about storage system pools in performance reports.

Information about storage system pools

You can create performance reports that include the following information:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Capacity and usage data

You can create performance reports that include the following information:

Storage Pool Capacity (GiB)
The total amount of storage space in a pool.

Storage Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
pool capacity - allocated space
```

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

Storage Pool Total Volume Capacity (GiB)
The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Pool Allocated Space (GiB)
The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

Storage Pool Used Volume Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Pool Extent Size (MiB)
The size of the extent that was specified when a pool was created. Smaller extent sizes limit the maximum size of the volumes that can be created in a pool. Smaller extent sizes minimize the amount of potentially wasted space per volume.

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage Pool Assigned Volume Space (GiB)
The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Pool Unassigned Volume Space (GiB)
The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Pool Number of Storage Volumes
The number of volumes in the storage pool.

Storage Pool Number of MDisks
The number of managed disks in a storage pool.

Storage Pool Physical Allocation Percentage
The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

`(allocated space ÷ pool capacity) × 100`

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Pool Virtual Allocation Percentage
The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

`(total volume capacity ÷ pool capacity) × 100`

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pool is 30 GiB.

This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, you can use the value for Storage Pool Shortfall Percentage to estimate how critical the shortage of space is for a pool.

This value is only available for pools with thin-provisioned volumes.

Storage Pool Shortfall Percentage
The percentage of the remaining unallocated volume space in a pool that is not available to be allocated.

The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(unallocatable space ÷ (virtual capacity − allocated space)) × 100`

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

Storage Pool Unallocated Volume Space (GiB)
The amount of the Total Volume Capacity in the pool that is not allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

`total volume capacity - allocated space`

The space that is allocated for thin-provisioned volumes is typically less than their virtual capacity. Therefore, the unallocated space represents the difference between the virtual capacity and the allocated space for all the volumes in the pool.

Storage Pool Unused Volume Space (GiB)
The amount of space that is allocated to the volumes in a pool and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

```
allocated space - used space
```

This property is available only for SAN Volume Controller and Storwize V7000 pools.

Storage Pool Unallocatable Volume Space (GiB)
The amount of space by which the Total Volume Capacity exceeds the physical capacity of a pool.

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool. If you create volumes with more virtual capacity than can be physically allocated in the pool, you can over commit storage in the pool.

This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.

Storage Pool Used Space Percentage
The percentage of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

Storage Pool Compressed Virtual Capacity (GiB)
The total virtual capacity of all the volumes that are compressed in a pool.

Storage Pool Compressed Capacity (GiB)
The amount of storage space that is used by compressed volumes in a pool. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 20.

Storage Pool Uncompressed Capacity (GiB)
The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Storage Pool Maximum I/O Capability
The projected maximum number of I/O operations per second for a pool. This value is calculated based on the value in the Storage Pool Back-End Storage Disks property, and on the values in following properties:

- Storage Pool Back-End Storage System Type
- Storage Pool Back-End Storage RAID Level
- Storage Pool Back-End Storage Disk Type

This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Storage Pool Back-End Storage Disks
The number of physical disks that contribute to the volumes on the back-end storage system. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Component properties

You can create performance reports that include the following information:

Storage Pool Is Encryptable
Shows whether the storage pool can be encrypted. If this value is `Yes`, the storage pool can be encrypted.

Storage Pool Is Encrypted
Shows whether the resource is encrypted. If this value is `Yes`, the resource is encrypted.

Storage Pool Is Solid State
Shows whether there are solid-state drives in the pool. This property can contain the following values:

Mixed The pool contains both hard disk drives and solid-state drives.

Non solid state
The pool contains no solid-state drives.

Solid state
The pool contains a least one solid-state drive.

Storage Pool Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Pool RAID Level
The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Pool Warning Level
The percentage of used capacity of the storage pool at which a warning is generated.

Storage Pool Is Compression Active
Shows whether the compression feature is enabled on the storage pool. If this value is `Yes`, the compression feature is enabled.

Storage Pool Custom Tag 1, 2, and 3
User-defined text that is associated with a storage pool. You can add or edit the custom tags for a storage pool on the **Properties** notebook of the pool in the Tivoli Storage Productivity Center web-based GUI.

Storage Pool Back-End Storage System Type
The type of storage system that provides storage space to a pool. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Storage Pool Back-End Storage RAID Level
The RAID level of the volumes on the back-end storage system that provide storage space to a pool. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Storage Pool Back-End Storage Disk Type
The class and speed of the physical disks that contribute to the volumes on the back-end storage system. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage Pool Status
The condition of the resource, for example normal, warning, or error.

Storage Pool Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Performance metrics for pools on SAN Volume Controller and Storwize systems:

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for volumes and back-end arrays for pools on these systems in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Maximum Read I/O Rate (ops/s)
The maximum read I/O rate of the volume in operations per second. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Pool Utilization
The average percentage of time that the volumes in the pool are busy.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Maximum Back-End Read I/O Rate (ops/s)
The maximum number of read operations per second that are issued to the back-end storage resources. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Back-End Write I/O Rate (ops/s)
The maximum number of write operations per second that are issued to the back-end storage resources. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Peak Back-End Read Response Time (ms)
The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)
The longest time for a back-end storage resource to respond to a write operation.

Maximum Back-End Read Response Time (ms/op)
The maximum number of milliseconds for the back-end storage resources to respond to a read operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Maximum Back-End Write Response Time (ms/op)
The maximum number of milliseconds for the back-end storage resources to respond to a write operation. This metric is only available for daily data for SAN Volume Controller and Storwize V7000 systems.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Read Queue Time (ms)
The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Performance metrics for pools on DS8000 and DS6000 storage systems:

You can include performance metrics for volumes, caches, back-end arrays, and other data for pools that are on DS8000 and DS6000 storage systems in performance reports.

Volume data

You can create performance reports that include the following information:

Normal Read I/O Rate (ops/s)
The average number of nonsequential read operations per second.

Sequential Read I/O Rate (ops/s)
The average number of sequential read operations per second.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Normal Write I/O Rate (ops/s)
The average number of nonsequential write operations per second.

Sequential Write I/O Rate (ops/s)
The average number of sequential write operations per second.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Normal I/O Rate (ops/s)
The average number of nonsequential I/O operations per second. This value includes both read and write operations.

Total Sequential I/O Rate (ops/s)
The average number of sequential I/O operations per second. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Record Mode Read I/O Rate (ops/s)
The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Normal Read Cache Hit Percentage
The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Sequential Read Cache Hit Percentage
The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Normal Write Cache Hit Percentage
The percentage of nonsequential write operations that are handled in the cache.

Sequential Write Cache Hit Percentage
The percentage of sequential write operations that are handled in the cache.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Normal Cache Hit Percentage
The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Sequential Cache Hit Percentage
The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Record Mode Read Cache Hit Percentage
The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

Other data

You can create performance reports that include the following information:

HPF Read I/O Rate (ops/s)
The average number of read operations per second that are issued by the High Performance FICON feature of the storage system.

HPF Write I/O Rate (ops/s)
The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

Total HPF I/O Rate (ops/s)
The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

HPF I/O Percentage
The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

PPRC Transfer Rate (ops/s)
The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

Performance metrics for pools on an XIV system:

You can include performance data for volumes, caches, and other data for pools on an XIV system in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Pool Utilization
The average percentage of time that the volumes in the pool are busy.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Read Data Cache Hit Percentage
The percentage of all read data that is read from the cache.

Write Data Cache Hit Percentage
The percentage of all write data that is written to cache slots that are marked as modified.

Overall Data Cache Hit Percentage
The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

Read Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a read-cache hit operation.

Write Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a write-cache hit operation.

Overall Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

Read Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a read-cache miss operation.

Write Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a write-cache miss operation.

Overall Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

Other data

You can create performance reports that include the following information:

Small Transfers I/O Percentage
The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers I/O Percentage
The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers I/O Percentage
The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers I/O Percentage
The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has

a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

Data for managed disks on storage systems in performance reports

You can include general information, capacity data, properties, and other information about managed disks on storage systems in performance reports.

Information about managed disks on storage systems

You can create performance reports that include the following information:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage System MDisk Name
The name that was assigned to the managed disk on a storage system when it was added to the system.

Capacity and usage data

You can create performance reports that include the following information:

Storage System MDisk Capacity (GiB)
The amount of storage space on the managed disk.

Storage MDisk Available Space (GiB)
The amount of unused storage space on the managed disk on a storage system.

Storage System MDisk Strip Size (KB)
The RAID strip size on a managed disk on a storage system.

Component properties

You can create performance reports that include the following information:

Storage MDisk RAID Level
The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

Storage MDisk Spare Goal
The number of spare drives that are required to maintain redundancy. Use spare drives to protect the system against drive failures in the array on a managed disk on a storage virtualizer.

Storage System MDisk Type
The type of managed disk on a storage system. For example, the disk on a storage system can be a local managed disk.

Storage System MDisk Mode
The access mode of a managed disk on a storage virtualizer. The access

mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

Array Extents are provided from local disks.

Managed
Extents are provided from a back-end storage volume.

Unmanaged
The managed disk is not used in the system.

Storage System MDisk Is Balanced
Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is `Yes`, the LUNs are balanced.

Storage System MDisk Fast Write State
Shows whether the cache for a volume on a disk that is managed by a storage system is empty, contains data, or is corrupted.

Storage System MDisk Write Verify
Shows whether all write operations on a managed disk on a storage system are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is `Yes`, all write operations are verified by a follow-up read operation.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage MDisk Status
The condition of the resource, for example normal, warning, or error.

Storage MDisk Native Status
Shows the level of access of nodes in the system to a managed disk or volume on a storage system. The level of access can be online, offline, degraded, or excluded.

Performance metrics for managed disks on SAN Volume Controller and Storwize systems:

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for back-end arrays on managed disks on these systems in performance reports.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Peak Back-End Read Response Time (ms)
The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)
The longest time for a back-end storage resource to respond to a write operation.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Read Queue Time (ms)
The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Data for storage system nodes in performance reports

You can include general information, capacity data, properties, and other information about storage system nodes in performance reports.

Information about storage system nodes

You can create performance reports that include the following information:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Node Name
The name that was assigned to the storage system node when it was added to the system.

Storage I/O Group Name
The name that was assigned to the I/O group when it was added to the system.

Component properties

You can create performance reports that include the following information:

Storage Node IP Address
The IP address of the resource.

Storage Node WWN
The worldwide name of the node.

Storage Node Configuration
The number of the node for the resource. This value can be either `0` or `1`.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage Node Status
Shows whether a storage system node is online.

Performance metrics for nodes on SAN Volume Controller and Storwize systems:

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports, front ends, caches, back-end arrays, and other data for nodes on these systems in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port to Host Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to host computers.

Port to Host Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from host computers.

Total Port to Host I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between host computers and the component.

Port to Disk Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to back-end storage resources.

Port to Disk Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from back-end storage resources.

Total Port to Disk I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

Port to Local Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

Total Port to Local Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

Port to Remote Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

Total Port to Remote Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port to Host Send Data Rate (MiB/s)
The average rate at which data is sent to host computers. The rate is measured in MiB per second.

Port to Host Receive Data Rate (MiB/s)
The average rate at which data is received from host computers. The rate is measured in MiB per second.

Total Port to Host Data Rate (MiB/s)
The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Disk Send Data Rate (MiB/s)
The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

Port to Disk Receive Data Rate (MiB/s)
The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

Total Port to Disk Data Rate (MiB/s)
The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Local Node Send Data Rate (MiB/s)
The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Receive Data Rate (MiB/s)
The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Local Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Send Data Rate (MiB/s)
The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Receive Data Rate (MiB/s)
The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Remote Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Local Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Overall Port to Local Node Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation with another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Remote Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Remote Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Overall Port to Remote Node Response Time (ms/op)
The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Local Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the

queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.

Overall Port to Local Node Queue Time (ms/op)
The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote SAN Volume Controller cluster.

Port to Remote Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote SAN Volume Controller cluster.

Overall Port to Remote Node Queue Time (ms/op)
The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote SAN Volume Controller cluster.

Front end data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Write Cache Overflow I/O Rate (ops/s)
The average number of write operations per second that are delayed because there is not enough space in the write cache.

Write Cache Flush Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in flush-through mode.

Write Cache Write Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in write-through mode.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Read Ahead Percentage of Cache Hits
The percentage of all read-cache hits that occur on pre-staged data.

Dirty Write Percentage of Cache Hits
The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

Write Cache Overflow Percentage
The percentage of write operations that are delayed because there is not enough space in the write cache.

Write Cache Flush Through Percentage
The percentage of tracks that is written to disk in flush-through mode.

Write Cache Write Through Percentage
The percentage of tracks that is written to disk in write-through mode.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Back-end data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Peak Back-End Read Response Time (ms)
The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)
The longest time for a back-end storage resource to respond to a write operation.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Read Queue Time (ms)
The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)
The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write Percentage
The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Global Mirror Overlapping Write I/O Rate (op/s)
The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Global Mirror Secondary Write Lag (ms/op)
The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

Peak Read Response Time (ms)
The longest time to complete a read operation.

Peak Write Response Time (ms)
The longest time to complete a write operation.

Link Failure Rate (count/s)
The average number of link errors per second.

Loss of Sync Rate (count/s)
The average number of times per second that synchronization is lost.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)
The average number of times per second that the signal is lost.

CRC Error Rate (count/s)
The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

Primitive Sequence Protocol Error Rate (count/s)
The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)
The average number of bit errors per second that are detected.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Overall Host Attributed Response Time Percentage
The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

The value is based on the following conditions:

- For read response time, the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes.
- For write response time, the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Zero Buffer-to-Buffer Credit Timer (microseconds)
The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

Data for storage system I/O groups in performance reports

You can include general information, capacity data, properties, and other information about storage system I/O groups in performance reports.

Information about storage system I/O groups

You can create performance reports that include the following information:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage I/O Group Name
The name that was assigned to the I/O group when it was added to the system.

Component properties

You can create performance reports that include the following information:

Storage I/O Group Mirroring Total Memory (MiB)
The total amount of memory that is allocated on the nodes in an I/O group to provide volume mirroring functions.

Storage I/O Group Mirroring Free Memory (MiB)
The amount of available memory on the nodes in an I/O group to provide volume mirroring functions.

Storage I/O Group FlashCopy Total Memory (MiB)
The total amount of memory that is allocated on the nodes in an I/O group to provide FlashCopy mirroring functions.

Storage I/O Group FlashCopy Free Memory (MiB)
The amount of available memory on the nodes in an I/O group to provide FlashCopy mirroring functions.

Storage I/O Group Remote Copy Total Memory (MiB)
The total amount of memory that is allocated on the nodes in an I/O group to provide mirroring functions by using a remote copy.

Storage I/O Group Remote Copy Free Memory (MiB)
The amount of available memory on the nodes in an I/O group to provide mirroring functions by using a remote copy.

Storage I/O Group RAID Total Memory (MiB)
The total amount of memory that is allocated on the nodes in an I/O group to provide RAID mirroring functions.

Storage I/O Group RAID Free Memory (MiB)
The amount of available memory on the nodes in an I/O group to provide RAID mirroring functions.

Status information

You can create performance reports that include the following information:

Storage I/O Group Maintenance
Shows whether the I/O group is in maintenance mode. If this value is `Yes`, the I/O group is in maintenance mode.

Storage I/O Group Is Compression Active
Shows whether the compression feature is enabled on the volumes in the storage I/O group. If this value is `Yes`, the compression feature is enabled.

Storage I/O Group Is Compression Supported
Shows whether the compression feature is available for volumes in the storage I/O group. If this value is `Yes`, the compression feature is available.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Performance metrics for I/O groups on SAN Volume Controller and Storwize systems:

A storage environment can include SAN Volume Controller systems, Storwize V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports, volumes, caches, back-end arrays, and other data for I/O groups on these systems in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port to Host Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to host computers.

Port to Host Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from host computers.

Total Port to Host I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between host computers and the component.

Port to Disk Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to back-end storage resources.

Port to Disk Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from back-end storage resources.

Total Port to Disk I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between back-end storage resources and the component. This value includes both send and receive operations.

Port to Local Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from other nodes that are in the local SAN Volume Controller cluster.

Total Port to Local Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send I/O Rate (ops/s)
The average number of exchanges per second that are sent to nodes that are in the remote SAN Volume Controller cluster.

Port to Remote Node Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from nodes that are in the remote SAN Volume Controller cluster.

Total Port to Remote Node I/O Rate (ops/s)
The average number of exchanges per second that are transmitted between the component and nodes that are in the remote SAN Volume Controller cluster.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port to Host Send Data Rate (MiB/s)
The average rate at which data is sent to host computers. The rate is measured in MiB per second.

Port to Host Receive Data Rate (MiB/s)
The average rate at which data is received from host computers. The rate is measured in MiB per second.

Total Port to Host Data Rate (MiB/s)
The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Disk Send Data Rate (MiB/s)
The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

Port to Disk Receive Data Rate (MiB/s)
The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

Total Port to Disk Data Rate (MiB/s)
The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Local Node Send Data Rate (MiB/s)
The average rate at which data is sent to other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Receive Data Rate (MiB/s)
The average rate at which data is received from other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Local Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and other nodes that are in the local SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Send Data Rate (MiB/s)
The average rate at which data is sent to nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Remote Node Receive Data Rate (MiB/s)
The average rate at which data is received from nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Total Port to Remote Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and nodes that are in the remote SAN Volume Controller cluster. The rate is measured in MiB per second.

Port to Local Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Local Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from

another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Overall Port to Local Node Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation with another node that is in the local SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Remote Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Remote Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from a node that is in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Overall Port to Remote Node Response Time (ms/op)
The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote SAN Volume Controller cluster. This value represents the external response time of the transfers.

Port to Local Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local SAN Volume Controller cluster.

Port to Local Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local SAN Volume Controller cluster.

Overall Port to Local Node Queue Time (ms/op)
The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local SAN Volume Controller cluster.

Port to Remote Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote SAN Volume Controller cluster.

Port to Remote Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote SAN Volume Controller cluster.

Overall Port to Remote Node Queue Time (ms/op)
The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote SAN Volume Controller cluster.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted

between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Peak Back-End Read Response Time (ms)
The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)
The longest time for a back-end storage resource to respond to a write operation.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Read Queue Time (ms)
The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Write Cache Flush Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in flush-through mode.

Write Cache Write Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in write-through mode.

Write Cache Overflow I/O Rate (ops/s)
The average number of write operations per second that are delayed because there is not enough space in the write cache.

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Read Ahead Percentage of Cache Hits
The percentage of all read-cache hits that occur on pre-staged data.

Dirty Write Percentage of Cache Hits
The percentage of all cache write operations that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

Write Cache Flush Through Percentage
The percentage of tracks that is written to disk in flush-through mode.

Write Cache Write Through Percentage
The percentage of tracks that is written to disk in write-through mode.

Write Cache Overflow Percentage
The percentage of write operations that are delayed because there is not enough space in the write cache.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)
The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write Percentage
The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Global Mirror Overlapping Write I/O Rate (op/s)
The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Peak Read Response Time (ms)
The longest time to complete a read operation.

Peak Write Response Time (ms)
The longest time to complete a write operation.

Global Mirror Secondary Write Lag (ms/op)
The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

Link Failure Rate (count/s)
The average number of link errors per second.

Loss of Sync Rate (count/s)
The average number of times per second that synchronization is lost.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)
The average number of times per second that the signal is lost.

CRC Error Rate (count/s)
The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

Processor Utilization Percentage
The average percentage of time that the processors are busy.

Overall Host Attributed Response Time Percentage
The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

The value is based on the following conditions:

- For read response time, the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes.
- For write response time, the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Nonpreferred Node Usage Percentage
The overall percentage of I/O operations that are completed, or data that is transferred over the non-preferred nodes of the virtual disks.

Primitive Sequence Protocol Error Rate (count/s)
The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)
The average number of bit errors per second that are detected.

Zero Buffer-to-Buffer Credit Timer (microseconds)
The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset.

Data for local disks on storage systems in performance reports

You can include general information, capacity data, properties, and other information about local disks that are on storage systems in performance reports.

Information about local disks on storage systems

You can create performance reports that include the following information:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage MDisk Name
The name that was assigned to the managed disk when it was added to the system.

Storage Local Disk Name
The name that was assigned to the local disk on the storage system when the local disk was created on the system.

Capacity and usage data

You can create performance reports that include the following information:

Storage Local Disk Capacity (GiB)
The amount of storage space on the local disk.

Component properties

You can create performance reports that include the following information:

Storage Local Disk Is Encryptable
Shows whether the resource can be encrypted. If this value is `Yes`, the resource can be encrypted.

Storage Local Disk Is Encrypted
Shows whether the resource is encrypted. If this value is `Yes`, the resource is encrypted.

Storage Local Disk Storage Class
The storage technology of the local disk on a storage system. For example, the storage class can be a serial-attached SCSI (SAS) or solid-state drive (SSD).

Storage Local Disk Is Solid State
Shows whether the disk is a solid-state drive.

Storage Local Disk Speed (RPM)
The speed of the local disk on a storage system.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage Local Disk Status
The condition of the resource, for example normal, warning, or error.

Vendor, model, and device information

You can create performance reports that include the following information:

Storage System Local Disk Firmware Version
The version number of the firmware that is running on the disk.

Storage System Local Disk Model
The model name or model number of the resource.

Storage System Local Disk Serial Number
The serial number of the resource.

Storage System Local Disk Vendor
The vendor who supplied the resource.

Performance metrics for local disks on storage systems:

You can include performance data for back-end arrays on local disks that are on storage systems in performance reports.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

For SAN Volume Controller models, this value represents the external response time of the managed disks.

Peak Back-End Read Response Time (ms)
The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)
The longest time for a back-end storage resource to respond to a write operation.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Read Queue Time (ms)
The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Data for storage host connections in performance reports

You can include general information, capacity data, properties, and other information about storage host connections in performance reports.

Information about storage host connections

You can create performance reports that include the following information:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Storage Host Connection Name
The user-defined name that describes the server or storage device that is assigned to a volume.

XIV Cluster
The name of the cluster, if any, as defined on the XIV system. If no cluster is defined for the host connection, this field is blank.

Performance metrics for host connections on SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified systems:

You can include performance metrics for volumes, caches, and for other data for host connections in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Write Cache Flush Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in flush-through mode.

Write Cache Write Through I/O Rate (ops/s)
The average number of tracks per second that are written to disk in write-through mode.

Write Cache Overflow I/O Rate (ops/s)
The average number of write operations per second that are delayed because there is not enough space in the write cache.

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Read Ahead Percentage of Cache Hits
The percentage of all read-cache hits that occur on pre-staged data.

Dirty Write Percentage of Cache Hits
The percentage of all cache write operations that occur on data in the cache

that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk.

Write Cache Flush Through Percentage
The percentage of tracks that is written to disk in flush-through mode.

Write Cache Write Through Percentage
The percentage of tracks that is written to disk in write-through mode.

Write Cache Overflow Percentage
The percentage of write operations that are delayed because there is not enough space in the write cache.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)
The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write I/O Rate (op/s)
The average number of overlapping write operations per second that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Peak Read Response Time (ms)
The longest time to complete a read operation.

Peak Write Response Time (ms)
The longest time to complete a write operation.

Global Mirror Secondary Write Lag (ms/op)
The average number of milliseconds to service a secondary write operation for Global Mirror. This value does not include the time to service the primary write operation.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

Overall Host Attributed Response Time Percentage
The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

The value is based on the following conditions:

- For read response time, the time that it takes for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes.
- For write response time, the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Global Mirror Overlapping Write Percentage
The percentage of overlapping write operations that are issued by the Global Mirror primary site.

For SAN Volume Controller software releases 4.3.1 and later, some overlapping writes are processed in parallel and are excluded from this value. For SAN Volume Controller software releases 4.3.0 and earlier, all overlapping writes are serialized.

Nonpreferred Node Usage Percentage
The overall percentage of I/O operations that are completed, or data that is transferred over the non-preferred nodes of the virtual disks.

Performance metrics for host connections on DS8000 and DS6000 storage systems:

You can include performance metrics for volumes, caches, and for other data for storage host connections on DS8000 and DS6000 storage systems in performance reports.

Volume data

You can create performance reports that include the following information:

Normal Read I/O Rate (ops/s)
The average number of nonsequential read operations per second.

Sequential Read I/O Rate (ops/s)
The average number of sequential read operations per second.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Normal Write I/O Rate (ops/s)
The average number of nonsequential write operations per second.

Sequential Write I/O Rate (ops/s)
The average number of sequential write operations per second.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Normal I/O Rate (ops/s)
The average number of nonsequential I/O operations per second. This value includes both read and write operations.

Total Sequential I/O Rate (ops/s)
The average number of sequential I/O operations per second. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Record Mode Read I/O Rate (ops/s)
The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Normal Read Cache Hit Percentage
The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Sequential Read Cache Hit Percentage
The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Normal Write Cache Hit Percentage
The percentage of nonsequential write operations that are handled in the cache.

Sequential Write Cache Hit Percentage
The percentage of sequential write operations that are handled in the cache.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Normal Cache Hit Percentage
The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Sequential Cache Hit Percentage
The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Record Mode Read Cache Hit Percentage
The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

Disk to Cache Transfer Rate (ops/s)
The average number of tracks that are transferred per second from the disks to the cache.

Cache to Disk Transfer Rate (ops/s)
The average number of tracks per second that are transferred from the cache to the disks.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because there is not enough space in the write cache.

The only operations that can be delayed are write operations, but the value is calculated as an average of all I/O operations.

Other data

You can create performance reports that include the following information:

HPF Read I/O Rate (ops/s)
The average number of read operations per second that are issued by the High Performance FICON feature of the storage system.

HPF Write I/O Rate (ops/s)
The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

Total HPF I/O Rate (ops/s)
The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

HPF I/O Percentage
The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

PPRC Transfer Rate (ops/s)
The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

Performance metrics for storage host connections on an XIV system:

You can include performance metrics for volumes, caches, and for other data for storage host connections on an XIV system in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiB per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and nonsequential write operations.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes sequential and nonsequential operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Read Data Cache Hit Percentage
The percentage of all read data that is read from the cache.

Write Data Cache Hit Percentage
The percentage of all write data that is written to cache slots that are marked as modified.

Overall Data Cache Hit Percentage
The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

Read Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a read-cache hit operation.

Write Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a write-cache hit operation.

Overall Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

Read Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a read-cache miss operation.

Write Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a write-cache miss operation.

Overall Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

Other data

You can create performance reports that include the following information:

Small Transfers I/O Percentage
The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers I/O Percentage
The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers I/O Percentage
The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers I/O Percentage
The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Data Percentage
The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has

a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Response Time (ms/op)
The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

Data for switches in performance reports

You can include properties and other information about switches in performance reports.

Information about switches

You can create performance reports that include the following information:

Switch Logical Name
The logical name of a switch or switch port.

Component properties

You can create performance reports that include the following information:

Switch WWN
The worldwide name of the switch.

Parent Switch WWN
The worldwide name of the parent switch.

Switch Management Telnet Address
The Telnet address that you access to manage the switch.

Switch Management SNMP Address
The Simple Network Management Protocol (SNMP) address that you access to manage the switch.

Switch Management URL Address
The URL that you access to manage the switch.

Switch Domain
The domain of the switch.

Switch Version
The version number of the operating system that is running on the switch.

Switch Location
The physical location of a switch. The location is defined when a switch is added to Tivoli Storage Productivity Center. You can add or edit the location of the switch in the **Properties** pane of the switch in the web-based GUI.

Switch Custom Tag 1, 2, and 3
User-defined text that is associated with a switch. You can add or edit the custom tags for a switch in the **Properties** pane of the switch in the Tivoli Storage Productivity Center web-based GUI.

Status information

You can create performance reports that include the following information:

Switch Status
The condition of the resource, for example normal, warning, or error.

Switch Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Vendor, model, and device information

You can create performance reports that include the following information:

Switch IP Address
The IP address of the resource.

Switch Serial Number
The serial number of the resource.

Switch Vendor
The vendor who supplied the resource.

Switch Model
The model name or model number of the resource.

Performance metrics for switches:

You can include performance data for switches, including switch error data, in performance reports.

Switch data

You can create performance reports that include the following information:

Port Send Frame Rate (frames/s)
The average number of frames per second that are sent through the port.

Port Receive Frame Rate (frames/s)
The average number of frames per second that are received by the port.

Total Port Frame Rate (frames/s)
The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Send Frame Size (KiB/op)
The average size of a frame, in KiB, that is sent through the port.

Port Receive Frame Size (KiB/op)
The average size of a frame, in KiB, that is received by the port.

Overall Port Frame Size (KiB/op)
The average frame transfer size. This value is measured in KiB and includes frames that are sent and frames that are received by the port.

Port Send Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage
The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Switch error data

You can create performance reports that include the following information:

Error Frame Rate (count/s)
The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

Discarded Frame Rate (count/s)
The average number of frames per second that are discarded because host buffers are unavailable for the port.

Link Failure Rate (count/s)
The average number of link errors per second.

Loss of Sync Rate (count/s)
The average number of times per second that synchronization is lost.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)
The average number of times per second that the signal is lost.

CRC Error Rate (count/s)
The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

Encoding Disparity Error Rate (count/s)
The average number of disparity errors per second that are received.

Discarded Class 3 Frame Rate (count/s)
The average number of class 3 frames per second that are discarded.

Primitive Sequence Protocol Error Rate (count/s)
The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)
The average number of bit errors per second that are detected.

Link Reset Transmitted Rate (count/s)
The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.

Link Reset Received Rate (count/s)
The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.

Short Frame Rate (count/s)
The average number of frames that are received per second that are shorter than 28 octets. This number excludes start-of-frame bytes and end-of-frame bytes.

The 28 octet limit is calculated based on the assumption that a frame has 24 header bytes, and 4 CRC bytes.

Long Frame Rate (count/s)
The average number of frames that are received per second that are longer than 2140 octets. This number excludes start-of-frame bytes and end-of-frame bytes.

The 2140 octet limit is calculated based on the assumption that a frame has 24 header bytes, 4 CRC bytes, and 2112 data bytes.

F-BSY Frame Rate (count/s)
The average number of F-BSY frames per second that are generated.

An F-BSY frame is issued by the fabric to indicate that a frame cannot be delivered because the fabric or destination N_port is busy.

F-RJT Frame Rate (count/s)
The average number of F-RJT frames per second that are generated.

An F-RJT frame is issued by the fabric to indicate that delivery of a frame was denied.

Data for switch ports in performance reports

You can include properties and other information about switch ports in performance reports.

Information about switch ports

You can create performance reports that include the following information:

Switch Logical Name
The logical name of a switch or switch port.

Switch Port WWPN
The worldwide port name of the port on the switch.

Component properties

You can create performance reports that include the following information:

Switch Port Number
The port number on the resource.

Switch Port Type
The type of port on the storage system, storage virtualizer, or switch. For example, the port type can be N_Port, F_Port, or another type of port.

Switch Port Speed (GiB/s)
The speed of a port, which is measured in GiB per second.

Switch Blade Slot Number
The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

Status information

You can create performance reports that include the following information:

Switch Port Status
The condition of the resource, for example normal, warning, or error.

Switch Port Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Switch Port Enabled State
Shows whether a port is enabled, disabled, or is enabled but offline.

Switch Port Operational Status
The current operational state of the switch port. For example, this value can be `Error` or `OK`.

Performance metrics for switch ports:

You can include performance data for switch ports, including port error data, in performance reports.

Switch port data

You can create performance reports that include the following information:

Port Send Frame Rate (frames/s)
The average number of frames per second that are sent through the port.

Port Receive Frame Rate (frames/s)
The average number of frames per second that are received by the port.

Total Port Frame Rate (frames/s)
The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.

Port Send Data Rate (MiB/s)
The average rate at which data is sent through the port. The rate is measured in MiB per second.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Peak Send Data Rate (MiB/s)
The fastest rate at which data is sent through this switch port.

Port Peak Receive Data Rate (MiB/s)
The fastest rate at which data is received through this switch port.

Port Send Frame Size (KiB/op)
The average size of a frame, in KiB, that is sent through the port.

Port Receive Frame Size (KiB/op)
The average size of a frame, in KiB, that is received by the port.

Overall Port Frame Size (KiB/op)
The average frame transfer size. This value is measured in KiB and includes frames that are sent and frames that are received by the port.

Port Send Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage
The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Data for errors on switch ports

You can create performance reports that include the following information:

Error Frame Rate (count/s)
The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

Discarded Frame Rate (count/s)
The average number of frames per second that are discarded because host buffers are unavailable for the port.

Link Failure Rate (count/s)
The average number of link errors per second.

Loss of Sync Rate (count/s)
The average number of times per second that synchronization is lost.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)
The average number of times per second that the signal is lost.

CRC Error Rate (count/s)
The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.

Encoding Disparity Error Rate (count/s)
The average number of disparity errors per second that are received.

Discarded Class 3 Frame Rate (count/s)
The average number of class 3 frames per second that are discarded.

Primitive Sequence Protocol Error Rate (count/s)
The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)
The average number of bit errors per second that are detected.

Link Reset Transmitted Rate (count/s)
The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.

Link Reset Received Rate (count/s)
The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.

Short Frame Rate (count/s)
The average number of frames that are received per second that are shorter than 28 octets. This number excludes start-of-frame bytes and end-of-frame bytes.

The 28 octet limit is calculated based on the assumption that a frame has 24 header bytes, and 4 CRC bytes.

Long Frame Rate (count/s)
The average number of frames that are received per second that are longer than 2140 octets. This number excludes start-of-frame bytes and end-of-frame bytes.

The 2140 octet limit is calculated based on the assumption that a frame has 24 header bytes, 4 CRC bytes, and 2112 data bytes.

F-BSY Frame Rate (count/s)
The average number of F-BSY frames per second that are generated.

An F-BSY frame is issued by the fabric to indicate that a frame cannot be delivered because the fabric or destination N_port is busy.

F-RJT Frame Rate (count/s)
The average number of F-RJT frames per second that are generated.

An F-RJT frame is issued by the fabric to indicate that delivery of a frame was denied.

Custom reports about capacity and relationships

You can use Tivoli Common Reporting to create custom reports about capacity and relationships. Capacity and relationship reports can contain detailed information about the properties of monitored resources, and the available space and capacity of those resources.

Capacity and relationship reports can contain information about the following resources and their components:

- Servers
- Hypervisors
- Network-attached storage servers
- Storage virtualizers
- Storage systems
- Storage resource groups

Creating custom capacity and relationship reports

You can create custom reports that show capacity and relationship information for resources. Capacity and relationship reports can contain detailed information about the properties of monitored resources, and the available space and capacity of those resources. You access custom capacity and relationship reports from the Tivoli Storage Productivity Centerweb-based GUI, and you create the reports in Tivoli Common Reporting.

Before you can create a custom report about capacity and relationships, ensure that data was collected from the resources. The following table shows where the data must be collected from for each resource:

Resource	Prerequisite
Server	Data must be collected from the server by the Storage Resource agent, or by an agentless server.
Hypervisor	Data must be collected from the VMware ESX server.
NAS	Data must be collected from any NetApp or IBM Scale Out Network Attached Storage resources that are in your storage environment.
Storage virtualizer	Data must be collected from the any of the following storage virtualizers that are in your storage environment: • Storwize V7000 • Storwize V7000 Unified • SAN Volume Controller • Hitachi Universal Storage Platform V
Storage system	Data must be collected from the storage system, or any of the following systems that are configured as back-end storage: • Storwize V7000 • Storwize V7000 Unified • SAN Volume Controller • Hitachi Universal Storage Platform V

Tip: Before you run a report, you can preview the report with no data.

1. Start the web-based GUI for Tivoli Storage Productivity Center.
2. Depending on the type of report that you want to see, complete one of the following steps:
   - To create a basic report with limited formatting about the capacity and relationships of resources, select **Reporting** > **Create capacity reports**.
   - To create a report about the performance of resources and apply advanced formatting features to the report, complete the following steps:
     a. Select **Reporting** > **Create capacity reports**.
     b. Click the **Return** icon .
     c. Click **Launch** > **Report Studio**.
3. Expand **Tivoli Storage Productivity Center** in the list.
4. Expand **Servers to Storage Systems** in the list.

5. Explore the folders in the list to locate the resources, properties, and statistics that you want to include in the report.
6. Depending on your browser, do one of the following steps to add properties and statistics to the report:
   - In Internet Explorer, drag the items that you want to see information about to the work area.
   - In Firefox, select the items that you want to see information about, and then click **Insert**.
7. Click **Save**.

Related tasks:

"Adding filters to reports"
You can use filters with capacity and relationship reports. Filters allow you to show only certain data that complies with your criteria. You add filters to your capacity and relationship reports when you create the reports.

Related information:

IBM Cognos Business Intelligence 10.2.0 Information Center
For more information about how to preview and run reports, see the IBM Cognos Business Intelligence 10.2.0 Information Center at pic.dhe.ibm.com/infocenter/cbi/v10r2m0/index.jsp. Search for "Run a Report".

Adding filters to reports

You can use filters with capacity and relationship reports. Filters allow you to show only certain data that complies with your criteria. You add filters to your capacity and relationship reports when you create the reports.

Filters for resources are in the same list as the resources, in the appropriate resource folder. For example, to view the filters for servers, select **Server Resources** > **Server Resource Filters** > **Server Filters**.

Filters are available for all resources, except network-attached storage resources.

1. To apply a filter to the data in the work area, locate the filter that you want to include in the folders in the list.
2. Depending on your browser, do one of the following steps to add a filter to the report:
   - In Internet Explorer, drag the filter that you want to use to the work area.
   - In Firefox, select the filter that you want to use, and then click **Insert**.

   The filter is applied to the report. If you did not add properties and statistics to the report, the filter is applied when you add the properties and statistics.
3. Optional: To delete the filter, right-click the filter in the work area, and then select **Delete**.

Related information:

IBM Cognos Business Intelligence 10.2.0 Information Center
For more information about how to use filters with reports, see the IBM Cognos Business Intelligence 10.2.0 Information Center at pic.dhe.ibm.com/infocenter/cbi/v10r2m0/index.jsp. Search for "Create a Simple Filter".

Filters for servers in capacity and relationship reports

You can filter capacity and relationship reports to include only the servers, server file systems, and server logical volumes that meet specific criteria.

Filters for servers

You can create capacity and relationship reports that include the following filters:

Servers Connected to Virtualizers or Storage Systems
Shows servers that are connected to storage virtualizers or storage systems. You can use this filter to identify servers that are connected directly to storage virtualizers, or connected directly or indirectly to storage systems. This filter excludes local disks from the report.

Servers Directly Connected to Storage Systems
Shows servers that are directly connected to a back-end storage system. You can use this filter to identify servers or server disks that are connected to a storage system, but are not connected through a storage virtualizer.

For example, you might want to migrate data from one storage system to another. In this case, you can use this filter to identify file systems or server disks that are not virtualized. These file systems or server disks might be affected if a storage system is removed.

Servers Directly Connected to Virtualizers
Shows servers that are directly connected to a storage virtualizer. You can use this filter to identify servers that are connected to a storage virtualizer, but are not connected directly to a back-end storage system.

For example, you can use this filter to identify only server disks that are on a SAN Volume Controller.

Filters for server file systems

You can create capacity and relationship reports that include the following filters:

Rows with File System Mount Points
Shows the file systems that have a mount point on the server.

Rows without File System Mount Points
Shows resources that do not have a mount point on the server. You can use this filter to identify disks that do not have file systems, that is, disks that are not in use.

For example, to identify unused space include the Server Name property and the Server Disk Path property in a report, and then include the Rows without File System Mount Points filter.

Filters for server logical volumes

You can create capacity and relationship reports that include the following filters:

Rows with Logical Volume Path
Shows the logical volumes that have a path to the server.

Rows without Logical Volume Path
Shows resources that do not have a logical volume path to the server. You can use this filter to identify volumes that do not have paths, that is, volumes that are not in use.

For example, to identify unused space include the Server Name property and the Server Disk Path property in a report, and then include the Rows without Logical Volume Path filter.

Filters for hypervisors in capacity and relationship reports

You can filter capacity and relationship reports to include only the hypervisors, hypervisor file systems, and hypervisor logical volumes that meet specific criteria.

Filters for hypervisors

You can create capacity and relationship reports that include the following filters:

Hypervisor with Probed Servers
Shows hypervisors that are connected to servers. You can use this filter to include only hypervisors and their associated servers from which data was collected. For example, you might collect data only from your critical production servers on a hypervisor, but not your test servers. In this case, you can use this filter to exclude information for the test servers from the report.

Hypervisors Connected to Virtualizers or Storage Systems
Shows hypervisors that are connected to storage virtualizers or storage systems. You can use this filter to identify hypervisors that are connected directly to storage virtualizers, or connected directly or indirectly to storage systems. This filter excludes local disks from the report.

Hypervisors Directly Connected to Storage Systems
Shows hypervisors that are directly connected to a back-end storage system. You can use this filter to identify hypervisors or hypervisor disks that are connected to a storage system, but are not connected through a storage virtualizer.

For example, you might want to migrate data from one storage system to another. In this case, you can use this filter to identify file systems or hypervisor disks that are not virtualized. These file systems or hypervisor disks might be affected if a storage system is removed.

Hypervisors Directly Connected to Virtualizers
Shows hypervisors that are directly connected to a storage virtualizer. You can use this filter to identify hypervisors that are connected to a storage virtualizer, but are not connected directly to a back-end storage system.

For example, you can use this filter to identify only hypervisor disks that are on a SAN Volume Controller.

Filters for hypervisor file systems

You can create capacity and relationship reports that include the following filters:

Hypervisors with File System Mount Points
Shows the file systems that have a mount point on the hypervisor.

Hypervisors without File System Mount Points
Shows resources that do not have a mount point on the hypervisor. You can use this filter to identify disks that do not have file systems, that is, disks that are not in use.

For example, to identify unused space include the Hypervisor Name property and the Hypervisor Disk Path property in a report, and then include the Hypervisors without File System Mount Points filter.

Filters for hypervisor logical volumes

You can create capacity and relationship reports that include the following filters:

Hypervisors with Logical Volume Path
Shows the logical volumes that have a path to the hypervisor.

Hypervisors without Logical Volume Path
Shows resources that do not have a logical volume path to the hypervisor. You can use this filter to identify volumes that do not have paths, that is, volumes that are not in use.

For example, to identify unused space include the Hypervisor Name property and the Hypervisor Disk Path property in a report, and then include the Hypervisors without Logical Volume Path filter.

Filters for storage virtualizers in capacity and relationship reports

You can filter capacity and relationship reports to include only the storage virtualizers and storage virtualizer volumes that meet specific criteria.

Storage virtualizer filters

You can create capacity and relationship reports that include the following filters:

Virtualizers Not Connected to Servers
Shows storage virtualizers that are not connected to servers.

Filters for storage virtualizer volumes

You can create capacity and relationship reports that include the following filters:

Rows with Virtualizer Volumes
Shows storage virtualizer disks that are connected to the storage virtualizer that have an associated volume.

Rows without Virtualizer Volumes
Shows storage virtualizer disks that are connected to the storage virtualizer that do not have an associated volume.

Virtualizer Primary Volume Information
Shows primary volumes that are connected to the storage virtualizer. You can use this filter if you want to get a total of the used space on primary volumes in your storage environment.

Virtualizer Volume Copy Information
Shows copy volumes that are connected to the storage virtualizer. You can use this filter if you want to get a total of the used space on volume copies in your storage environment.

Filters for storage systems in capacity and relationship reports

You can filter capacity and relationship reports to include only the storage systems or storage volumes that meet specific criteria.

Filters for storage systems

You can create capacity and relationship reports that include the following filters:

Storage Systems Directly Connected to Servers
Shows back-end storage systems that are directly connected to a server. You can use this filter to identify storage systems that are connected to a server, but are not connected through a storage virtualizer.

For example, you might want to migrate data from one storage system to another. In this case, you can use this filter to identify file systems or

server disks that are not virtualized. These file systems or server disks might be affected if a storage system is removed.

Storage Systems Connected to Servers through Virtualizers
Shows storage systems that are indirectly connected to a server through a storage virtualizer. You can use this filter to identify storage systems that are connected to a storage virtualizer, but are not connected directly to a server.

For example, you can use this filter to identify only server disks that are on a SAN Volume Controller.

Storage Systems not Connected to Servers
Shows storage systems that are not connected to a server either directly, or indirectly through a storage virtualizer.

Rows with Storage Volumes
Shows storage resources that have an associated storage volume.

Rows without Storage Volumes
Shows storage resources that do not have an associated storage volume.

Filters for storage resource groups in capacity and relationship reports

You can filter capacity and relationship reports to include only the storage resource groups that meet specific criteria.

Filters for storage resource groups

You can create capacity and relationship reports that include the following filters:

Storage Resource Group for Volumes
Shows only the volumes that are in a storage resource group.

Storage Resource Group for Pools
Shows only the pools that are in a storage resource group.

Data and properties in capacity and relationship reports

You can include capacity data, properties, relationship information, and other information about resources in your storage environment in capacity and relationship reports.

Data for clusters in capacity and relationship reports

You can include general information about clusters in capacity and relationship reports.

Information about clusters

You can create capacity and relationship reports that include the following information:

Cluster Name
The name that was assigned to the cluster when it was created.

Cluster Type
The type of the cluster. The cluster can be a Microsoft Cluster Server or a High-Availability Cluster Multiprocessing (HACMP) server.

Cluster IP Address
The IP address of the resource.

Cluster Domain Name
The domain name of the cluster.

Data for servers in capacity and relationship reports

You can include general information, capacity data, properties of components, and other information about servers in capacity and relationship reports.

Information about servers

You can create capacity and relationship reports that include the following information:

Server Name
The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Server Domain Name
The domain name of the server.

Server Short Name
The host name from the fully qualified domain name of the server. For example, if the fully qualified domain name of the server is `mycomputer.example.com`, the host name is `mycomputer`.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Server Swap Space (GiB)
The amount of disk space on a resource that is available to store some of the contents of the RAM when the RAM is full.

Server Total Disk Capacity (GiB)
The amount of storage space that is on a server.

Server Total Disk Available Space (GiB)
The amount of unused storage space on all disks on a server.

Server Total File System Available Space (GiB)
The amount of unused storage space that is available on the file systems on a server.

Component properties

You can create capacity and relationship reports that include the following information:

Server Cluster Current Hosting Node
The name of the server that is hosting the cluster resource group.

Server Cluster Resource Type
Shows whether the cluster is a node or a cluster resource group.

Server IP Address
The IP address of the resource.

Server Is Virtual Machine
Shows whether a resource is a virtual machine. If this value is `Yes`, the resource is a virtual machine.

Server OS Type
The operating system that is running on the server.

Server OS Version
The version number of the operating system that is running on the resource.

Server Processor Architecture
The architecture of the processor on the server or hypervisor. For example, the architecture of a processor might be Intel 64 bit (IA64) or Intel 32 bit (IA32).

Server Processor Count
The number of processors on the server or hypervisor.

Server Processor Speed (MHz)
The speed of the processor on the server or hypervisor.

Server Processor Type
Shows information about the processor, such as the family and model of the processor.

Server RAM (GiB)
The amount of RAM on the server or hypervisor.

Server Time Zone
The time zone in which a resource is located.

Server Location
The physical location of a server. The location is defined when a server is added to Tivoli Storage Productivity Center. You can add or edit the location of the server in the **Properties** pane of the server in the web-based GUI.

Server Custom Tag 1, 2, and 3
User-defined text that is associated with a server. You can add or edit the custom tags for a server in the **Properties** pane of the server in Tivoli Storage Productivity Center web-based GUI.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Server Model
The model name or model number of the resource.

Server Serial Number
The serial number of the resource.

Server Vendor
The vendor who supplied the resource.

Status information

You can create capacity and relationship reports that include the following information:

Server Status
The status of a server, such as `Normal`, `Warning`, `Error`, `Unreachable`,

Unknown, or Agentless. Use the status to determine the condition of the server, and if any actions must be taken. If this value is Agentless, the server is an unmanaged server.

Server Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is Yes, the resource was detected.

Server Last Data Collection
The date and time when storage statistics were last collected from the resource.

Server Last Start Time
The last time that the resource was started.

Data for file systems on servers in capacity and relationship reports

You can include general information, capacity data, and status information about file systems on servers in capacity and relationship reports.

Information about file systems on servers

You can create capacity and relationship reports that include the following information:

File System Mount Point
The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be c:\ or d:\. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be /opt or /export/home.

File System Export Name
The name of the exported file system.

File System Type
The type of file system that the resource uses.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

File System Capacity (GiB)
The amount of storage space on the file system of the resource.

File System Directory Count
The number of directories in the file systems that are on a server.

File System File Count
The number of files in the file system.

File System Available Space (GiB)
The amount of unused storage space in the file system of the resource.

File System Available Inodes
The number of unused inodes in file systems on the operating system.

File System Used Inodes
The number of used inodes in file systems on the operating system.

Tip: For Microsoft Windows systems, this property is blank.

File System Maximum Files
The maximum number of files that the file system on the resource can contain.

File System Physical Size (GiB)
The amount of physical storage space on the file systems on a resource. The physical size is the size of all the clusters that the file system uses.

File System Used Space (GiB)
The amount of used storage space in the file system of the resource.

Component properties

You can create capacity and relationship reports that include the following information:

File System Is Remote Mount
Shows whether the file system is remotely mounted on another server. If this value is `Yes`, the file system is remotely mounted.

Status information

You can create capacity and relationship reports that include the following information:

File System Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

File System Last Scan
The date and time when file statistics were last collected from the resource.

Data for logical volumes on servers in capacity and relationship reports

You can include general information, capacity data, and status information about logical volumes on servers in capacity and relationship reports.

Information about logical volumes on servers

You can create capacity and relationship reports that include the following information:

Logical Volume Path
The path to a logical volume on a resource, for example `/dev/hd1`.

Logical Volume Is Swap Space
Shows whether a logical volume is used to store some of the contents of the RAM when the RAM is full. If this value is `Yes`, the logical volume is used to store some of the contents of the RAM.

Logical Volume Mirror Count
The number of mirrors that keep a synchronous copy of the resource.

Logical Volume Type
The type of server logical volume. For example, the logical volume type can be concatenation or stripe set.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Logical Volume Allocation on Server Disk (GiB)
The amount of disk storage space that is allocated to the logical volume.

Logical Volume Capacity (GiB)
The amount of storage space that is on a logical volume.

Logical Volume Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Status information

You can create capacity and relationship reports that include the following information:

Logical Volume Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Data for volume groups on servers in capacity and relationship reports

You can include general information and capacity data about volume groups on servers in capacity and relationship reports.

Information about volume groups on servers

You can create capacity and relationship reports that include the following information:

Volume Group Name
The name that was assigned to the volume group when it was added to the system.

Volume Group Type
The type of the volume manager that manages the volume group. The volume manager can be a logical volume manager (LVM) or a Veritas volume manager.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Volume Group Capacity (GiB)
The amount of storage space that is on a volume group.

Volume Group Available Space (GiB)
The amount of unused storage space on a server volume group.

Volume Group Number of Logical Volumes
The number of logical volumes on a resource.

Volume Group Number of Server Disks
The number of server disks in a volume group.

Data for server disks in capacity and relationship reports

You can include general information, capacity data, properties of components, and other information about server disks in capacity and relationship reports.

Information about server disks

You can create capacity and relationship reports that include the following information:

Server Disk Path
The identifier that specifies a path to a server disk, for example `/dev/hdisk0`.

A server disk might have one or more disk paths. To identify unique server disks, add the Server Disk Serial Number property to the report, and group the server disks by the serial number.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Server Disk Capacity (GiB)
The amount of storage space that is on a server disk.

Server Disk Available Space (GiB)
The amount of unused storage space that is on a server disk.

Component properties

You can create capacity and relationship reports that include the following information:

Server Disk Is Removable
Shows whether the disk can be removed from the server. If this value is `Yes`, the disk can be removed from the server.

Status information

You can create capacity and relationship reports that include the following information:

Server Disk Status
The condition of the resource, for example normal, warning, or error.

Server Disk Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Server Disk Firmware Version
The version number of the firmware that is running on the disk.

Server Disk Model
The model name or model number of the resource.

Server Disk Serial Number
The serial number of the resource.

Server Disk Vendor
The vendor who supplied the resource.

Path information

You can create capacity and relationship reports that include the following information:

Server Disk LUN
The logical unit within a small computer system interface target on a disk.

Server Disk SCSI Target
The bus address for the small computer system interface of the target on a disk.

Server Disk Target WWPN
The worldwide port name of the storage port that the volume is mapped to.

Data for server groups in capacity and relationship reports

You can include general information and properties of server groups in capacity and relationship reports.

Information about server groups

You can create capacity and relationship reports that include the following information:

Server Group Name
The name that was assigned to the server group when it was added to the system.

Component properties

You can create capacity and relationship reports that include the following information:

Server Group Creator
The user name of the user who created the group.

Server Group Last Modified
The date and time that the resource was last modified.

Server Group Last Modified by User
The user name of the user who last modified the resource.

Data for file system groups on servers in capacity and relationship reports

You can include general information about file system groups on servers, and properties of file system groups in capacity and relationship reports.

Information about file system groups on servers

You can create capacity and relationship reports that include the following information:

File System Group Name
The name that was assigned to the file system group when it was added to the system.

Component properties

You can create capacity and relationship reports that include the following information:

File System Group Creator
The user name of the user who created the group.

File System Group Last Modified
The date and time that the resource was last modified.

File System Group Last Modified by User
The user name of the user who last modified the resource.

Data for hypervisors in capacity and relationship reports

You can include general information, capacity data, properties, and other information about hypervisors in capacity and relationship reports.

Information about hypervisors

You can create capacity and relationship reports that include the following information:

Hypervisor Name
The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

Hypervisor Short Name
The host name from the fully qualified domain name of the hypervisor. For example, if the fully qualified domain name of the hypervisor is `mycomputer.example.com`, the host name is `mycomputer`.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Hypervisor Kernel Memory (GiB)
The amount of memory on a hypervisor that is used for kernel memory. Kernel memory is used to perform internal operating system processes.

Hypervisor Service Console Memory (GiB)
The amount of memory on a hypervisor that is used for the VMware ESX Service Console. Service Console is the operating system that is used to interact with ESX and the virtual machines that run on the hypervisor.

Hypervisor Swap Space (GiB)
The amount of disk space on a resource that is available to store some of the contents of the RAM when the RAM is full.

Hypervisor Total Disk Capacity (GiB)
The amount of storage space on disks on a hypervisor.

Hypervisor Total Disk Available Space (GiB)
The amount of unused storage space on disks on the hypervisor.

Hypervisor Total File System Available Space (GiB)
The amount of unused storage space that is available on the file systems on a hypervisor.

Component properties

You can create capacity and relationship reports that include the following information:

Hypervisor IP Address
The IP address of the resource.

Hypervisor OS Type
The software that runs the hypervisor, that is, VMware ESX.

Hypervisor OS Version
The version number of the operating system that is running on the resource.

Hypervisor Processor Architecture
The architecture of the processor on the server or hypervisor. For example, the architecture of a processor might be Intel 64 bit (IA64) or Intel 32 bit (IA32).

Hypervisor Processor Count
The number of processors on the server or hypervisor.

Hypervisor Processor Speed (MHz)
The speed of the processor on the server or hypervisor.

Hypervisor Processor Type
Shows information about the processor, such as the family and model of the processor.

Hypervisor RAM (GiB)
The amount of RAM on the server or hypervisor.

Hypervisor Software API Version
The version number of the software API that runs on the hypervisor.

Hypervisor Software Full Name
The name, version number, and build number of the software that runs the hypervisor.

Hypervisor Time Zone
The time zone in which a resource is located.

Hypervisor Location
The physical location of a hypervisor. The location is defined when a hypervisor is added to Tivoli Storage Productivity Center. You can add or edit the location of the hypervisor in the **Properties** pane of the hypervisor in the web-based GUI.

Hypervisor Custom Tag 1, 2, and 3
User-defined text that is associated with a hypervisor. You can add or edit the custom tags for a hypervisor in the **Properties** pane of the hypervisor in the Tivoli Storage Productivity Center web-based GUI.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Hypervisor Model
The model name or model number of the resource.

Hypervisor Serial Number
The serial number of the resource.

Hypervisor Vendor
The vendor who supplied the resource.

Status information

You can create capacity and relationship reports that include the following information:

Hypervisor Status
The condition of the resource, for example normal, warning, or error.

Hypervisor Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Hypervisor Last Data Collection
The date and time when storage statistics were last collected from the resource.

Hypervisor Last Start Time
The last time that the resource was started.

Data for clusters on hypervisors in capacity and relationship reports

You can include general information about clusters on hypervisors in capacity and relationship reports.

Information about clusters on hypervisors

You can create capacity and relationship reports that include the following information:

Hypervisor Cluster Name
The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

Hypervisor Cluster Type
The type of a cluster that is monitored in your storage environment. For example, the type might be VMware Cluster.

Data for file systems on hypervisors in capacity and relationship reports

You can include general information, capacity data, and other information about file systems on hypervisors in capacity and relationship reports.

Information about file systems on hypervisors

You can create capacity and relationship reports that include the following information:

Hypervisor File System Mount Point
The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be `c:\` or `d:\`. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be `/opt` or `/export/home`.

Hypervisor File System Export Name
The name of the exported file system.

Hypervisor File System Type
The type of file system that the hypervisor uses.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Hypervisor File System Capacity (GiB)
The amount of storage space on the file system of the resource.

Hypervisor File System Directory Count
The number of directories in the file systems that are on a server.

Hypervisor File System File Count
The number of files in the file system.

Hypervisor File System Available Space (GiB)
The amount of unused storage space in the file system of the resource.

Hypervisor File System Available Inodes
The number of unused inodes in file systems on the operating system.

Hypervisor File System Used Inodes
The number of used inodes in file systems on the operating system.

Tip: For Microsoft Windows systems, this property is blank.

Hypervisor File System Maximum Files
The maximum number of files that the file system on the resource can contain.

Hypervisor File System Physical Size (GiB)
The amount of physical storage space on the file systems on a resource. The physical size is the size of all the clusters that the file system uses.

Hypervisor File System Used Space (GiB)
The amount of used storage space in the file system of the resource.

Status information

You can create capacity and relationship reports that include the following information:

Hypervisor File System Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Hypervisor File System Last Scan
The date and time when file statistics were last collected from the resource.

Data for data stores on hypervisors in capacity and relationship reports

You can include general information, capacity data, and other information about data stores on hypervisors in capacity and relationship reports.

Information about data stores on hypervisors

You can create capacity and relationship reports that include the following information:

Hypervisor Data Store
The path to the VMware ESX data store on the hypervisor, for example `datastore_svc3c`.

Hypervisor Data Store Cluster Name
The name of the cluster, if the data store is a member of a data store cluster. If the data store is not a member of a data store cluster, no information is displayed.

Hypervisor Data Store Type
The type of file system that the hypervisor data store uses. For example, the type can be NFS or VMFS.

Hypervisor NAS Name
The name of the network-attached storage that the data store is obtained from.

Hypervisor NAS Path Name
The path name on the network-attached storage of the data store.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Hypervisor Data Store Capacity (GiB)
The amount of storage space that is on the hypervisor data store.

Hypervisor Data Store Available Space (GiB)
The amount of unused storage space on a hypervisor data store.

Hypervisor Data Store Allocation on Hypervisor Disk (GiB)
The amount of storage space that is allocated to the hypervisor data store on hypervisor server disks.

Hypervisor Data Store VM File System Version
The version number of the virtual machine file system on the hypervisor data store.

Status information

You can create capacity and relationship reports that include the following information:

Hypervisor Data Store Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Data for server disks on hypervisors in capacity and relationship reports

You can include general information, capacity data, path information, and other information about server disks on hypervisors in capacity and relationship reports.

Information about server disks

You can create capacity and relationship reports that include the following information:

Hypervisor Disk Path
The identifier that specifies a path to a hypervisor server disk, for example `vmhba1:0:0`.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Hypervisor Disk Capacity (GiB)
The amount of storage space that is on the hypervisor disk.

Hypervisor Disk Available Space (GiB)
The amount of unused storage space on a hypervisor server disk.

Status information

You can create capacity and relationship reports that include the following information:

Hypervisor Disk Status
The condition of the resource, for example normal, warning, or error.

Hypervisor Disk Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Hypervisor Disk Firmware Version
The version number of the firmware that is running on the disk.

Hypervisor Disk Model
The model name or model number of the resource.

Hypervisor Disk Serial Number
The serial number of the resource.

Hypervisor Disk Vendor
The vendor who supplied the resource.

Path information

You can create capacity and relationship reports that include the following information:

Hypervisor Disk LUN
The logical unit within a small computer system interface target on a disk.

Hypervisor Disk SCSI Target
The bus address for the small computer system interface of the target on a disk.

Hypervisor Disk Target WWPN
The worldwide port name of the storage port that the volume is mapped to.

Data for virtual machines managed by hypervisors in capacity and relationship reports

You can include general information, and file and disk information for virtual machines in capacity and relationship reports.

Information about virtual machines

You can create capacity and relationship reports that include the following information:

Virtual Machine Name
The name that was assigned to the virtual machine when it was added to the system.

Virtual Machine Configuration File
The path to the virtual machine configuration file for a hypervisor virtual machine, for example `[datastore1] example/example.vmx`.

Virtual Machine Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Virtual Machine Number of CPU
The number of processors on the virtual machine.

Virtual Machine Status
The condition of the resource, for example normal, warning, or error.

Virtual Machine Total Memory (GiB)
The amount of memory on a hypervisor virtual machine.

Virtual machine file information

You can create capacity and relationship reports that include the following information:

Virtual Machine File Path
The path to the virtual machine configuration file, or to the virtual machine disk file for the virtual machine. The virtual machine configuration file is a `.vmx` file, and the virtual machine disk file is a `.vmdk` file.

Virtual Machine File Logical Size (GiB)
The logical size of a virtual machine disk file for a hypervisor virtual machine.

Virtual Machine File Type
The type of the file for the virtual machine. For example, the file can be a virtual machine configuration file, a virtual machine disk file, an ISO image file, or other types of virtual machine file.

Virtual machine disk information

You can create capacity and relationship reports that include the following information:

Virtual Machine Disk Name
The name that was assigned to the virtual machine disk when it was added to the system.

Virtual Machine Disk SCSI Target
The bus address for the small computer system interface of the target on a disk.

Virtual Machine Disk Type
The type of virtual machine disk, for example, collection mirroring or logical disk mirroring.

Data for network-attached storage systems in capacity and relationship reports

You can include general information, capacity data, properties, and other information about network-attached storage (NAS) systems in capacity and relationship reports.

Information about NAS systems

You can create capacity and relationship reports that include the following information:

NAS Name
The fully qualified domain name of the NAS device. For example, the name of a NAS device might be `nas.example.com`.

NAS Domain Name
The domain name of the cluster, server, or NAS server.

NAS Short Name
The host name from the fully qualified domain name of the NAS device. For example, if the fully qualified domain name of the NAS device is `mycomputer.example.com`, the host name is `mycomputer`.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

NAS Total Disk Capacity (GiB)
The amount of storage space on all disks on a NAS resource.

NAS Total Disk Available Space (GiB)
The amount of unused storage space on all disks on a NAS resource.

NAS Total File System Available Space (GiB)
The amount of unused storage space on the file systems on a NAS resource.

Component properties

You can create capacity and relationship reports that include the following information:

NAS OS Type
The operating system that is running on the server.

NAS OS Version
The version number of the operating system that is running on the resource.

NAS Location
The physical location of a NAS system that is probed as a filer. You can add or edit the location of the NAS system on the **File Storage** tab in the **Storage Systems** pane in the web-based GUI.

NAS Custom Tag 1, 2, and 3
User-defined text that is associated with a NAS system that is probed as a filer. You can add or edit the custom tags for a NAS system on the **File Storage** tab in the **Storage Systems** pane in the Tivoli Storage Productivity Center web-based GUI.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

NAS Model
The model name or model number of the resource.

NAS Vendor
The vendor who supplied the resource.

NAS Serial Number
The serial number of the resource.

Status information

You can create capacity and relationship reports that include the following information:

NAS Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

NAS Last Data Collection
The date and time when storage statistics were last collected from the resource.

NAS Last Start Time
The last time that the resource was started.

NAS system components information

You can create capacity and relationship reports that include the following information:

NAS System Name
The name that was assigned to the NAS system when it was added to the system.

NAS System IP Address
The IP address of the resource.

NAS System Location
The physical location of a NAS system that is probed as a storage system. The location is defined when a NAS system is added to Tivoli Storage Productivity Center. You can add or edit the location of the NAS system on the **File Storage** tab in the **Storage Systems** pane in the web-based GUI.

NAS System Custom Tag 1, 2, and 3
User-defined text that is associated with a NAS system that is probed as a

storage system. You can add or edit the custom tags for a NAS system on the **File Storage** tab in the **Storage Systems** pane in the Tivoli Storage Productivity Center web-based GUI.

Data for file systems on network-attached storage systems in capacity and relationship reports

You can include general information, capacity data, and other information about file systems on network-attached storage (NAS) systems in capacity and relationship reports.

Information about file systems on NAS systems

You can create capacity and relationship reports that include the following information:

NAS File System Mount Point
The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be `c:\` or `d:\`. On operating systems such as AIX, Linux, Solaris, or HP-UX, the mount point might be `/opt` or `/export/home`.

NAS File System Export Name
The name of the exported file system.

NAS File System Type
The type of file system that the resource uses.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

NAS File System Capacity (GiB)
The amount of storage space on the file system of the resource.

NAS File System Directory Count
The number of directories in the file systems that are on a server.

NAS File System File Count
The number of files in the file system.

NAS File System Available Space (GiB)
The amount of unused storage space in the file system of the resource.

NAS File System Available Inodes
The number of unused inodes on the file system on a NAS server.

NAS File System Used Inodes
The number of used inodes on the file system on a NAS server.

NAS File System Maximum Files
The maximum number of files that the file system on the resource can contain.

NAS File System Physical Size (GiB)
The amount of physical storage space on the file systems on a resource. The physical size is the size of all the clusters that the file system uses.

NAS File System Used Space (GiB)
The amount of used storage space in the file system of the resource.

Component properties

You can create capacity and relationship reports that include the following information:

NAS File System Custom Tag 1, 2, and 3
User-defined text that is associated with a NAS file system. You can add or edit the custom tags for a NAS file system in the **Properties** pane of the resource in the Tivoli Storage Productivity Center web-based GUI.

Status information

You can create capacity and relationship reports that include the following information:

NAS File System Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

NAS File System Last Scan
The date and time when file statistics were last collected from the resource.

Data for network-attached storage exports in capacity and relationship reports

You can include general information and status information about network-attached storage (NAS) exports in capacity and relationship reports.

Information about network-attached storage exports

You can create capacity and relationship reports that include the following information:

NAS Export Name
The name of the network-attached storage export. Use this name to mount the network-attached storage export on a server.

NAS Export Path
The path to a logical volume on a resource, for example `/dev/hd1`.

NAS Export State
Shows whether the network-attached storage export is active.

Status information

You can create capacity and relationship reports that include the following information:

NAS Export Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Data for logical volumes on network-attached storage systems in capacity and relationship reports

You can include general information, capacity data, and other information about logical volumes on network-attached storage (NAS) systems in capacity and relationship reports.

Information about logical volumes on NAS systems

You can create capacity and relationship reports that include the following information:

NAS Logical Volume Name
The name that was assigned to the NAS logical volume when it was added to the system.

NAS Logical Volume Type
The type of NAS logical volume. The logical volume type can be a NetApp traditional volume or a NetApp flexible volume.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

NAS Logical Volume Capacity (GiB)
The amount of storage space on the NAS logical volume.

NAS Logical Volume Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Status information

You can create capacity and relationship reports that include the following information:

NAS Logical Volume Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

NAS Logical Volume Last Scan
The date and time when file statistics were last collected from the resource.

Data for network-attached storage disks in capacity and relationship reports

You can include general information, capacity data, properties, and other information about network-attached storage (NAS) disks in capacity and relationship reports.

Information about NAS disks

You can create capacity and relationship reports that include the following information:

NAS Disk Name
The name that was assigned to the NAS disk when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

NAS Disk Total Capacity (GiB)
The amount of storage space that is on the NAS disk.

NAS Disk Available Space (GiB)
The amount of unused storage space on a NAS disk.

Status information

You can create capacity and relationship reports that include the following information:

NAS Disk Status
The condition of the resource, for example normal, warning, or error.

NAS Disk Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

IBM Scale Out Network Attached Storage and Storwize V7000 Unified information

You can create capacity and relationship reports that include the following information:

NSD Available Fragments
The number of fragments of disk space that are unused on the Network Shared Disk (NSD).

NSD Failover Group
The failover group that the NSD belongs to. A failover group is a group of resources that are grouped together to ensure that some of those resources are available if one or more of the resources fail.

NSD Type
The type of the NSD. For example, the type can be:

- Data and metadata
- Metadata only
- Data only

NSD Last Data Collection
The date and time when storage statistics were last collected from the resource.

NSD Custom Tag 1, 2, and 3
User-defined text that is associated with a Network Shared Disk (NSD). You can add or edit the custom tags for an NSD in the **Properties** pane of the resource in the Tivoli Storage Productivity Center web-based GUI.

NetApp information

You can create capacity and relationship reports that include the following information:

NAS Disk Flag
Shows whether the NAS disk is a spare device. If this value is `Spare device`, the NAS disk is a spare device.

NAS Disk Is Encrypted
Shows whether the resource is encrypted. If this value is `Yes`, the resource is encrypted.

NAS Disk Is Removable
Shows whether the disk can be removed from the server. If this value is `Yes`, the disk can be removed from the server.

NAS Disk Is Solid State
Shows whether the disk is a solid-state drive.

NetApp paths information

You can create capacity and relationship reports that include the following information:

NAS Disk LUN
The logical unit within a small computer system interface target on a disk.

NAS Disk SCSI Target
The bus address for the small computer system interface of the target on a disk.

NAS Disk Target WWPN
The worldwide port name of the storage port that the volume is mapped to.

NetApp vendor, model, and device information

You can create capacity and relationship reports that include the following information:

NAS Disk Firmware Version
The version number of the firmware that is running on the disk.

NAS Disk Model
The model name or model number of the resource.

NAS Disk Serial Number
The serial number of the resource.

NAS Disk Vendor
The vendor who supplied the resource.

Data for pools on IBM SONAS and Storwize V7000 Unified systems in capacity and relationship reports

You can include general information and capacity data about pools on IBM Scale Out Network Attached Storage and Storwize V7000 Unified systems in capacity and relationship reports.

Information about pools on network-attached storage systems

You can create capacity and relationship reports that include the following information:

NAS Pool Name
The name that was assigned to the NAS pool when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

NAS Pool Capacity (GiB)
The amount of storage space that is in a NAS pool.

NAS Pool Available Space (GiB)
The amount of unused storage space in a NAS pool.

NAS Pool Number of Logical Volumes
The number of logical volumes on a resource.

Data for IBM SONAS and Storwize V7000 Unified file sets in capacity and relationship reports

You can include general information, capacity data, and other information about IBM Scale Out Network Attached Storage and Storwize V7000 Unified file sets in capacity and relationship reports.

Information about network-attached storage file sets

You can create capacity and relationship reports that include the following information:

NAS File Set Name
The name that was assigned to the network-attached storage (NAS) file set when it was added to the system.

NAS File Set Comment
A comment that is associated with the NAS file set.

NAS File Set Path
The path to a logical volume on a resource, for example `/dev/hd1`.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

NAS File Set Used Space (GiB)
The amount of used storage space in the file set of the NAS resource.

NAS File Set Used Inodes
The number of used inodes in file systems on operating systems such as AIX, Linux, Solaris, or HP-UX.

Status information

You can create capacity and relationship reports that include the following information:

NAS File Set Status
Shows whether the network-attached storage file set is linked or unlinked.

NAS File Set Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Data for storage virtualizers in capacity and relationship reports

You can include general information, capacity data, properties, and other information about storage virtualizers in capacity and relationship reports.

Information about storage virtualizers

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage virtualizer was added for monitoring.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Total Pool Capacity (GiB)
The total amount of storage space in pools on the storage virtualizer.

Storage Virtualizer Pool Used Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can pre-allocate thin-provisioned volume space when the volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. For pools with compressed volumes on SAN Volume Controller and Storwize V7000, the Storage Virtualizer Pool Used Space reflects the size of compressed data written to disk. As the data changes the Storage Virtualizer Pool Used Space may at times be less than the Storage Virtualizer Pool Allocated Space.

For pools with volumes that are not thin provisioned or compressed in SAN Volume Controller, Storwize V7000, and other storage systems, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a pool. Because data collection is run on a set schedule and the used space on volumes can change rapidly, the value in this column might not be 100% accurate for the current state of volumes.

Storage Virtualizer Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Storage Virtualizer Total Volume Capacity (GiB)
The amount of virtual volume space on all volumes on the storage virtualizer.

Storage Virtualizer Volume Allocated Space (GiB)
The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000
- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

Storage Virtualizer Volume Assigned Space (GiB)
The amount of virtual volume space that is assigned to the server or storage virtualizer.

Storage Virtualizer MDisk Available Space (GiB)
The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

Storage Virtualizer MDisk Capacity (GiB)
The amount of storage space on the managed disk.

Storage Virtualizer Disk Capacity (GiB)
The amount of local disk storage on the storage virtualizer. The local disks are the internal disk drives.

Storage Virtualizer Number of MDisks
The number of managed disks on a resource.

Storage Virtualizer Number of Volumes
The number of volumes on a resource.

Storage Virtualizer Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Virtualizer Tier Available Space (GiB)
The amount of unused tier storage of a particular type. The type is identified in the Storage Virtualizer Tier Type property.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Tier Capacity (GiB)
The amount of tier storage of a particular type. The type is identified in the Storage Virtualizer Tier Type property.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Physical Allocation Percentage
The percentage of physical space in storage virtualizer pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

`(allocated space ÷ pool capacity) × 100`

For example, the physical allocation percentage is 25% for a total storage virtualizer pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Virtualizer Virtual Allocation Percentage
The percentage of physical space in storage virtualizer pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(total volume capacity ÷ pool capacity) × 100`

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB.

This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage Virtualizer Shortfall Percentage property to estimate how critical the shortage of space is for storage virtualizer pools.

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Shortfall Percentage
The percentage of the remaining unallocated volume space in storage virtualizer pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

```
(unallocatable space ÷ (volume space - used volume space)) × 100
```

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space.

This Storage Virtualizer Shortfall Percentage represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB ÷ 100 GiB × 100).

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

This property is available only for storage virtualizer pools with thin-provisioned volumes.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Access Information
The URL from which you can access the storage system or storage virtualizer.

Storage Virtualizer Alias
Shows an alternative name for a storage virtualizer.

Storage Virtualizer Code Level
The microcode or firmware level of the storage virtualizer.

Storage Virtualizer Time Zone
The time zone in which a resource is located.

Storage Virtualizer Location
The physical location of a storage virtualizer. The location is defined when a storage virtualizer is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage virtualizer in the **Properties** pane for the storage system in the web-based GUI.

Storage Virtualizer Custom Tag 1, 2, and 3
User-defined text that is associated with a storage virtualizer. You can add or edit the custom tags for a storage virtualizer on the **Properties** page of the storage system in the Tivoli Storage Productivity Center web-based GUI.

Storage Virtualizer User Provided Name
The name that was specified for the storage virtualizer in the storage environment.

Storage Virtualizer Tier Type
The type of tier that is used on the storage virtualizer. For example, the tier type can be hard disk drive (HDD) or solid-state drive (SSD).

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Is Compression Active
Shows whether the compression feature is enabled on the storage virtualizer. If this value is `Yes`, the compression feature is enabled.

Storage Virtualizer Configuration
Shows that a storage virtualizer is configured to virtualize back-end storage. This value is always **Storage Virtualizer**.

Status information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Status
The condition of the resource, for example normal, warning, or error.

Storage Virtualizer Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Storage Virtualizer Last Data Collection
The date and time when storage statistics were last collected from the resource.

Storage Virtualizer Last Data Collection Status
The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Type
The type of storage virtualizer. For example, the storage virtualizer can be an **IBM Storwize V7000** system, an **IBM SAN Volume Controller** system, or another type of storage virtualizer.

Storage Virtualizer Vendor
The vendor who supplied the resource.

Storage Virtualizer Model
The model name or model number of the storage virtualizer.

Storage Virtualizer Machine Type
The machine type of a storage virtualizer. For example, the storage virtualizer can be **Storwize V7000 - 2076**, **SAN Volume Controller - 2145**, or another type of storage virtualizer.

Storage Virtualizer IP Address
The IP address of the resource.

Storage Virtualizer Serial Number
The serial number of the resource.

Data for storage virtualizer pools in capacity and relationship reports

You can include general information, capacity data, properties, and other information about storage virtualizer pools in capacity and relationship reports.

Information about storage virtualizer pools

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Pool Name
The name that was assigned to the pool when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Pool Capacity (GiB)
The total amount of storage space in a pool.

Storage Virtualizer Pool Available Space (GiB)
The amount of unused storage space in a storage virtualizer pool.

Storage Virtualizer Pool Total Volume Capacity (GiB)
The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Virtualizer Pool Allocated Space (GiB)
The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes.

The space that is allocated for thin-provisioned volumes is less than their virtual capacity. The Storage Virtualizer Pool Total Volume Capacity (GiB) property shows the virtual capacity of the volumes. If a pool does not contain thin-provisioned volumes, this value is the same as the value of the Storage Virtualizer Pool Total Volume Capacity (GiB) property.

This value is equal to Storage Virtualizer Used Space (GiB) for volumes that are not thin provisioned.

Storage Virtualizer Pool Used Volume Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Virtualizer Pool Extent Size (MiB)
The size of the extents that were specified when a pool was created. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Assigned Volume Space (GiB)
The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Virtualizer Pool Unassigned Volume Space (GiB)
The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Virtualizer Pool Tier Capacity (GiB)
The amount of pool tier storage of a particular type. The type is identified in the Storage Virtualizer Pool Tier Type property.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Tier Available Space (GiB)
The amount of pool tier storage of a particular type that is unused. The type is identified in the Storage Virtualizer Tier Type property.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Tier Number of MDisks
The number of managed disks on a storage virtualizer pool. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Number of Storage Volumes
The number of volumes in the storage virtualizer pool.

Storage Virtualizer Pool Number of MDisks
The number of managed disks on a resource.

Storage Virtualizer Pool Maximum I/O Capability
The projected maximum number of I/O operations per second for a pool. This value is calculated based on the value in the Storage Virtualizer Pool Back-End Storage Disks property, and on the values in following properties:

- Storage Virtualizer Pool Back-End Storage System Type
- Storage Virtualizer Pool Back-End Storage RAID Level
- Storage Virtualizer Pool Back-End Storage Disk Type

This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Storage Virtualizer Pool Back-End Storage Disks
The number of physical disks that contribute to the volumes on the back-end storage system. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Storage Virtualizer Pool Physical Allocation Percentage
The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

```
(allocated space ÷ pool capacity) × 100
```

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Virtualizer Pool Virtual Allocation Percentage
The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

```
(total volume capacity ÷ pool capacity) × 100
```

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. In this case, the virtual capacity that is committed to the volumes in the pool is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, use the Storage Virtualizer Pool Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

This value is only available for pools with thin-provisioned volumes.

Storage Virtualizer Pool Shortfall Percentage
The percentage of the remaining unallocated volume space in a pool that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

```
(unallocatable space ÷ (total volume capacity − allocated space)) × 100
```

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Unallocated Volume Space (GiB)
The amount of the Total Volume Capacity in the pool that is not allocated.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

```
total volume capacity − allocated space
```

The space that is allocated for thin-provisioned volumes is typically less than their virtual capacity. Therefore, the unallocated space represents the difference between the virtual capacity and the allocated space for all the volumes in the pool.

Storage Virtualizer Pool Unused Volume Space (GiB)
The amount of space that is allocated to the volumes in a pool and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
allocated space − used pool space
```

This property is available only for SAN Volume Controller and Storwize V7000 pools.

Storage Virtualizer Pool Unallocatable Volume Space (GiB)
The amount of space by which the Total Volume Capacity exceeds the physical capacity of a pool.

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool. If you create volumes with more virtual capacity than can be physically allocated in the pool, you can over commit storage in the pool.

This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.

Storage Virtualizer Pool Used Space Percentage
The percentage of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

Storage Virtualizer Pool Compressed Virtual Capacity (GiB)
The total virtual capacity of all the volumes that are compressed in a storage virtualizer pool.

Storage Virtualizer Pool Compressed Capacity (GiB)
The amount of storage space that is used by compressed volumes in a storage virtualizer pool. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 20.

Storage Virtualizer Pool Uncompressed Capacity (GiB)
The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Pool Is Encryptable
Shows whether the storage virtualizer pool can be encrypted. If this value is `Yes`, the storage virtualizer pool can be encrypted.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Is Encrypted
Shows whether the storage virtualizer pool is encrypted. If this value is `Yes`, the storage virtualizer pool is encrypted.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Is Solid State
Shows whether there are solid-state drives in the pool. This property can contain the following values:

Mixed The pool contains both hard disk drives and solid-state drives.

Non solid state
The pool contains no solid-state drives.

Solid state
The pool contains a least one solid-state drive.

Storage Virtualizer Pool Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Virtualizer Pool RAID Level
The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Virtualizer Pool Warning Level
The percentage of used capacity of the resource at which a warning is generated. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Tier
The tier that a managed disk is assigned to by the internal cluster automatic detection for internal disks or by the user for external disks. This classification is used by the storage virtualizer to determine when to activate the Easy Tier function for certain pools.

Storage Virtualizer Pool Tier Type
The type of tier that is used on the storage virtualizer pool. For example, the tier type can be hard disk drive (HDD) or solid-state drive (SSD).

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Is Easy Tier
Shows whether the IBM System Storage Easy Tier feature is enabled. The IBM System Storage Easy Tier feature can be set to `auto`, `on`, or `off`.

The IBM System Storage Easy Tier feature automatically helps optimize solid-state storage deployments in multitier systems.

Storage Virtualizer Pool Back-End Storage System Type
The type of storage system that is providing storage space to a pool. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Storage Virtualizer Pool Back-End Storage RAID Level
The RAID level of the volumes on the back-end storage system that are providing storage space to a pool. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Storage Virtualizer Pool Back-End Storage Disk Type
The class and speed of the physical disks that contribute to the volumes on the back-end storage system. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Storage Virtualizer Pool Custom Tag 1, 2, and 3
User-defined text that is associated with a storage virtualizer pool. You can add or edit the custom tags for a storage virtualizer pool on the **Properties** notebook of the pool in the Tivoli Storage Productivity Center web-based GUI.

Storage Virtualizer Pool Is Compression Active
Shows whether the compression feature is enabled on the storage virtualizer pool. If this value is `Yes`, the compression feature is enabled.

Status information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Pool Status
The condition of the resource, for example normal, warning, or error.

Storage Virtualizer Pool Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Storage Virtualizer Pool Easy Tier Status
Shows whether the Easy Tier feature is active or inactive. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Data for storage virtualizer volumes in capacity and relationship reports

You can include general information, capacity data, properties, information about copies of volumes, and other information about storage virtualizer volumes in capacity and relationship reports.

Information about storage virtualizer volumes

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Volume Name
The name that was assigned to the storage virtualizer volume when it was added to the system.

Storage Virtualizer Volume Assigned Host Connection
The host connection to which the storage virtualizer volume is assigned. The host connection is a definition in the storage virtualizer that contains the WWPN for the server. The storage virtualizer uses the WWPN to assign volumes to servers.

Storage Virtualizer Volume Host Connection OS Type
The operating system type of the server or hypervisor that the virtualizer volume is assigned to.

Storage Virtualizer Volume Preferred Node Name
The number of the storage virtualizer node that is used for I/O operations for the storage virtualizer volume.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Volume Capacity (GiB)
The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

Storage Virtualizer Volume Allocated Space (GiB)
The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following storage systems:

- Storage systems other than SAN Volume Controller and Storwize V7000

- SAN Volume Controller and Storwize V7000 storage systems that are not thin-provisioned

Storage Virtualizer Volume Used Space (GiB)
The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Virtualizer Volume Used Space (GiB) might be different than the Storage Virtualizer Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Virtualizer Volume Used Space (GiB) and Storage Virtualizer Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Virtualizer Volume Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Virtualizer Volume Used Allocated Space Percentage
The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

```
(used space ÷ allocated space) × 100
```

This property is available only for SAN Volume Controller and Storwize V7000 volumes.

Storage Virtualizer Volume Unallocated Space (GiB)
The amount of space in a pool that is not reserved for a volume.

Tivoli Storage Productivity Center uses the following formula to determine the value:

```
capacity - allocated space
```

Storage Virtualizer Volume Physical Allocation Percentage
The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

```
(allocated space ÷ capacity) × 100
```

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore the space that is reserved for volumes is 50 GiB.

Storage Virtualizer Volume Unused Space (GiB)
The amount of space that is allocated to a volume and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
allocated space - used Space
```

This property is available only for SAN Volume Controller and Storwize V7000 pools.

Storage Virtualizer Volume Shortfall Percentage
The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume. The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(unallocatable volume space ÷ unallocated volume space) × 100`

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level for the volume. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

`(100 GiB unallocated volume space - 70 GiB remaining pool space) ÷ 100 GiB unallocated volume space × 100`

This value is only available for pools with thin-provisioned volumes.

Storage Virtualizer Volume Unallocatable Space (GiB)
The unallocatable storage space in GiB of a thin-provisioned volume. Unallocatable space cannot be supplied by the pool to which the volume belongs.

Storage Virtualizer Volume Uncompressed Capacity (GiB)
The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Volume Type
The type of storage virtualizer volume. For example, the storage virtualizer volume can be striped, sequential, or another type of storage virtualizer volume. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Volume WWN
The worldwide name of the volume.

Storage Virtualizer Volume Fast Write State
Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted.

Storage Virtualizer Volume Grain Size (KiB)
The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Volume Is Assigned
Shows whether the volume is assigned to a server. If this value is `Yes`, the volume is assigned to a server.

Storage Virtualizer Volume Is Auto Expand Enabled
Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is `Yes`, the volume or volume copy can allocate new extents from a pool automatically.

Storage Virtualizer Volume Is Encryptable
Shows whether the resource can be encrypted. If this value is `Yes`, the resource can be encrypted.

Storage Virtualizer Volume Is Encrypted
Shows whether the resource is encrypted. If this value is `Yes`, the resource is encrypted.

Storage Virtualizer Volume Is Formatted
Shows whether a storage virtualizer volume is formatted. If this value is `Yes`, the storage virtualizer volume is formatted.

Storage Virtualizer Volume Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Virtualizer Volume Metro Mirror Name
The name of the Metro Mirror that keeps the synchronous copy of the storage virtualizer volume.

Storage Virtualizer Volume Mirror Count
The number of mirrors that keep a synchronous copy of the resource.

Storage Virtualizer Volume Throttle
The maximum number of commands that the storage virtualizer volume can queue.

Storage Virtualizer Volume Warning Level
The percentage of used capacity of the resource at which a warning is generated.

Storage Virtualizer Volume FlashCopy Relationship
Shows whether a volume on a storage virtualizer is in a FlashCopy relationship. This property can contain the following values:

Source Relationship
The volume is the source of the relationship.

Target Relationship
The volume is the target of the relationship.

Both Source and Target Relationship
The volume is both the source and target of the relationship.

Not in a FlashCopy Relationship
The volume is not in a FlashCopy relationship.

Storage Virtualizer Volume Remote Copy Relationship
Shows whether a volume on a storage virtualizer is in a remote copy relationship. This property can contain the following values:

Not in a Copy Relationship
The volume is not in a remote copy relationship.

Source Relationship
The volume is the source of the relationship.

Target Relationship
The volume is the target of the relationship.

Both Source and Target Relationship
The volume is both the source and target of the relationship.

Storage Virtualizer Volume RAID Level
The RAID level of a storage virtualizer volume, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Virtualizer Volume Is Compressed
Shows whether the storage virtualizer volume is compressed. If this value is `Yes`, the volume is compressed.

Status information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Volume Status
The condition of the resource, for example normal, warning, or error.

Storage Virtualizer Volume Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Storage Virtualizer Volume Native Status
Shows the level of access of nodes in the system to the storage virtualizer managed disk or volume. The level of access can be online, offline, degraded, or excluded.

Volume copy information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Volume Copy ID
The identifier of a volume copy. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Volume Copy Grain Size (KiB)
The grain size with which a thin-provisioned volume was created. This

value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

Storage Virtualizer Volume Copy Is Auto Expand Enabled
Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is `Yes`, the volume or volume copy can allocate new extents from a pool automatically.

Storage Virtualizer Volume Copy Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Storage Virtualizer Volume Copy Is Primary
Shows whether the volume copy is the primary volume copy. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Volume Copy Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Virtualizer Volume Copy Total Number of Extents
The number of extents that are used by the storage virtualizer volume copy.

Storage Virtualizer Volume Copy Type
The type of copy of the storage virtualizer volume. For example, the copy of the storage virtualizer volume can be sequential, a mirror copy, or another type of copy.

Storage Virtualizer Volume Copy Warning Level
The percentage of used capacity of the resource at which a warning is generated.

Storage Virtualizer Volume Copy Used Allocated Space Percentage
The percentage of reserved space for a volume copy that is being used. This value is always less than or equal to 100% because a volume copy cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

`(used space ÷ allocated space) × 100`

This property is available only for SAN Volume Controller and Storwize V7000 pools.

Storage Virtualizer Volume Copy Used Space (GiB)
The amount of allocated space that is used by a volume copy.

Storage Virtualizer Volume Copy Allocated Space (GiB)
The amount of space that is reserved for a volume copy.

The space that is allocated for a thin-provisioned volume copy is less than its virtual capacity, which is shown in Storage Virtualizer Capacity (GiB).

This value is equal to Storage Virtualizer Volume Copy Used Space (GiB) for volume copies that are not thin provisioned.

Storage Virtualizer Volume Copy Unallocated Space (GiB)
The amount of space in a pool that is not reserved for a volume copy.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
capacity - allocated space
```

Storage Virtualizer Volume Copy Physical Allocation Percentage
The percentage of physical space that is reserved for a volume copy. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
(allocated Space ÷ capacity) × 100
```

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volume copies is 50 GiB.

Storage Virtualizer Volume Copy Unused Space (GiB)
The amount of space that is allocated to a volume copy and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
allocated Space - used space
```

This property is available only for SAN Volume Controller and Storwize V7000 pools.

Storage Virtualizer Volume Copy Shortfall Percentage
The percentage of the remaining unallocated volume copy space in a pool that is not available to be allocated to a volume copy.

The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
(unallocatable space ÷ (capacity – real capacity)) × 100
```

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level for the volume. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

```
(100 GiB unallocatable volume copy space - 70 GiB
remaining pool space) ÷ 100 GiB unallocated volume space × 100
```

Available only for thin-provisioned volume copies.

Storage Virtualizer Volume Copy Unallocatable Space (GiB)
The unallocatable storage space in GiB of a thin-provisioned volume copy. Unallocatable space cannot be supplied by the pool to which the volume copy belongs.

Storage Virtualizer Volume Copy Is Compressed
Shows whether the volume copy is compressed. If this value is `Yes`, the volume copy is compressed.

Storage Virtualizer Volume Copy Uncompressed Capacity (GiB)
The amount of storage space that is used if the compressed space on the volume copy is uncompressed. For example, if 100 GiB of data is compressed, the value is 100.

Data for managed disks on storage virtualizers in capacity and relationship reports

You can include general information, capacity data, properties, and other information about managed disks on storage virtualizers in capacity and relationship reports.

Information about managed disks on storage virtualizers

You can create capacity and relationship reports that include the following information:

Storage Virtualizer MDisk Name
The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Virtualizer MDisk Capacity (GiB)
The amount of storage space on the managed disk.

Storage Virtualizer MDisk Available Space (GiB)
The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

Storage Virtualizer MDisk Strip Size (KB)
The RAID strip size on a managed disk on a storage virtualizer. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Virtualizer MDisk RAID Level
The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Spare Goal
The number of spare drives that are required to maintain redundancy. Use spare drives to protect against drive failures in the array on a managed disk on a storage virtualizer.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Type
The type of managed disk on a storage virtualizer. For example, the storage virtualizer can be a local or virtual managed disk.

Storage Virtualizer MDisk Mode
The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

Array Extents are provided from local disks.

Managed
Extents are provided from a back-end storage volume.

Unmanaged
The managed disk is not used in the system.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Is Balanced
Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is `Yes`, the LUNs are balanced.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Fast Write State
Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Write Verify
Shows whether all write operations on a managed disk on a storage virtualizer are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is `Yes`, all write operations are verified by a follow-up read operation.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Status information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer MDisk Status
The condition of the resource, for example normal, warning, or error.

Storage Virtualizer MDisk Native Status
Shows the level of access of nodes in the system to a managed disk or volume on a storage virtualizer. The level of access can be online, offline, degraded, or excluded.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Data for local disks on storage virtualizers in capacity and relationship reports

You can include general information, capacity data, properties, and other information about local disks on storage virtualizers in capacity and relationship reports.

Information about local disks on storage virtualizers

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Local Disk Name
The name that was assigned to the local disk on the storage virtualizer when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Local Disk Capacity (GiB)
The amount of storage space on the local disk.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Local Disk Is Encryptable
Shows whether the local disk can be encrypted. If this value is `Yes`, the local disk can be encrypted.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Local Disk Is Encrypted
Shows whether the local disk is encrypted. If this value is `Yes`, the local disk is encrypted.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Local Disk Storage Class
The storage technology of the local disk on a storage virtualizer. For example, the storage class can be a serial-attached SCSI (SAS) or solid-state drive (SSD). This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Local Disk Is Solid State
Shows whether the disk is a solid-state drive. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Local Disk Speed (RPM)
The speed of the local disk on a storage virtualizer. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Local Disk Tag
A number that identifies the array to which a local disk belongs.

Status information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Local Disk Status
The condition of the resource, for example normal, warning, or error.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Local Disk Firmware Version
The version number of the firmware that is running on the local disk. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Local Disk Model
The model name or model number of the resource.

Storage Virtualizer Local Disk Serial Number
The serial number of the resource.

Storage Virtualizer Local Disk Vendor
The vendor who supplied the resource.

Data for storage virtualizer groups in capacity and relationship reports

You can include the name of storage virtualizer groups and properties for storage virtualizer groups in capacity and relationship reports.

Information about storage virtualizer groups

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Group Name
A user-defined name of the storage virtualizer group.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Group Creator
The user name of the user who created the storage virtualizer group.

Storage Virtualizer Group Last Modified
The date and time that the storage virtualizer group was last modified.

Storage Virtualizer Group Last Modified by User
The user name of the user who last modified the storage virtualizer group.

Data for storage systems in capacity and relationship reports

You can include general information, capacity data, properties of components, and other information about storage systems in capacity and relationship reports.

Information about storage systems

You can create capacity and relationship reports that include the following information:

Storage System Name
A user-defined name of the storage system. If a name was not defined, Tivoli Storage Productivity Center shows the name that was defined when the storage system was added for monitoring.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage System Pool Capacity (GiB)
The amount of storage space in pools that are on the storage system. For an XIV system, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage System Used Pool Space (GiB)
The amount of space that is in use in all pools on a storage system.

Storage System Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

Storage System Volume Capacity (GiB)
The amount of space on all volumes on the storage system.

Storage System Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage System Assigned Volume Space (GiB)
The amount of space on the storage system that is on volumes that are assigned to a server or storage virtualizer.

Storage System Unallocated Disk Space (GiB)
The amount of disk space that can be added to a pool.

Storage System Number of Disks
The number of disks on the storage system.

Storage System Number of Volumes
The number of volumes on a resource.

Storage System Cache (GiB)
The size of the cache on the storage system. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

Storage System Real Available Pool Space (GiB)
The amount of unused space in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System Real Configured Pool Space (GiB)
The amount of storage space that is in pools that have an associated soft size on an XIV system. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System Volume Capacity for z/OS (GiB)
The amount of space on all volumes on the storage system that the z/OS operating system can use.

Storage System Volume Capacity Assigned to MDisks (GiB)
The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

Storage System Disk Capacity (GiB)
The amount of storage space on physical disks in a storage system, excluding spare disks. This value is not shown for Storwize V7000 storage systems that are configured as back-end storage.

Storage System Total Disk Capacity (GiB)
The amount of space on physical disks on a storage system, including spare disks.

Storage System Physical Allocation Percentage
The percentage of physical space in storage system pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(allocated pool space ÷ pool capacity) × 100`

For example, the physical allocation percentage is 25% for a 200 GiB storage pool. Therefore, the space that is reserved for volumes is 50 GiB.

Storage System Virtual Allocation Percentage
The percentage of physical space in storage system pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(total volume capacity ÷ pool capacity) × 100`

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage System Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

Storage System Shortfall Percentage
The percentage of the remaining unallocated volume space in storage system pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(unallocatable space ÷ (volume space - used volume space)) × 100`

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

Storage System MDisk Available Space (GiB)
The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage System Tier Available Space (GiB)
The amount of unused space in a tier. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage System Tier Capacity (GiB)
The total amount of storage space in a tier.This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Component properties

You can create capacity and relationship reports that include the following information:

Storage System Access Information
The URL from which you can access the storage system or storage virtualizer.

Storage System Code Level
The Shared Ethernet Adapter level of a DS8000 storage system. For other storage systems, this value is the firmware version.

Storage System Location
The physical location of a storage system. The location is defined when a storage system is added to Tivoli Storage Productivity Center. You can add or edit the location of the storage system in the **Properties** pane of the storage system in the web-based GUI.

Storage System Custom Tag 1, 2, and 3
User-defined text that is associated with a storage system. You can add or edit the custom tags for a storage system in the **Properties** pane of the storage system in the Tivoli Storage Productivity Center web-based GUI.

Storage System User Provided Name
The name that was specified for the storage system in the storage environment.

Storage System Tier Type
The type of tier that is used on a storage system. For example, the tier type can be hard disk drive (HDD) or solid-state drive (SSD).

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is `Storage System`, the storage system is configured as back-end storage. If this value is `Storage Virtualizer`, the storage system is configured to virtualize back-end storage.

Storage System Is Compression Active
Shows whether the compression feature is enabled on the storage system. If this value is `Yes`, the compression feature is enabled.

Status information

You can create capacity and relationship reports that include the following information:

Storage System Status
The condition of the resource, for example normal, warning, or error.

Storage System Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Storage System Last Data Collection
The date and time when storage statistics were last collected from the resource.

Storage System Last Data Collection Status
The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Storage System Type
The type of storage system. For example, the storage system can be an **IBM System Storage DS8800** system, an **IBM System Storage DS8700** system, an **IBM System Storage XIV** system, or another type of storage system.

Storage System Vendor
The vendor who supplied the resource.

Storage System Model
The model name or model number of the storage system.

Storage System Machine Type
The machine type of a storage system. For example, the storage system can be **DS8800**, **DS8700**, **XIV - 2812**, or another type of storage system.

Storage System IP Address
The IP address of the resource.

Storage System Serial Number
The serial number of the resource.

Data for storage system pools in capacity and relationship reports

You can include general information, capacity data, properties of components, and other information about storage system pools in capacity and relationship reports.

Information about storage systems pools

You can create capacity and relationship reports that include the following information:

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Pool Capacity (GiB)
The total amount of storage space in a pool.

Storage Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
pool capacity - allocated space
```

For XIV systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

Storage Pool Assigned Volume Space (GiB)
The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Pool Unassigned Volume Space (GiB)
The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Pool Total Volume Capacity (GiB)
The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Pool Real Configured Space (GiB)
The amount of storage space that is in an XIV system pool that has an associated soft size. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage Pool Used Volume Space (GiB)
The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when volumes are created. For these storage systems, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Pool Real Available Space (GiB)
The amount of unused space that is in an XIV system pool that has an associated soft size. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage Pool Repository Capacity (GiB)
The amount of space on all extents in the repository of a pool. This space can be used to allocate track space-efficient volumes. This attribute applies only to DS8000 storage systems.

Storage Pool Used Repository Space (GiB)
The amount of space on all extents in the repository of a pool that are allocated. This space can be used to allocate track space-efficient volumes. This attribute applies only to the DS8000 storage systems.

Storage Pool Physical Allocation Percentage
The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

`(allocated space ÷ pool capacity) × 100`

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Pool Virtual Allocation Percentage
The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

Tivoli Storage Productivity Center uses the following formula to determine the allocation percentage:

`(total volume capacity ÷ pool capacity) × 100`

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pool is 30 GiB.

This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for

the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, you can use the value for Storage Pool Shortfall Percentage to estimate how critical the shortage of space is for a pool.

This value is only available for pools with thin-provisioned volumes.

Storage Pool Shortfall Percentage
The percentage of the remaining unallocated volume space in a pool that is not available to be allocated.

The higher the percentage, the more critical the shortfall of pool space.

Tivoli Storage Productivity Center uses the following formula to determine this percentage:

`(unallocatable space ÷ (virtual capacity − allocated space)) × 100`

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

Storage Pool Allocated Space (GiB)
The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

Storage Pool Extent Size (MiB)
The size of the extent that was specified when a pool was created. Smaller

extent sizes limit the maximum size of the volumes that can be created in a pool. Smaller extent sizes minimize the amount of potentially wasted space per volume.

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage Pool Tier Capacity (GiB)
The information that this property shows depends on whether the report includes the Storage Pool Tier Type property:

- Report includes Storage Pool Tier Type: Shows the amount of space on solid-state drives and hard disk drives in a pool. The report shows the storage space on solid-state drives and on hard disk drives on different rows.
- Report does not include Storage Pool Tier Type: Shows the total storage space on all of the solid-state drives and hard disk drives in a pool.

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage Pool Tier Available Space (GiB)
The information that this property shows depends on whether the report includes the Storage Pool Tier Type property:

- Report includes Storage Pool Tier Type: Shows the amount of unused space on solid-state drives and hard disk drives in a pool. The report shows the unused space on solid-state drives and on hard disk drives on different rows.
- Report does not include Storage Pool Tier Type: Shows the total unused space on all of the solid-state drives and hard disk drives in a pool.

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage Pool Tier Number of MDisks
The information that this property shows depends on whether the report includes the Storage Pool Tier Type property:

- Report includes Storage Pool Tier Type: Shows the number of managed disks on the tier that are solid-state drives or hard disk drives. The report shows the number of managed disks that are solid-state drives and the number that are hard disk drives on different rows.
- Report does not include Storage Pool Tier Type: Shows the total number of managed disks on all of the solid-state drives and hard disk drives in a tier.

This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage Pool Unallocated Volume Space (GiB)
The amount of the Total Volume Capacity in the pool that is not allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
total volume capacity - allocated space
```

The space that is allocated for thin-provisioned volumes is typically less than their virtual capacity. Therefore, the unallocated space represents the difference between the virtual capacity and the allocated space for all the volumes in the pool.

Storage Pool Unallocatable Volume Space (GiB)
The amount of space by which the Total Volume Capacity exceeds the physical capacity of a pool.

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool. If you create volumes with more virtual capacity than can be physically allocated in the pool, you can over commit storage in the pool.

This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.

Storage Pool Number of MDisks
The number of managed disks in a storage pool.

Storage Pool Compressed Virtual Capacity (GiB)
The total virtual capacity of all the volumes that are compressed in a pool.

Storage Pool Compressed Capacity (GiB)
The amount of storage space that is used by compressed volumes in a pool. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 20.

Storage Pool Uncompressed Capacity (GiB)
The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Storage Pool Maximum I/O Capability
The projected maximum number of I/O operations per second for a pool. This value is calculated based on the value in the Storage Pool Back-End Storage Disks property, and on the values in following properties:

- Storage Pool Back-End Storage System Type
- Storage Pool Back-End Storage RAID Level
- Storage Pool Back-End Storage Disk Type

This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Storage Pool Back-End Storage Disks
The number of physical disks that contribute to the volumes on the back-end storage system. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Pool RAID Level
The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Pool Is Solid State
Shows whether there are solid-state drives in the pool. This property can contain the following values:

Mixed The pool contains both hard disk drives and solid-state drives.

Non solid state
The pool contains no solid-state drives.

Solid state
The pool contains a least one solid-state drive.

Storage Pool Is Encrypted
Shows whether the resource is encrypted. If this value is `Yes`, the resource is encrypted.

Storage Pool Format
Shows the format of the volume or pool. The format can be a Count Key Data (CKD) format, or a fixed block format.

Storage Pool Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Pool Logical Subsystem
The logical subsystem (LSS) to which a volume or pool belongs.

Storage Pool Rank Group
The rank group to which a pool is assigned, that is, 0 or 1. This property applies only to the DS8000 and DS6000 storage systems.

Storage Pool Lock Behavior
Shows whether the pool is locked for write operations, or is disabled for both read and write operations when storage space is limited. This property applies only to an XIV system.

Storage Pool Tier
The tier that a managed disk is assigned to by the internal cluster automatic detection for internal disks or by the user for external disks. This classification is used by the storage system to determine when to activate the Easy Tier function for certain pools.

Storage Pool Custom Tag 1, 2, and 3
User-defined text that is associated with a storage pool. You can add or edit the custom tags for a storage pool on the **Properties** notebook of the pool in the Tivoli Storage Productivity Center web-based GUI.

Storage Pool Is Compression Active
Shows whether the compression feature is enabled on the storage pool. If this value is `Yes`, the compression feature is enabled.

Storage Pool Back-End Storage System Type
The type of storage system that provides storage space to a pool. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Storage Pool Back-End Storage Disk Type
The class and speed of the physical disks that contribute to the volumes on the back-end storage system. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Storage Pool Back-End Storage RAID Level
The RAID level of the volumes on the back-end storage system that provide storage space to a pool. This property applies only to SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified pools.

Status information

You can create capacity and relationship reports that include the following information:

Storage Pool Status
The condition of the resource, for example normal, warning, or error.

Storage Pool Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Storage Pool Tier Type
The type of tier that is used on a storage pool. For example, the tier type can be hard disk drive (HDD) or solid-state drive (SSD).

This value is only available for Storwize V7000 storage systems that have internal storage.

Data for storage system volumes in capacity and relationship reports

You can include general information, capacity data, properties, information about copies of volumes, and other information about storage system volumes in capacity and relationship reports.

Information about storage system volumes

You can create capacity and relationship reports that include the following information:

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Storage Volume Assigned Host Connection
The host connection to which the storage volume is assigned. The host connection is a definition in the storage system that contains the WWPN for the server. The storage system uses the WWPN to assign volumes to servers.

Storage Volume Host Connection OS Type
The operating system type of the server or hypervisor that the volume is assigned to.

Storage System Volume Group ID
The unique identifier for a collection of volumes on DS8000 and DS6000 storage systems. The storage system generates this identifier.

Storage System Volume Group Name
The user-defined name for a collection of volumes on DS8000 and DS6000 storage systems.

Storage System Volume Preferred Node Name
The name of the node that is used for I/O operations for the storage system volume.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Volume Capacity (GiB)
The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

For XIV systems, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage Volume Overhead (GiB)
The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Volume Allocated Space (GiB)
The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following storage systems:

- Storage systems other than Storwize V7000 systems that are configured as back-end storage
- Storwize V7000 storage systems that are configured as back-end storage, and that are not thin-provisioned

Storage Volume Used Space (GiB)
The amount of allocated space that is used by a volume.

For SAN Volume Controller and Storwize V7000, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that Tivoli Storage Productivity Center collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Volume Used Allocated Space Percentage
The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
(volume used space ÷ volume allocated space) × 100
```

Storage Volume Unallocated Space (GiB)
The amount of space in a pool that is not reserved for a volume.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
capacity - allocated space
```

Available only for thin provisioned volumes.

Storage Volume Physical Allocation Percentage
The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
(allocated space ÷ capacity) × 100
```

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Volume Unused Space (GiB)
The amount of space that is allocated to a volume and is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
Allocated Space - Used Space
```

This property is available only for SAN Volume Controller and Storwize V7000 pools.

Storage Volume Shortfall Percentage
The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume.

The higher the percentage, the more critical the shortfall of space.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
(unallocatable volume Space ÷ volume unallocated space) × 100
```

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This Storage Volume Shortfall Percentage represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

```
100 GiB unallocated volume space - 70 GiB remaining
pool space ÷ 100 GiB unallocated volume space × 100
```

Storage Volume Unallocatable Space (GiB)
The unallocatable storage space in GiB of a thin-provisioned volume. Unallocatable space cannot be supplied by the pool to which the volume belongs.

Storage Volume Uncompressed Capacity (GiB)
The amount of storage space that is used if the compressed volume space

is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Volume Type
The type of storage volume. For example, the storage volume can be striped or sequential. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

Storage Volume WWN
The worldwide name of the volume.

Storage Volume Fast Write State
Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted.

Storage Volume Grain Size (KiB)
The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

Storage Volume Is Assigned
Shows whether the volume is assigned to a server. If this value is `Yes`, the volume is assigned to a server.

Storage Volume Is Auto Expand Enabled
Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is `Yes`, the volume or volume copy can allocate new extents from a pool automatically.

Storage Volume Is Encryptable
Shows whether the resource can be encrypted. If this value is `Yes`, the resource can be encrypted.

Storage Volume Is Encrypted
Shows whether the resource is encrypted. If this value is `Yes`, the resource is encrypted.

Storage Volume Is Formatted
Shows whether a storage volume is formatted. If this value is `Yes`, the volume is formatted. This value is only available for SAN Volume Controller and Storwize V7000 storage systems that are configured as back-end storage.

Storage Volume Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Volume Metro Mirror Name
The name of the Metro Mirror that keeps the synchronous copy of the storage virtualizer volume.

Storage Volume Mirror Count
The number of mirrors that keep a synchronous copy of the resource.

Storage Volume Throttle
The maximum number of commands that the storage volume can queue.

Storage Volume Warning Level
The percentage of used capacity of the resource at which a warning is generated.

Storage Volume FlashCopy Relationship
Shows whether a volume on a storage system is in a FlashCopy relationship. This property can contain the following values:

Source Relationship
The volume is the source of the relationship.

Target Relationship
The volume is the target of the relationship.

Both Source and Target Relationship
The volume is both the source and target of the relationship.

Not in a FlashCopy Relationship
The volume is not in a FlashCopy relationship.

Storage Volume Remote Copy Relationship
Shows whether a volume on a storage system is in a remote copy relationship. This property can contain the following values:

Not in a Copy Relationship
The volume is not in a remote copy relationship.

Source Relationship
The volume is the source of the relationship.

Target Relationship
The volume is the target of the relationship.

Both Source and Target Relationship
The volume is both the source and target of the relationship.

Storage Volume RAID Level
The RAID level of a storage volume, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Volume Is Compressed
Shows whether the volume is compressed. If this value is `Yes`, the volume is compressed.

Storage Volume Is Volume Control Manager Database
Shows whether a volume is used as the Volume Control Manager Database. If this value is `Yes`, the volume is used as the Volume Control Manager Database.

Storage Volume Format
Shows the format of the volume or pool. The format can be a Count Key Data (CKD) format, or a fixed block format.

Storage Volume Serial Number
The serial number of the resource.

Storage Volume Number
The number of the volume that is assigned by the system.

Storage Volume Logical Subsystem
The logical subsystem (LSS) to which a volume or pool belongs.

Status information

You can create capacity and relationship reports that include the following information:

Storage Volume Status
The condition of the resource, for example normal, warning, or error.

Storage Volume Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Volume copy information

You can create capacity and relationship reports that include the following information:

Storage Volume Copy ID
The identifier of a volume copy.

Storage Volume Copy Grain Size (KiB)
The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

Storage Volume Copy Is Auto Expand Enabled
Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is `Yes`, the volume or volume copy can allocate new extents from a pool automatically.

Storage Volume Copy Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Storage Volume Copy Is Primary
Shows whether the volume copy is the primary volume copy.

Storage Volume Copy Is Thin Provisioned
Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is `Yes`, the resource is thin-provisioned.

Storage Volume Copy Total Number of Extents
The number of extents that are used by the storage volume copy.

Storage Volume Copy Type
The type of copy of the storage volume. For example, the copy of the storage volume can be sequential, a mirror copy, or another type of copy.

Storage Volume Copy Warning Level
The percentage of used capacity of the resource at which a warning is generated.

Storage Volume Copy Used Allocated Space Percentage
The percentage of reserved space for a volume copy that is being used. This value is always less than or equal to 100% because a volume copy cannot use more space than is allocated.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
(used space ÷ allocated space) × 100
```

This property is available only for SAN Volume Controller and Storwize V7000 pools.

Storage Volume Copy Used Space (GiB)
The amount of allocated space that is used by a volume copy.

Storage Volume Copy Allocated Space (GiB)
The amount of space that is reserved for a volume copy.

The space that is allocated for a thin-provisioned volume copy is less than its virtual capacity, which is shown in the Storage Virtualizer Capacity (GiB) property.

For volume copies that are not thin provisioned, this value is equal to the value of the Storage Volume Copy Used Space (GiB) property.

Storage Volume Copy Unallocated Space (GiB)
The amount of space in a pool that is not reserved for a volume copy.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
capacity - allocated space
```

Storage Volume Copy Physical Allocation Percentage
The percentage of physical space that is reserved for a volume copy. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
(allocated space ÷ capacity) × 100
```

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volume copies is 50 GiB.

Storage Volume Copy Unused Space (GiB)
The amount of space that is allocated to a volume copy and that is not yet used.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
allocated space - used space
```

This property is available only for SAN Volume Controller and Storwize V7000 pools.

Storage Volume Copy Shortfall Percentage
The percentage of the remaining unallocated volume copy space in a pool that is not available to be allocated to a volume copy.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

The higher the percentage value is, the more critical the shortfall of space is.

Tivoli Storage Productivity Center uses the following formula to determine this value:

```
(unallocatable space ÷ (capacity – real capacity)) × 100
```

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level for the volume. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space for a volume.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume uses 50 GiB, 100 GiB is, nevertheless, committed to that volume (150 GiB minus 50 GiB). There is a shortfall of 30 GiB when you use the following calculation:

```
(100 GiB remaining commitment of volume space to
the volume - 70 GiB remaining pool space)
```

The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

```
(100 GiB unallocatable volume copy space - 70 GiB
remaining pool space) ÷ 100 GiB unallocated volume space × 100
```

This property is available only for volume copies that are thin provisioned.

Storage Volume Copy Unallocatable Space (GiB)
The unallocatable storage space in GiB of a thin-provisioned volume copy. Unallocatable space cannot be supplied by the pool to which the volume copy belongs.

Storage Volume Copy Is Compressed
Shows whether the volume copy is compressed. If this value is `Yes`, the volume copy is compressed.

Storage Volume Copy Uncompressed Capacity (GiB)
The amount of storage space that is used if the compressed space on the volume copy is uncompressed. For example, if 100 GiB of data is compressed, the value is 100.

Data for storage system disks in capacity and relationship reports

You can include properties, and other data about storage system disks in capacity and relationship reports.

Information about storage system disks

You can create capacity and relationship reports that include the following information:

Storage Disk Name
The name that was assigned to the storage disk when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Disk Capacity (GiB)
The amount of storage space that is on a physical disk on the storage system.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Disk Flags
Shows whether a device is a spare device. This property applies only to DS8000 storage systems.

Storage Disk Is Encryptable
Shows whether the resource can be encrypted. If this value is `Yes`, the resource can be encrypted.

Storage Disk Is Encrypted
Shows whether the resource is encrypted. If this value is `Yes`, the resource is encrypted.

Storage Disk Is Solid State
Shows whether the disk is a solid-state drive.

Storage Disk Speed (RPM)
The speed of the disk.

Storage Disk Tag
A number that identifies the array to which a disk belongs. This property applies only to DS8000, DS6000, and TotalStorage Enterprise Storage Server storage systems.

Storage Disk Class
The storage technology of the disk. This property applies only to DS8000, DS5000, and DS4000 storage systems. This property can contain the following values:

FC The disk is a Fibre Channel disk.

Nearline ATA
The disk is a nearline Advanced Technology Attachment disk.

Nearline SAS
The disk is a nearline serial-attached SCSI disk.

SAS The disk is a serial-attached SCSI disk.

SATA The disk is a Serial Advanced Technology Attachment disk.

SSD The disk is a solid-state drive.

Unavailable
The disk is on a DS5000 or DS4000 series system, and is not a solid-state drive.

Unsupported
The disk is on a system that is not a DS8000, DS5000, or DS4000 series system.

Status information

You can create capacity and relationship reports that include the following information:

Storage Disk Status
The condition of the resource, for example normal, warning, or error.

Storage Disk Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is `Yes`, the resource was detected.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Storage Disk Firmware Version
The version number of the firmware that is running on the disk.

Storage Disk Model
The model name or model number of the resource.

Storage Disk Serial Number
The serial number of the resource.

Storage Disk Vendor
The vendor who supplied the resource.

Data for storage system groups in capacity and relationship reports

You can include general information about, and properties of, storage system groups in capacity and relationship reports.

Information about storage system groups

You can create capacity and relationship reports that include the following information:

Storage System Group Name
The name that was assigned to the storage system group when it was added to the system.

Component properties

You can create capacity and relationship reports that include the following information:

Storage System Group Creator
The user name of the user who created the group.

Storage System Group Last Modified
The date and time that the resource was last modified.

Storage System Group Last Modified by User
The user name of the user who last modified the resource.

Data for managed disks on storage systems in capacity and relationship reports

You can include general information, capacity data, properties, and other information about managed disks on storage systems in capacity and relationship reports.

Information about managed disks on storage systems

You can create capacity and relationship reports that include the following information:

Storage System MDisk Name
The name that was assigned to the managed disk on a storage system when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage System MDisk Capacity (GiB)
The amount of storage space on the managed disk.

Storage System MDisk Available Space (GiB)
The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.

Storage System MDisk Strip Size (KB)
The RAID strip size on a managed disk on a storage system.

Component properties

You can create capacity and relationship reports that include the following information:

Storage MDisk RAID Level
The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

Storage MDisk Spare Goal
The number of spare drives that are required to maintain redundancy. Use spare drives to protect the system against drive failures in the array on a managed disk on a storage virtualizer.

Storage System MDisk Type
The type of managed disk on a storage system. For example, the disk on a storage system can be a local managed disk.

Storage System MDisk Mode
The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

Array Extents are provided from local disks.

Managed
Extents are provided from a back-end storage volume.

Unmanaged
The managed disk is not used in the system.

Storage System MDisk Is Balanced
Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is `Yes`, the LUNs are balanced.

Storage System MDisk Fast Write State
Shows whether the cache for a volume on a disk that is managed by a storage system is empty, contains data, or is corrupted.

Storage System MDisk Write Verify
Shows whether all write operations on a managed disk on a storage system are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is `Yes`, all write operations are verified by a follow-up read operation.

Status information

You can create capacity and relationship reports that include the following information:

Storage System MDisk Status
The condition of the resource, for example normal, warning, or error.

Storage MDisk Native Status
Shows the level of access of nodes in the system to a managed disk or volume on a storage system. The level of access can be online, offline, degraded, or excluded.

Data for storage resource groups in capacity and relationship reports

You can include general information about storage resource groups in capacity and relationship reports.

Information about storage resource groups

You can create capacity and relationship reports that include the following information:

Storage Resource Group Name
The name that was assigned to the storage resource group when it was created.

Storage Resource Group Last Modified by User
The user name of the user who last modified the resource.

Storage Resource Group Custom Tag 1, 2, and 3
User-defined text that is associated with a storage resource group.

Troubleshooting reports

Find answers to questions about resolving issues with reports in Tivoli Common Reporting.

Problem
When you run reports, Tivoli Common Reporting times out.

Solution

Based on your operating system, complete the following steps to change the timeout:

Windows

1. Open a command-line interface.
2. Stop Tivoli Common Reporting:

   ```
   c:\program files\IBM\JazzSM\profiles\bin\stopServer.bat
   server_name -username JazzSM_username -password JazzSM_password
   ```

3. Open the `configuration` directory:

   ```
   c:\program files\IBM\JazzSM\reporting\cognos\configuration
   ```

4. Copy the `cogstartup.xml` file and, then change the name of the copied file to `cogstartup.orig`.
5. Open the `cogstartup.xml` file:

   ```
   c:\program files\IBM\JazzSM\reporting\cognos\
   configuration\cogstartup.xml
   ```

6. In the `cogstartup.xml` file, enter the new timeout value in seconds for the *inactivityTimeout* setting.
7. In the `cogstartup.xml` file, enter new timeout value in minutes for the *mdmSessionTimeout* setting.
8. Save the `cogstartup.xml` file.
9. Start Tivoli Common Reporting:

   ```
   c:\program files\IBM\JazzSM\profiles\bin\startServer.bat
   server_name
   ```

AIX or Linux

1. Open a command-line interface.
2. Stop Tivoli Common Reporting:

   ```
   /opt/IBM/JazzSM/profile/bin/stopServer.sh server_name
    -username JazzSM_username -password JazzSM_password
   ```

3. Open the `configuration` directory:

   ```
   /opt/IBM/JazzSM/reporting/cognos/configuration
   ```

4. Copy the `cogstartup.xml` file and, then change the name of the copied file to `cogstartup.orig`.
5. Open the `cogstartup.xml` file:

   ```
   /opt/IBM/JazzSM/reporting/cognos/configuration/cogstartup.xml
   ```

6. In the `cogstartup.xml` file, enter the new timeout value in seconds for the *inactivityTimeout* setting.
7. In the `cogstartup.xml` file, enter new timeout value in minutes for the *mdmSessionTimeout* setting.
8. Save the `cogstartup.xml` file.
9. Start Tivoli Common Reporting:

   ```
   /opt/IBM/JazzSM/profile/bin/startServer.sh server_name
   ```

Problem

You can set output format options for reports either when you schedule a report or run a report. If you choose PDF format for your report and the report contains many columns, all of the columns might not fit on one page. For example, if a report contains 20 columns, then the row for a specific resource might be shown on more than one page: 10 columns on the first page and 10 columns on the second page.

Solution

To fit the maximum number of columns on one page, set the PDF options in one of the following ways:

When you schedule a report

1. Navigate to a report in Tivoli Common Reporting.
2. Click the **Schedule** icon .
3. Select **Override the default values**.
4. Select **PDF**, and then click **Set**.
5. From the **Orientation** list, select **Landscape**.
6. From the **Paper size** list, select **11 × 17**.

When you run a report

1. Navigate to a report in Tivoli Common Reporting.
2. Click the **Run with options** icon .
3. Click the **advanced options** link in the tip.
4. Click **Run in the background**.
5. Select **PDF**, and then click **Set**.
6. From the **Orientation** list, select **Landscape**.
7. From the **Paper size** list, select **11 × 17**.

Problem

When you create a custom report, you cannot find online help that describes the columns that you can include in your reports.

Solution

To view the online help, follow these steps:

1. Go to the dashboard of the Tivoli Storage Productivity Center web-based GUI.
2. Select **Reporting** > **View predefined reports**.
3. Click any predefined report.
4. Click the **Help** link at the upper right of the page.
5. Go to the custom reports help content.

To view the custom reports help content in the information center, go to the Tivoli Storage Productivity Center Information Center at publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp. Search for "Custom reports".

Problem

If the Tivoli Storage Productivity Center database (TPCDB) is not accessible, you are prompted to reenter your user ID and password although you previously entered the correct user ID and password.

Solution

Check that the server that hosts TPCDB is available. Check that the firewall settings on the server are not preventing you from accessing TPCDB.

Problem

If the password to access the Tivoli Storage Productivity Center database (TPCDB) is changed, you must also change the password in Tivoli Common Reporting.

Solution

In Tivoli Common Reporting, complete the following steps:

1. In the **Public Folders** tab, click **Launch** > **Administration**.
2. Click the **Configuration** tab.
3. Click **TPCDB**, and then click **TPCDB** again.
4. In the Actions column, click the **Set properties** icon, and then click the **Signon** tab.
5. Click **Edit the signon**.
6. Make your changes, and then click **OK**.

Problem

When you use particular combinations of properties from the Storage Volume section in the Performance package in a custom report, a `No data available` message might be displayed. This occurs when you use the following combinations of properties in a custom report:

- Server Name, Hypervisor Name, and any other property from the Storage Volume section
- Server Name, Hypervisor Name, Hypervisor Cluster Name, and any other property from the Storage Volume section
- Server Name, Hypervisor Cluster Name, and any other property from the Storage Volume section

Solution

Do not use these combinations of properties in custom reports.

Problem

When you try to log in to Tivoli Common Reporting, the following error message is displayed:

```
DPR-ERR-2109 The dispatcher cannot service the request at this time. The
dispatcher is still initializing. Contact your administrator if this
problem persists.
```

The Jazz for Service Management server starts Tivoli Common Reporting. This problem occurs if Tivoli Common Reporting starts when DB2 is not running.

Solution

Stop the Jazz for Service Management server, restart DB2, and restart the Jazz for Service Management server. To do this, complete the following steps:

1. Open a command-line interface.
2. Stop the Jazz for Service Management server:

   Windows

   JazzSM_install_directory\profile\bin\stopServer.bat *server_name* -username *JazzSM_username* -password *JazzSM_user_password*

   AIX or Linux

   JazzSM_install_directory/profile/bin/stopServer.sh *server_name* -username *JazzSM_username* -password *JazzSM_user_password*
3. Restart DB2.
4. Start the Jazz for Service Management server:

   Windows

   JazzSM_install_directory\profile\bin\startServer.bat *server_name*

   AIX or Linux

   JazzSM_install_directory/profile/bin/startServer.sh *server_name*

Problem

In Internet Explorer 8, 9, or 10 when you open a report in Report Studio for the first time, a CM-REQ-4158 error might be displayed.

Solution

For information about how to resolve this problem, see the IBM Tivoli Systems Management Information Center at publib.boulder.ibm.com/infocenter/tivihelp/v3r1/index.jsp. Search for "CM-REQ-4158 Jazz".

Problem

When you add a filter to a property in Query Studio, not all of the resources are listed in the **Filter (Pick values from a list)** area. For example, when you add a filter to the Storage Volume Name property, not all of the volumes are listed.

This problem occurs because the maximum number of items in a list in Query Studio is 5000. If there are more than 5000 items in the list, only the first 5000 values are displayed, in alphabetical order.

Solution

Click **Search for values** in the **Filter (Pick values from a list)** area. Type a keyword, and then click **Search** to search for volumes on a resource. For example, type `a%` to find volume names that begin with 'a' or 'A'.

Problem

When you run a report, it takes a long time for the report output to display. It might take a long time to run a report for some combinations of the interval and reporting period. For example, the report might take a long time if you select the Sample interval or Hourly interval with one of the following periods:

- Last 30 days
- Last 90 days
- All

Solution

When you run a report, use shorter reporting periods with the Sample interval or Hourly interval.

Problem

When you run a report, the output that is displayed contains no data. If you select the Daily interval and a recent reporting period, then no data might be available. For example, if you select Last 24 hours or Today, the data for those reporting periods is unlikely to be collected yet.

Solution

When you run a report with a Daily interval, use a reporting period other than Last 24 hours or Today.

Problem

If you log off from a remote connection, you must manually restart Jazz for Service Management and Tivoli Common Reporting.

Solution

To restart Jazz for Service Management and Tivoli Common Reporting, complete the following steps for your operating system:

Windows

JazzSM_install_directory\profile\bin\startServer.bat *server_name*

AIX or Linux

JazzSM_install_directory/profile/bin/startServer.sh *server_name*

Problem

You set the default language of your operating system or your browser to Portuguese (Brazilian). However, you see English text in the report prompt pages or in the report output. This problem is a known problem with Cognos 10.2 and Tivoli Common Reporting Version 3.1.0.1.

Solution

To set the language of the report prompt pages and report output to Portuguese (Brazilian), complete the following steps:

1. In Tivoli Common Reporting, click the **My Area Options** icon .
2. Click **My Preferences**.
3. In the **Regional options** area, under **Content language**, select **Use the following language**.
4. Select **Portuguese (Brazil)**.
5. Click **OK**.

Problem

If you create a custom performance report in Query Studio, in particular circumstances, the report takes a long time to complete. This problem might occur in the following circumstances:

- The report is for more than 10,000 volumes.
- The interval for the report is Hourly, and the data for the report was collected for more than 15 days.
- The interval for the report is Sample, and the data for the report was collected for more than 7 days.

Solution

When this problem occurs, complete the following steps for the report in Query Studio:

1. Click **Run Report** in the menu pane.
2. Click **Advanced Options**.
3. Clear each of the following check boxes:
   - **Automatically generate footer summaries for measures**
   - **Automatically summarize detail values, suppressing duplicates**
   - **Enable drill up and drill down in the report output**
   - **Enable drill through from a package in the report output**
4. Click **OK**.

Showing the package version number and build ID

Access information about the packages for custom reports and predefined reports, such as the version number and the build ID.

You can identify the version number and build ID for each of the following packages:

- Capacity and Relationships
- Performance

1. Start the web-based GUI for Tivoli Storage Productivity Center.
2. Depending on the package whose version number and build ID you want to see, complete one of the following steps:

- To identify the version number and build ID of the Capacity and Relationships package, select **Reporting** > **Create capacity reports**.
- To identify the version number and build ID of the Performance package, select **Reporting** > **Create performance reports**.

3. In the navigation tree, position the cursor over the item under the name of the package.

The version number and the build ID for the package is shown.

Reporting in the stand-alone GUI

Use Tivoli Storage Productivity Center reports to view detailed information about the storage resources in your environment. These reports are organized into different types and categories and provide both summary and detailed information depending on your needs.

Data in reports is based on statistical information accumulated by data collection schedules and stored in the database repository. You can select the specific systems, storage subsystems, switches, fabrics, NAS filers, clusters, disks, file systems, and even users upon which you want to report.

Reports are organized into report types and categories that you can expand to reveal additional subcategories of reporting. Many reports are also accessible through the topology viewer.

- **Reporting types:** system reports, batch reports, user ID reports, user-defined online reports
- **Report categories:** asset, availability, capacity, usage, usage violation, backup, storage subsystem performance, switch performance

Consider the following steps when determining how to best use Tivoli Storage Productivity Center reports:

1. Determine how you want to generate and view reports. See "Choosing a reporting type" for more information.
2. Identify which reports contain the most important information about your storage environment. See "Choosing a report category" on page 1297 for more information.

Choosing a reporting type

Use pre-defined (system reports), saved reports, batch reports, rollup reports, and user-defined reports, tables, and charts to view the information collected by monitoring and quota jobs. Tivoli Storage Productivity Center includes these different reporting types to provide you with the flexibility to view data about your storage resources according to the needs of your environment.

The following table describes the reporting types available in Tivoli Storage Productivity Center. Use this table to help determine how to best deliver and view reports about your storage resources.

Reporting type	Description	Location in the navigation tree
User-defined online reports	These reports can be generated at any time during a product session and are only available for that session. Each manager within Tivoli Storage Productivity Center has a set of reports specific to that manager. For example, to view detailed reports about storage subsystems, access the reporting function within Disk Manager. See the following topics for more information about online reports: • "Working with user-defined online reports" on page 1320 • "Data Manager reports" on page 1331 • "Using Data Manager for Databases reports" on page 1462 • "Disk Manager reports" on page 1490 • "Fabric Manager reports" on page 1502	Data Manager > Reporting Data Manager for Databases > Reporting Disk Manager > Reporting Fabric Manager > Reporting
System reports	System reports are automatically generated by Tivoli Storage Productivity Center and contain some of the most common information about your monitored storage resources. These pre-defined reports enable you to view your storage enterprise data quickly and efficiently with one click. Data for these system reports are gathered every time that monitoring jobs are run against your storage resources. See the following topics for more information about online reports: • "System reports - Data" on page 1510 • "System reports - Data Manager for Databases" on page 1525 • "System Reports - Disk Manager" on page 1537 • "System reports - Fabric Manager" on page 1554	IBM Tivoli Storage Productivity Center > Reporting > System Reports
User ID (saved) reports	Create saved reports when you want specific reports generated for you whenever you log in to a product session. This enables you to save reports that you can generate at any time without having to define your report criteria again. Saved reports are only visible to the user ID that created them. See "User ID (saved) reports" on page 1574 for more information.	IBM Tivoli Storage Productivity Center > Reporting > My Reports > *user ID* Reports where *user ID* represents the IBM Tivoli Storage Productivity Center ID of the user that saved the report

Reporting type	Description	Location in the navigation tree
Batch reports	A batch report represents any Tivoli Storage Productivity Center report that you want to run on a regular basis and save its data to a file. You can view the information in a generated batch report file directly or import it into an application of your choice. To create a batch report, you must determine which Tivoli Storage Productivity Center report you want to generate, the schedule on which to generate the report, and the file to which to save its data. You can save the data generated for a batch report to any of the following file formats: CSV, HTML, formatted text, PDF chart, and HTML chart. See "Batch reports" on page 1562 for more information.	IBM Tivoli Storage Productivity Center > Reporting > My Reports > Batch Reports
Rollup Reports	Use rollup reports to view the storage information collected by the master Tivoli Storage Productivity Center server from its associated subordinate servers. Rollup reports enable you to have a network-wide perspective of storage usage in an environment where multiple Tivoli Storage Productivity Center servers are deployed and collecting storage metrics. See the following topics for more information about online reports: • "Rollup reports" on page 1575 • "Tivoli Storage Productivity Center servers" on page 246 • "Collect storage statistics from Tivoli Storage Productivity Center servers" on page 331	IBM Tivoli Storage Productivity Center > Reporting > Rollup Reports
Data Source reports	Use these reports to view detailed information about the agents from which Tivoli Storage Productivity Center collects information about storage resources. Data sources include CIMOM agents, Storage Resource agents, Fabric agents, VMWare, and subordinate Tivoli Storage Productivity Center servers. See Data Source reports for more information.	IBM Tivoli Storage Productivity Center > Reporting > Data Source Reports

Choosing a report category

IBM Tivoli Storage Productivity Center provides a number of categories for reports that help you identify which reports contain the information you want to view about your storage environment. Each category of report contains information gathered by different data collection and quota jobs.

The following table describes the reporting categories. Use this table to help determine which reports are most useful to you and which data collection jobs to run for those reports.

Report Category	Description	What data collection job should I run for this report?
Asset reports	View itemized statistics about your storage resources in a hierarchical view that enables you to drill down to view those resources in greater detail. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > Asset • Data Manager for Databases > Reporting > Asset • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) • Data Manager for Databases > My Reports **Example:** View information about a storage subsystem and drill down into its related storage pools, disks, volumes, and disk groups. See "Asset reports" on page 1334 for more information about Asset reports.	Probes, Discovery Data for System-wide Users and Monitored Directory reports are gathered by Scans.
Availability reports	View availability statistics about your storage resources, both from a network point of view and from a computer uptime perspective. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > Availability • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) **Example:** Use a Computer Uptime report to determine the percentage of time that a computer was up, the number of times it was rebooted, and the total time it was down within a user-defined period of time. See "Availability Reports" on page 1358 for more information about Availability reports.	Ping data: Pings Computer Uptime: Probes

Report Category	Description	What data collection job should I run for this report?
TPC-wide Storage Space	Examine storage capacities of storage resources at a file system, computer, or network level. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > TPC-wide Storage Space • Data Manager for Databases> Reporting > Capacity • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) • Data Manager for Databases > My Reports **Example:** Identify how much storage space is being used on the file systems within your environment, and how much unused storage space is available. See "TPC-wide Storage Space Reports" on page 1360 for more information about Capacity reports.	Probes
Usage reports	View information about the usage and growth of consumption of your storage resources, as well as select files against which to run an Tivoli Storage Productivity Center archive/backup job. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > Usage • Data Manager for Databases > Reporting > Usage • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) • Data Manager for Databases > My Reports (if you save the definition of an online report it appears in this location) **Example:** Determine how much of the storage being used in your environment can be considered wasted space (orphaned, obsolete, or forbidden files). See "Usage Reports" on page 1420 for more information about Usage reports.	Scans

Report Category	Description	What data collection job should I run for this report?
Usage Violation reports	View quota and constraint violations as well as select files for archive/backup jobs. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > Usage Violations • Data Manager for Databases > Reporting > Usage Violations • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) • Data Manager for Databases > My Reports (if you save the definition of an online report it appears in this location) **Example:** Identify the users or groups of users that are using storage space on machines where they should not. See "Usage Violation reports" on page 1427 for more information about Usage Violation reports.	Scans, Quotas, Constraints
Backup reports	Investigate the sizes of data that would be required to backup to achieve full backups, determine the amount of space required to perform an incremental backup of any or all of your data at any given point in time, and select files for archive/backup jobs. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > Backup • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) • Data Manager for Databases > My Reports (if you save the definition of an online report it appears in this location) **Example:** Determine if you have enough storage space to perform a backup of your data (incremental or full). See "Backup reports" on page 1429 for more information about Usage Violation reports.	Scans

Report Category	Description	What data collection job should I run for this report?
Monitored Computer Storage Space	Use these host-based reports to view the data that is collected by probes for computers, Network Attached Storage, and hypervisors. These reports contain detailed information about the consumption of file system space and enable you to relate that space back to external or internal storage without counting shared space multiple times. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > Monitored Computer Storage Space • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) See "Monitored Computer Storage Space reports" on page 1431 for more information about Monitored Computer Storage Space reports.	Probes
Storage subsystem reports	View storage capacity for a computer, file system, storage subsystem, volumes, array sites, ranks, and disks. These reports also enable you to view the relationships among the components of a storage subsystem. **Location:** Disk Manager > Reporting > Storage Subsystems **Example:** Determine if you have enough storage space to perform a backup of your data (incremental or full). See "Disk Manager reports" on page 1490 for more information about Storage Subsystem reports.	Probes
Storage Subsystem Performance reports	View performance data for storage subsystems. **Location:** Disk Manager > Reporting > Storage Subsystem Performance **Example:** Check the performance rate for storage subsystem activities, such as read I/O rate, read data rate, total port I/O rate, etc. See "Disk Manager reports" on page 1490 for more information about Storage Subsystem Performance Monitor reports.	Subsystem Performance Monitors

Report Category	Description	What data collection job should I run for this report?
Switch Performance reports	View performance data for fabric switches. **Location:** Fabric Manager > Reporting > Switch Performance **Example:** Check the performance rate for switch activities, such as port send packet rate, port send data rate, total port data rate, etc. See "Fabric Manager reports" on page 1502 for more information about Switch Performance Monitor reports.	Switch Performance Monitors
Rollup reports: Asset, Capacity, Database Asset, Database Capacity	View the storage information that has been collected by IBM Tivoli Storage Productivity Center servers through a single interface. These rollup reports enable you to have a consolidated perspective of storage usage in an environment where multiple IBM Tivoli Storage Productivity Center servers are deployed and collecting storage metrics. **Location:** IBM Tivoli Storage Productivity Center > Reporting > Rollup Reports **Example:** Check the disk capacity of all the computers monitored by different servers in your environment. See "Rollup reports" on page 1575 for more information about rollup reports.	IBM Tivoli Storage Productivity Center Server Probes
Data Source reports	Use data source reports to view detailed information about the agents from which Tivoli Storage Productivity Center collects information about storage resources. Data sources include CIMOM agents, Storage Resource agents, Fabric agents, VMWare, and subordinate Tivoli Storage Productivity Center servers. **Location:** IBM Tivoli Storage Productivity Center > Reporting > Data Source Reports **Example:** View detailed information about the CIMOM Agents that have been added to Tivoli Storage Productivity Center as data sources. See "Data source reports - overview" on page 1633 for more information about data source reports.	These reports are available for data sources that have been associated with Tivoli Storage Productivity Center.

What can I find out using reports?

Learn about some of the key storage questions that you can answer by generating Tivoli Storage Productivity Center reports.

Use the following table to view a list of reports that can be generated to answer key questions about storage resources that are monitored by Tivoli Storage Productivity Center.

What do you want to know?	Which report answers this question?	Which schedule can I run to generate this report?
What hardware is currently being used within the enterprise?	Data Manager > Reporting > Asset > By <storage_resource_type>	Probes
How much storage space is being used?	Data Manager > Reporting > TPC-wide Storage Space > Consumed File System Space	Probes
Of the space that is being used, how much is wasted space? (orphaned, obsolete, or duplicate files)	Data Manager > Reporting > Usage > Files Data Manager > Reporting > Usage Violations > Constraint Violations	Scans
How much unused or unallocated storage space is available?	Data Manager > Reporting > TPC-wide Storage Space > Available File System Space	Probes
Is there enough space to do a backup of our data? (incremental or full)	Data Manager > Reporting > Backup > Backup Storage Requirements	Scans
Which files in the environment are at risk if a disaster hits?	Data Manager > Reporting > Backup > Most at Risk Files	Scans
What "forbidden files" are taking up storage space on the network?	Data Manager > Reporting > Usage Violations > Constraint Violations	Scans
Which user or group of users is consuming space on computers that they should not?	Data Manager > Reporting > Usage Violations > Quota Violations	Scans
How fast are storage resources being consumed?	Data Manager > Reporting > TPC-wide Storage Space > Consumed File System Space	Probes
Which user or business unit is consuming the most storage resources?	Data Manager > Reporting > Usage > Users Data Manager > Reporting > Usage > OS User Groups	Scans
Where is disk space reaching a critical shortage?	Data Manager > Reporting > TPC-wide Storage Space > Disk Space	Probes
Which devices have been experiencing the most use/traffic? Which devices have been experiencing the least use/traffic?	Data Manager > Reporting > Usage > Access Load	Scans
Based on historical consumption, how much storage space should be added and when should it be added?	Data Manager > Reporting > TPC-wide Storage Space > Available File System Space	Probes
What files have not been used within a certain time period?	Data Manager > Reporting > Usage > Files > Most Obsolete Files	Scans
What are the largest files or directories within the environment?	Data Manager > Reporting > Usage > Files > Largest Files	Scans

What do you want to know?	Which report answers this question?	Which schedule can I run to generate this report?
What are the most modified files within the environment?	Data Manager > Reporting > Usage > Access Load > Modification Time	Scans
Which computers are "up" or "down" within the environment?	Data Manager > Reporting > Availability > Ping	Pings
What hardware is most prone to error? (for example, disk defects)	Data Manager > Reporting > Asset > System-wide > Disks	Probes
Which file owners no longer exist in the environment?	Data Manager > Reporting > Usage > Files > Orphan files	Scans
How much storage space is an instance consuming?	• **Oracle:** Data Manager for Databases > Reporting > Capacity > Oracle > Total Instance Storage • **SQL/Server:** Data Manager for Databases > Reporting > Capacity > SQL/Server > Total Instance Storage • **Sybase:** Data Manager for Databases > Reporting > Capacity > Sybase > Total Instance Storage • **UDB:** Data Manager for Databases > Reporting > Capacity > UDB > Total Instance Storage	Probes

What do you want to know?	Which report answers this question?	Which schedule can I run to generate this report?
Which user or business unit is consuming the most storage resources and on which instances are they consuming those resources?	**Oracle:** • Data Manager for Databases > Reporting > Usage > Oracle > Users > Group by User Name, *or* • Data Manager for Databases > My Reports > System Reports > Oracle - User Database Space Usage **SQL/Server:** • Data Manager for Databases > Reporting > Usage > SQL/Server > Users > Group by User Name, *or* • Data Manager for Databases > My Reports > System Reports > SqlServer - User Database Space Usage **Sybase:** • Data Manager for Databases > Reporting > Usage > Sybase > Users > Group by User Name, *or* • Data Manager for Databases > My Reports > System Reports > Sybase - User Database Space Usage **UDB:** • Data Manager for Databases > Reporting > Usage > UDB > Users > Group By User Name, *or* • Data Manager for Databases > My Reports > System Reports > UDB - User Database Space Usage	Scans
Which tablespaces are reaching a critical shortage of free space?	• **Oracle:** Data Manager for Databases > Reporting > Capacity > Oracle > Data File Free Space > By Tablespace • **UDB:** Data Manager for Databases > Reporting > Capacity > UDB > Container Free Space > By Tablespace	Probes
Which tablespaces contain the most free space?	• **Oracle:** Data Manager for Databases > Reporting > Capacity > Oracle > Data File Used Space > By Tablespace • **UDB:** Data Manager for Databases > Reporting > Capacity > UDB > Container Used Space > By Tablespace	Probes

What do you want to know?	Which report answers this question?	Which schedule can I run to generate this report?
Which databases are reaching a critical shortage of free space?	• **SQL/Server:** Data Manager for Databases > Reporting > Capacity > SQL/Server > Data File Free Space > By Database • **Sybase:** Data Manager for Databases > Reporting > Capacity > Sybase > Fragment Free Space > By Database	Probes
Which databases contain the most used space?	• **SQL/Server:** Data Manager for Databases > Reporting > Capacity > SQL/Server > Data File Used Space > By Database • **Sybase:** Data Manager for Databases > Reporting > Capacity > Sybase > Fragment Used Space > By Database	Probes
Based on historical consumption of storage space for an RDBMS, how much storage space should be added and when should it be added?	• **Oracle:** Data Manager for Databases > Reporting > Capacity > Oracle > Data File Free Space > By Database > History Chart:Free Space % for selected • **SQL/Server:** Data Manager for Databases > Reporting > Capacity > SQL/Server > Data File Free Space > By Instance > History Chart:Free Space % for selected • **Sybase:** Data Manager for Databases > Reporting > Capacity > Sybase > Device Free Space > By Device > History Chart:Free Space % for selected • **UDB:** Data Manager for Databases > Reporting > Capacity > UDB > Container Free Space > By Instance > History Chart:Free Space % for selected	Probes
What are the largest tables within the environment?	• **Oracle:** Data Manager for Databases > Reporting > Usage > Oracle > Tables > Largest Tables • **SQL/Server:** Data Manager for Databases > Reporting > Usage > SQL/Server > Tables > Largest Tables • **Sybase:** Data Manager for Databases > Reporting > Usage > Sybase > Tables > Largest Tables • **UDB:** Data Manager for Databases > Reporting > Usage > UDB > Tables > Largest Tables	Scans
Which segments in an instance contain the most extents?	• **Oracle:** Data Manager for Databases > Reporting > Usage > Oracle > Segments > Most Extents	Scans
How much allocated space for an object is being wasted?	• **Oracle:** Data Manager for Databases > My Reports > System Reports > Oracle - Segments with Wasted Space	Scans

What do you want to know?	Which report answers this question?	Which schedule can I run to generate this report?
How many free extents does a tablespace have?	• **Oracle:** Data Manager for Databases > Reporting > Asset > By DBMS Type > Oracle > instance_name > By Tablespace	Probes
How many tables, indexes, and datafiles/containers are associated with a tablespace?	• **Oracle:** Data Manager for Databases > Reporting > Asset > System-wide > Oracle > Tablespaces • **UDB:** Data Manager for Databases > Reporting > Asset > System-wide > UDB > Tablespaces	Scans
How many tables, indexes, and datafiles/fragments are associated with a database?	• **SQL/Server:** Data Manager for Databases > Reporting > Asset > System-wide > SQL/Server > Databases • **Sybase:** Data Manager for Databases > Reporting > Asset > System-wide > Sybase > Databases	Scans
How can I determine which users have violated a quota on a specific instance?	• **All DBMSs:** Data Manager for Databases> Reporting > Usage Violations > Quota Violations > Instance Quotas > By Instance	Scans
What are the sizes and locations of Control Files and log files for an instance within the environment?	• **Oracle:** Data Manager for Databases > Reporting > Asset > By DBMS Type > Oracle > instance_name • **SQL/Server:** Data Manager for Databases > Reporting > Asset > By DBMS Type > SQL/Server > instance_name > Databases	Probes
How can I determine how close a user is to exceeding a quota?	• **All DBMSs:** Data Manager for Databases > Usage Violations > Quota Violations > All Quotas > By User	Scans

Report considerations and limitations

Keep in mind these considerations and limitations when viewing reports in Tivoli Storage Productivity Center.

Data collection for reports

Some columns within reports might not be populated with data for the storage resources that you are monitoring. Columns are empty if you did not configure Tivoli Storage Productivity Center to discover or collect the information intended for that column. Use discovery, probe, ping, or scan schedules to ensure that you collect information for all the columns that appear in reports.

When you specify a start time for using performance monitor data in a report, the start time might not coincide with the start of a performance monitor data collection interval. If the requested start time is after the start of a data collection interval and the duration of the data collection interval is long enough so that it ends at or after the requested start time, Tivoli Storage Productivity Center includes all data from the collection interval that began before the requested start time. This data might include data collected before the start time that you specified. This processing ensures

that the maximum amount of available performance data is used. For example, if a performance monitor collects data every 20 minutes starting at 01:00, but you request performance data that starts at 1:15, Tivoli Storage Productivity Center still analyzes performance data that starts at 01:00 that has a duration of 20 minutes.

Report charts that contain long path names
If you generate charts for reports and these charts contain long path names or many segments, the text on the charts might overlap and make them difficult to read.

Information about directories
Before information about directories is displayed in reports, you must complete the following tasks:

1. Include directories that were discovered by Tivoli Storage Productivity Center in a directory group.
2. Include the directory group in a data collection scan.

After a scan schedule collects data about the directories in a directory group, information about the directories is displayed in By Directory reports

Storage totals for Network-wide reports
Tivoli Storage Productivity Center provides a number of network-wide reports that display storage totals based on the data collected from the monitored computers and storage subsystems in your network. Click the magnifying glass icon next to a row in a network-wide report to view the details about the computers that are included in those storage totals. However, the reports that you drill down into display information about computers only. Storage information about storage subsystems is not included in the values for reports that are shown when you drill down from a network-wide report.

For example, the Data Manager > Reporting > Capacity > Disk Capacity > Network-wide report displays disk capacity totals for the computers and storage subsystems in your network that are monitored by Tivoli Storage Productivity Center. However, when you click the magnifying glass icon next to a row in that report, the resulting report displays disk totals for computers only. This action might cause the totals displayed in the network-wide report to be larger than the totals reflected in the report that you drill down in to (if you have disks that are part of both computers and storage subsystems within your network).

Report values for General Parallel File Systems (GPFS) file systems
For reports that display information collected by scans, Tivoli Storage Productivity Center does not include statistics for GPFS snapshot directories and the files in those directories. Because of this issue, the values in these columns might be smaller than the actual numbers for the file systems.

For example, when viewing data collected by scans of GPFS file systems in the Data Manager > TPC-wide Storage Space > File System Space report, the values for the File Count, Directory Count, and File System Space columns might be smaller than the actual numbers for the file systems.

IBM XIV Storage System reporting limitation for disk drives
Tivoli Storage Productivity Center does not collect physical serial number information about the disk drives that are associated with the XIV system

devices. Therefore, any reports that display information about the XIV system do not include the physical serial numbers of the disk drives.

Hitachi Data Systems storage subsystems limitations

- There is a limitation within reports when Tivoli Storage Productivity Center collects information about Hitachi Data Systems (HDS) storage subsystems that are configured as storage virtualizers. Specifically, any reports that display a value for the Disk Unallocated Space column do not count the storage used by volumes that exist on external disks when the HDS is used as a virtualizer.
- Tivoli Storage Productivity Center does not support Logical Unit Size Expansion (LUSE) with Hitachi storage subsystems. Because of this limitation, Tivoli Storage Productivity Center might report a different number of volumes for a subsystem than is reflected in the HiCommand Device Manager (HDvM) when LUSE is used.

Time zones for servers that are monitored by Storage Resource agents

The time zones of servers that are monitored by Storage Resource agents are shown as Greenwich mean time (GMT) offsets in Tivoli Storage Productivity Center reports. For example, a server in Los Angeles shows the following time zones in the By Computer report in Asset reporting:

```
(GMT-8:00) GMT-8:00
```

SQL access to data collected by Tivoli Storage Productivity Center for use in external reports and applications

You can use the data collected by data collection jobs to create custom reports in external tools. Tivoli Storage Productivity Center provides a set of DB2® views that represent key data that was stored in the database repository. Use Structured Query Language (SQL) commands to retrieve the data in these views. Use this data in external tools such as Business Intelligence and Reporting Tools (BIRT), Microsoft Excel, and Tivoli Common Reporting. See *Planning to retrieve information from views in the database repository* topic in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for more information about how to use SQL to retrieve data for use in custom reports and tools.

Storage Pools in EMC storage subsystems

The data that Tivoli Storage Productivity Center collects about EMC storage subsystems is based on SMI-S standards. Within SMI-S model, different types of storage pools are defined and several subsystem attributes can be modeled as special pool types. Beside the pools that are available for volume creation, additional storage space from the EMC storage subsystem might be represented as a pool within the SMI-S model. Because of this representation, Tivoli Storage Productivity Center reports and CLI commands might display more pools for EMC storage subsystems than what is reported through the EMC Navisphere agent.

Computers in a Windows domain

Once computers are discovered in a domain, they remain in the database until three discoveries do not report them. If you remove the computer that is doing the discovery in a domain, no further discovery results are processed by the server. Therefore, any remaining domain computers that were previously discovered continue to show their unreachable value in Unmanaged Computers reports.

UDB reports

For UDB reports, the "User Name" and "Creator" columns represent the

schema (not the operating system user) which is using the storage space. Each schema is associated with one operating system user. The reports show usage at the schema level.

Reports display incorrect number of folders for Windows computers

On monitored computers that run Windows, the number of folders that is reported by Tivoli Storage Productivity Center might be different from the number of folders that is indicated by Windows Explorer. This situation occurs when settings on the computer indicate that files, folders, and drives are hidden. For example, the Data Manager > Reporting > Usage > files > File Summary > by Computer report might display a number of folders for a Windows computer that is less than the number of folders provided by Windows Explorer on that computer.

To ensure that the number of folders reported by Tivoli Storage Productivity Center more closely matches the number reported by Windows Explorer, complete the following steps:

1. On each of the Windows computers that you are monitoring, start Windows Explorer.
2. Navigate to the file system whose settings that you want to change.
3. Select **Folder and search options...** or **Folder options...** from the **Organize** or **View** menus.
4. Click the **View** tab.
5. In the **Advanced Settings** section, make the following changes:
   - Select **Show hidden files, folders, and drives**.
   - Clear **Hide protected operating system files**.
6. Run probe and scan data collection schedules on the computers.
7. View the reports that display the number of folders for Windows computers.

Information about space efficient volumes

Keep in mind the following information about space-efficient volumes:

- A volume name or virtual disk in a report that is marked with an asterisk (*) indicates that the volume or virtual disk is a space-efficient volume. The * is not part of the actual volume or virtual disk name.
- Tivoli Storage Productivity Center does not include space efficient volumes when calculating the values for consumed and available space on DS8000 storage systems. Therefore, the values for these columns in Tivoli Storage Productivity Center reports might be different from the space that is available on DS8000 storage systems with space efficient volumes.
- Tivoli Storage Productivity Center does not display information about space-efficient volumes for IBM DS8000 storage systems in the following reports:
  - Disk Manager > Reporting > Storage Subsystems > Computer Views > By Filesystem/Logical Volume
  - Disk Manager > Reporting > Storage Subsystems > Storage Subsystem Views > Disks

The mapping of volumes to an ESX hypervisor might not display correctly

When assigning and unassigning volumes to an ESX hypervisor, the mapping of the volumes might not display correctly in Tivoli Storage Productivity Center reports. For example, if you assign a volume to a hypervisor, unassign it, and then assign a different volume to the

hypervisor, information about the original volume might still be displayed in the Correlated Storage Subsystem volume report that is available under Data Manager > Reporting > Asset > By Hypervisor.

To ensure that Tivoli Storage Productivity Center reports display the correct volume information, complete the following tasks before generating the reports:

- After unassigning a storage system volume from a hypervisor, use VMware vSphere Client to rescan the HBA.
- After assigning a new storage system volume to the hypervisor, use VMware vSphere Client to rescan the HBA.
- In Tivoli Storage Productivity Center, rerun probes for the computer, hypervisor, and storage system.

IBM Storwize V7000 Unified/IBM SONAS cluster names in reports

In reports that display file level information for IBM Storwize V7000 Unified/IBM SONAS clusters, the storage system name is included with the cluster name in the Computer or Device column. The format of the column is `storagesystem_name: cluster_name`, where storagesystem_name is the name of the IBM Storwize V7000 Unified/IBM SONAS device that the cluster belongs to. For example, tpcserver1.tpc.example.com: kq98n5d.acct displays storage system tpcserver1.tpc.example.com that cluster kq98n5d.acct belongs to.

Pool space for NetApp filers

For NetApp filers that are configured with block data, the values for pool space that are shown in reports are different from the values shown in a storage management tool for NetApp filers, such as FilerView. The values that are different include total pool space (capacity) and available pool space.

This discrepancy occurs because the pool space information that is collected by SMI-S agents for Tivoli Storage Productivity Center includes Write Anywhere File Layout (WAFL) and Snapshot reserve space, whereas tools like FilerView do not include this additional space.

Working with reports

Learn how to drill down through reports, navigate report tabs, view historical trends in reports, and view reports as tables of data or as graphical charts. The ability to customize your view of reports provides you with the flexibility to get a very detailed look at your information or a very high level overview.

Generating and viewing reports

Learn how to generate and view reports.

Each of the report types available in IBM Tivoli Storage Productivity Center are generated in different ways to provide you with the flexibility to meet your environment's changing needs. Click the following links to learn more about how to generate the different report types and determine what works best for your organization.

Report type	How to generate
User-defined online reports	"Generating reports" on page 1320
System reports	"Generating and viewing system reports" on page 1510
User ID (saved) reports	"Creating a user-defined report" on page 1574

Report type	How to generate
Batch reports	"Creating batch reports" on page 1562
Rollup reports	"Generating rollup reports" on page 1576

Sorting report columns

Learn how to sort the columns in reports.

Reports appear in a tabular format and are comprised of rows and columns. You can scroll the report to view the entire report. Use the **View** menu to hide/show the navigation tree to increase the viewable area of the report. You can also drag the divider bar in the middle of the screen back and forth to reallocate the amount of space that is given each pane.

To sort the columns or change the order of the columns in a report:

- **Re-sorting columns.** To resort a column, click on the column header. The column will resort in ascending or descending order.
- **Moving columns.** To change a column order, press <Alt>, and then click and drag the column you want to move.

Printing reports or charts

Learn how to print reports to a printer, a PDF, an HTML file, a CSV file, or a formatted text file.

You can print any IBM Tivoli Storage Productivity Center report or chart displayed in the content pane.

To print a report or chart, complete the following steps:

1. Generate the report you want.
2. Click **File > Print** or click on the tool bar. A print window opens.
3. Select the output that you want:
   - Select **Printer** to send the output directly to a printer
   - Select **PDF file** to save the output as a PDF file
   - Select **HTML file** to save the output as an HTML file
   - Select **CSV file** to save the output as a comma delimited file. Select **Include Headers** to include column headers in the report output file. Select **Include Totals** to include the totals column in the report output file.
   - Select **Formatted file** to save the output as a formatted text file
4. Click **Continue**.
   - If you selected **Printer**, the Page Setup window opens. Specify any custom settings for the print job on the Page Setup window and click **OK**. The report or chart is printed to your default printer.
   - If you selected **PDF file, HTML file, CSV file** or **Formatted file**, the Save As window opens. Select the location where you want to store the file and enter a name for the file in the **File Name** field. Click **Save**. Depending on the file type you selected, the report or chart is saved as a PDF or HTML file.

Including charts in reports

Learn how to create charts from the reports that you generate.

There are a large number of charts and chart types that can be generated from reports. All reports that support charting will have a button at the start of the report. Any row in a report that supports charting on that row will also have a .

Many charts can be generated simultaneously from the same report. After generating a chart, return to the **Report** tab and choose another chart or another set of objects to generate a chart upon. Another **Chart** tab will be generated in addition to the existing tabs and you will be positioned to the new chart tab.

See "Generating charts" for information on how to generate charts from reports.

You can also customize the appearance of charts to determine how you want to present data. For example, you can create pie chart, bar chart, stacked bar chart, determine how many charts appear per page, specify time increments, and so on. See "Customizing charts" on page 1314 for information on how to customize charts.

Generating charts:

Learn how to generate charts.

To view a chart of all of the rows in the table:

1. Click , which is at the start of the report. All reports that support charting will have this multi-colored pie chart icon at the start of the report.
2. Select the type of chart you would like to generate. The types of charts available will vary depending on the report. After you select the chart type, a new tab for that chart will appear in the content pane.

To generate a chart for just the element in that row:

- Click the chart icon for that report row.

To generate a chart for multiple, non-consecutive rows in a report:

1. Select rows in the report by holding down the control key (Ctrl key) on the keyboard while clicking on the required rows.
2. Click a button for any of the selected rows and select the type of chart you would like to generate. After you select the chart type, a new tab for that chart will appear in the content pane.

To generate a chart for a range of rows in a report:

1. Click on the first row in the range and then hold down the Shift and Ctrl while clicking on the end row in the range.
2. Click a button for any of the selected rows and select the type of chart you would like to generate. After you select the chart type, a new tab for that chart will appear in the content pane.

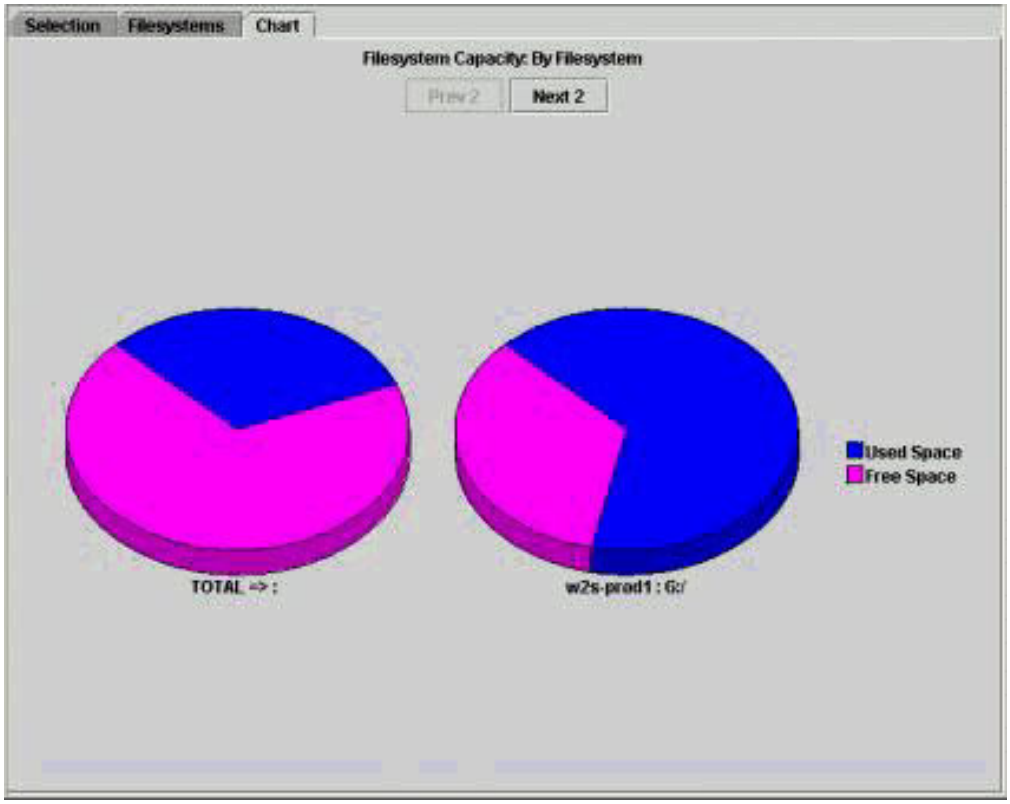

Figure 34. Chart Generation

Customizing charts:

Learn how to customize the appearance of charts.

While viewing a chart, you can customize its appearance using the Chart Customization window. This includes:

- Entering the title of a chart
- Determining the maximum number of charts or series that appear per screen
- Specifying how to show values for data points or areas
- Selecting a type for a chart such as a pie, bar, stacked bar, and 100% stacked bar
- Determining the orientation (vertical or horizontal) of a chart
- Specifying history chart settings such as trends and time increment values

To access this window from a chart, right-click on the chart and select **Customize this chart** from the pop-up menu. The Chart Customization window appears.

Use this window to customize the chart generated from a report, a report created in the Reporting facility, a batch report, user reports, and dashboard reports. You can use the settings on this window to:

- Customize the display of the current chart.
- Customize the display of the current chart and save those changes. If you save the customization settings for a report chart, that chart will use those settings whenever it is generated.

If a field on this window does not apply to a selected report it will appear as grayed out and cannot be edited. The following list describes the fields:

Chart title
Enter the name you want to give to the generated chart.

Maximum number of charts or series per screen
Select the maximum number of charts or series per page that should appear in the generated report from this list. For example, how many pie charts will appear on one page. When the number of charts that appear on one page exceeds the number in this field, the additional charts will be shown on a separate page. The most charts or series you can define to appear on a page is 30.

Show values for data points or areas
Determine how to display the data points/areas on a chart. You can choose to display the data points areas:

- Only when hovering - display values for data points on an online chart when the mouse hovers over an appropriate place on that chart
- Always - always display and print values for data points

Graph charts
Use the fields in this section to customize the appearance of graphs and charts.

Chart type
Select a type for the generated chart including:

- Pie Chart
- Bar Chart
- Stacked Bar
- 100% Stacked Bar

Chart orientation
Determine whether the chart graphic will appear vertically or horizontally on a page. You can set the orientation for the following chart types only: Bar Chart, Stacked Bar, and 100% Stacked Bar.

Maximum number of series per graph
Select the maximum number of series that should appear in one graph or chart. For example, how many slices will appear in one pie chart. When the number of series that appear in one graph exceeds the number in this field, the additional charts will be shown on a separate page. The default maximum is 12. The most series that can appear in one chart is 30.

Note: Disk Capacity batch reports always display one series per chart, and one chart per page.

History Charts
Use the fields in this section to customize the appearance of history charts.

Show Trends
Check this box if you want IBM Tivoli Storage Productivity Center to show a trend for the data in a history chart. The trend attempts to show future storage usage in a chart based on earlier usage.

Time increment for x-axis
Select which value to use for the time increment that appears on the x-axis of a history chart. You can select By Scan, Weekly, Monthly. For example, if you choose Monthly, the x-axis will measure storage usage in monthly increments.

Restore Defaults
Reset the values in the fields on this window to their original default settings.

Apply Apply the settings defined in this window to the current chart.

Save Apply the settings defined in this window to the current chart and save those settings. When the chart is generated in the future, it will automatically use the saved settings.

Cancel
Exit the window without applying or saving the settings in the window.

Drilling down a report row

Learn how to drill down into report rows to view more detailed information about your storage resources.

Many of the rows that are presented in the generated reports allow you to drill down and get more reports of more detailed information. For example, if you are looking at a report on computers, you might be able to drill down on a row for a specific computer and generate a report for the file systems contained on that computer.

When detailed drill down is allowed within a report, there is a icon for that row. To drill down, simply click on and a new tab with the drill down report will be generated and added to the tab dialog. You will then be positioned on that tab viewing the new detail report that you have generated.

Navigating tabs

Learn how to navigate report tabs to view detailed sub reports.

As you generate reports and drill down on rows to generate detailed sub reports, the number of tabs in the Reporting window will grow. You can move about from tab to tab simply by clicking on the tab name.

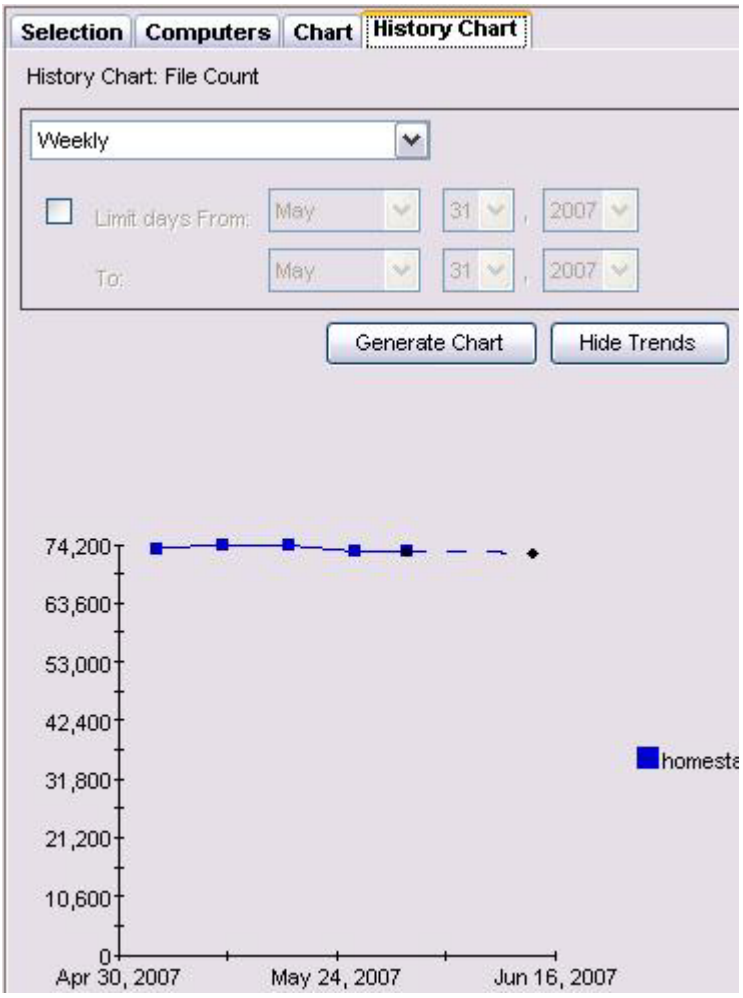

Figure 35. Tabs

You can also return to the **Selection** tab and change your selection criteria to start all over again with a newly generated report at any time. If you change your selection criteria and generate a new report, all subsequent or dependent tabs (for example, drill down and chart tabs) are closed.

To close a tab, right-click on the tab name and select **Close** from the pop-up menu.

Trending with historical data

Learn how to view trends related to the capacity and usage of your storage resources.

Use the trending of historical data to view patterns of your data across your entire network. For example, if you set up a graph that shows the disk capacity of a computer between January 1 and December 31, you can view "trended" data that will predict the disk capacity for that computer over the next 6 months. The trended data appears as a dashed line that is appended to the last probed or scanned data point.

The length of time for how far in the future that trended data will show is determined by 50% of the length of time displayed in the graph. For example, if the data in the non-trended portion of the graph displays data between January 1 and December 31, the predicted trend for that data will be shown for the next 6

months (January 1 to December 31 = 12 months, 50% of 12 months = 6 months trended). You can change the amount of time trended by changing the length of time of the graphed data using the **Limit days from** and **To** fields.

Before viewing trends in the data collected by IBM Tivoli Storage Productivity Center, use the history aggregator function to define and run jobs to sum data in the database repository for historical reporting purposes. For example, after running the history aggregator, you can view the sum of usage across multiple resources, by database.

To view a historical data, generate a report, right-click on the row or rows for which you want to view the data, and select an option from the menu.

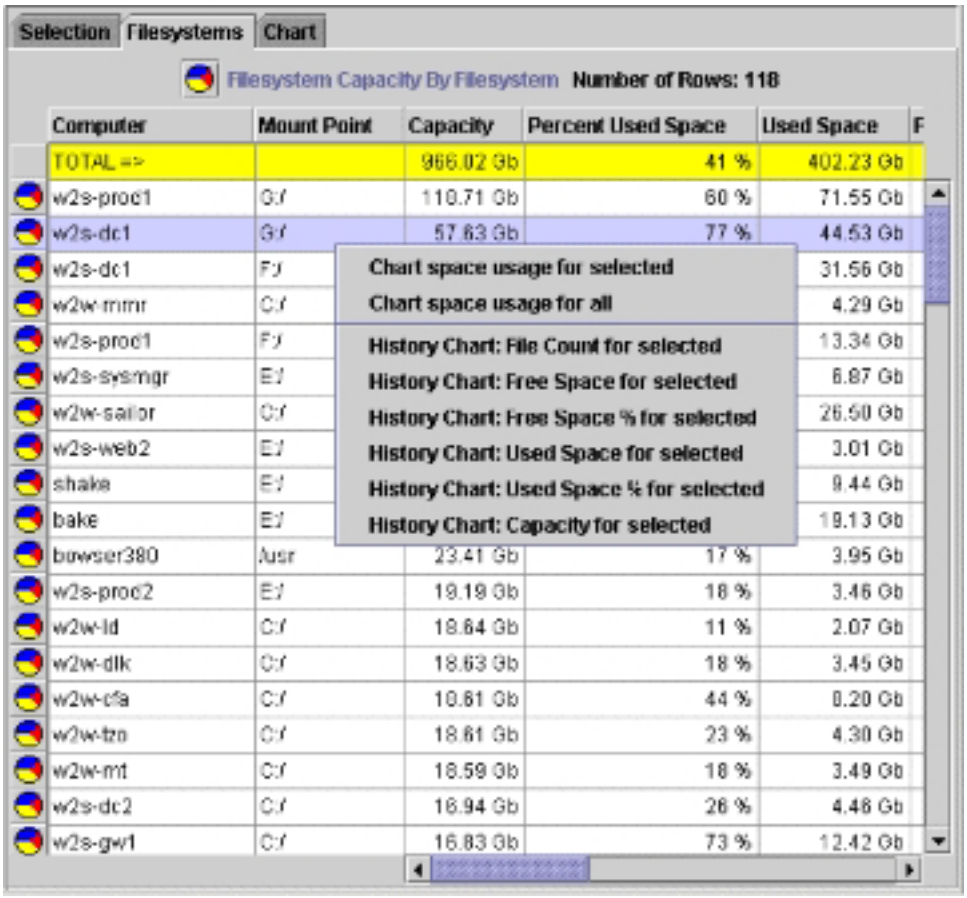

Figure 36. Trending with Historical Data

Depending on the type of report you generated, you can create the following history charts for the storage resources that you select:

History chart	Description
Free Space for Selected	Shows the free space available over a specified period of time.[1]
Free Space % for Selected	Shows the percentage of free space available over a specified period of time.[1]
Used Space for Selected	Shows used space over a specified period of time.[1]
Used Space % for Selected	Shows percentage of used space over a specified period of time.[1]

History chart	Description
Capacity for Selected	Shows capacity for selected over a specified period of time.[1]
Unallocated Space for Selected	Shows the unallocated space available for the selected resources over a specified period of time.[1]
Unallocated Space for All	Shows the unallocated space available for all the storage resources in the report over a specified period of time.[1]
File Count	Shows the number of files accessed or modified over a period of: • 1 day • 1 day - 1 week • 1 week - 1 month • 1 month - 1 year • > year
File Count %	Shows the percentage of files accessed or modified over a period of: • 1 day • 1 day - 1 week • 1 week - 1 month • 1 month - 1 year • > year
Size	Shows the size of files accessed or modified over a period of: • 1 day • 1 day - 1 week • 1 week - 1 month • 1 month - 1 year • > year
Size %	Shows the percentage of file size accessed or modified over a period of: • 1 day • 1 day - 1 week • 1 week - 1 month • 1 month - 1 year • > year

[1] The time options for data points in a chart are **By Aggregator**, **Weekly**, and **Monthly**. You can select one of these options from the list menu on the **History Chart** page for the report. (See Figure 37 on page 1320.)

- Select **By Aggregator** to generate a chart in which its data points correspond to each time a history aggregator job is run. By default, the history aggregator runs daily at 3:30 a.m., so when you generate a history chart with this default setting there is a point in the chart for every day in the selected time range. If you change the schedule of the history aggregator, the points in the history chart that you generate change accordingly.
- Select **Weekly** to generate a chart that includes one data point for each week in the selected time range.
- Select **Monthly** to generate a chart that includes one data point for each month in the selected time range.

Figure 37 on page 1320 shows a history chart for used space for a file system:

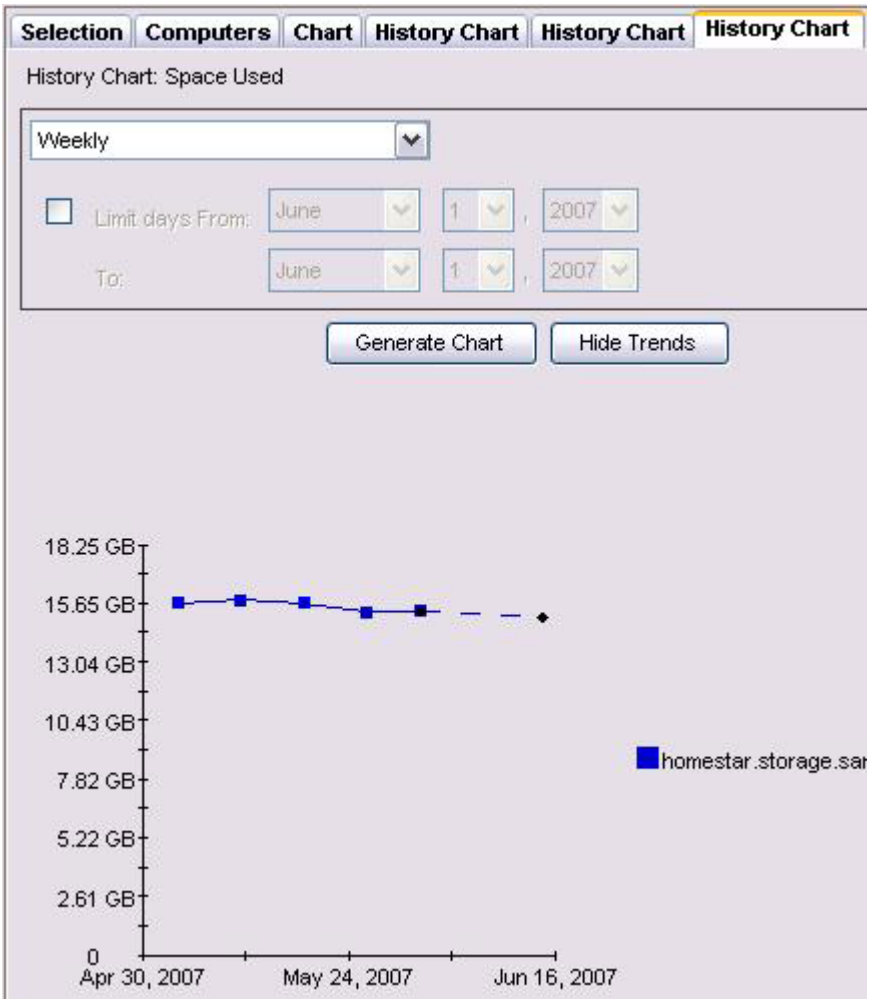

Figure 37. History Chart for File System Used Space

Select **File → Export data** to export history data to a comma delimited file (.csv). You can include or not include headers in the comma delimited file that you create. The column headings within the header reflect the type of report from which you are exporting history data.

Working with user-defined online reports

Use the information in this section to learn how to generate and view online reports. These reports are available in the **Reporting** nodes under Data Manager, Data Manager for Databases, Disk Manager, and Fabric Manager.

Generating reports

Learn how to generate and view online reports in Data Manager, Data Manager for Databases, Disk Manager, and Fabric Manager.

1. Expand the navigation tree to display the type of report you want to generate. For example, if you want to generate a disk capacity report, expand **Data Manager** > **Reporting** > **Capacity** > **Disk Capacity**.
2. Highlight a By Clause to determine how you want to view a report. For example, to view capacity information according to computer group, click **Reporting** > **Capacity** > **Disk Capacity** > **By Computer Group**. The **Selection** page is displayed.

3. Use the **Selection** page to select the profile to use when generating a report and determine what columns appear within a report. See "Selecting resources to display on a report" on page 1324 for information.
4. Click **Selection...** to select the objects that you want to report upon from the Select Resources window.
5. Click **Filter...** to further filter the objects that appear in a report. Filters enable you to apply general rules to the report based on the rows in that report. See "Filtering the resources that appear on a report" on page 1326 for more information on filtering.
6. Use the **Available Columns** and **Included Columns** list boxes to determine what columns are displayed in a generated report.
7. Click **Generate Report**. A new tab will be added to the tab dialog representing the report that you generated.
8. Click the new tab to view the report. Reports are tabular in format and composed of rows and columns. You can scroll the report to view the entire report. Use the **View** menu to hide/show the navigation tree to increase the viewable area of the report or drag the divider bar in the middle of the screen back and forth to reallocate the amount of space that is given each pane.

Defining options for a report:

Learn how to customize a report before you generate it.

Depending on the report that you want to generate and the level upon which you choose to report (by computer, by file system, network wide, and so on), you are typically presented with a **Selection** page. For example, if you want to generate the Most Obsolete Files report, expand **Data Manager** > **Reporting** > **Usage** > **Files** > **Most Obsolete Files** > **By Computer**. Choose items in the **Selection** page.

- Select the profile to use when you generate a report
- Select the resources (computers, computer groups, storage subsystems, and so on) to display in a generated report
- Apply filters to the columns within a report to determine what resources are displayed within the rows of a report. For example, show only the systems that contain over 80% used space.
- Determine what columns display within a report.

The following table describes the interface elements that might be displayed on the **Selection** page (depending on the report that you want to generate):

Interface element	Description
Profile (Data Manager only)	This field indicates the profile to use when the report is generated. If you want to use the supplied default profile, then leave the field as is and generate the report. However, if you want to use a customized profile that you set up for this reporting, choose that profile from the list. Keep in mind that profiles control what statistics are gathered and also what files are included in that statistical analysis.

Interface element	Description
Return maximum of __ rows per *storage_resource*	This field controls how many rows are displayed in a generated report. The value that displays as the *storage_resource* in this field depends on the selected By clause. For example, if you select **Data Manager** > **Reporting** > **Usage** > **Files** > **Most Obsolete Files** > **By Computer**, you can enter 10 in the **Return maximum of __ rows per computer** field and to limit the number of rows in the report to only 10 obsolete files per computer. If you set up scans to collect many more rows than that per computer, this selection is an effective and efficient way to control the size of the generated report. With this field, you can control the size of the report that is generated for different By clauses.
Limit to this time range:	The fields determine the date range from which to gather availability information for the objects that you select.
Relate Computers to: (Disk Manager - Storage Subsystem reports only)	This field generates a report that shows the relationship between the computer and one of these options: storage subsystems, volumes, or disks. • Storage Subsystem. Choose this option to view the Storage Subsystems that the logical volumes within a computer are using. • • Volumes. Choose this option to show the volumes within a storage subsystem that the logical volumes within a computer are using. • Disks. Choose this option to show the disks within a disk array that the logical volumes within a computer are using.
Relate File systems/Logical Volumes to: (Disk Manager - Storage Subsystem reports only)	This field generates a report that shows the relationship between the file systems/logical volumes and one of these options: storage subsystems, volumes, or disks. • Storage Subsystem. Choose this option to view the storage subsystems on which a file system/logical volume is allocated. • Volumes. Choose this option to view the volumes on which the file system/logical volume is allocated. • Disks. Choose this option to show the disks on which a file system/logical volume is allocated.
Relate Storage Subsystems to: (Disk Manager - Storage Subsystem reports only)	This field generates a report that shows the relationship between a storage subsystem and one of these options: computer or file system/logical volume. • Computer. Choose this option to view the computers that are using a storage subsystem. • File system/Logical Volume. Choose this option to view the file systems/logical volumes that are using a storage subsystem.
Relate Volumes to: (Disk Manager - Storage Subsystem reports only)	This field generates a report that shows the relationship between a volume and one of these options: computer or file system/logical volume. • Computer. Choose this option to view the computers that are using a volume. • File system/Logical Volume. Choose this option to view the file systems/logical volumes that are using a volume.

Interface element	Description
Relate Disks to: (Disk Manager - Storage Subsystem reports only)	This field generates a report that shows the relationship between a disk and one of these options: computer and file system/logical volume. • Computer. Choose this option to view the computers that are using a disk. • File system/Logical Volume. Choose this option to view the file systems/logical volumes that are using a disk.
Selection...	This button is used to determine the storage resources (computers, file systems, logical volumes, storage subsystems, and so on) upon which you want to report. When you click **Selection...**, the Select Resources window opens. If you clear the check box for a resource on the Select Resources window, a push-pin icon displays on the **Selection...** button. If you are generating a Network-wide report, the **Selection...** button is disabled because you chose to generate the report network-wide, for all objects.
Filter...	Use this button to further select or restrict the values within the columns you want to include on a report. You can apply general rules to the selected columns. When you click **Filter...**, the Edit Filter window opens. For example, select all computers with the name beginning with `w2s-*`, or select all file systems with percent used space >= to a specific value. If a filter is defined for this report, a pushpin icon displays on the **Filter...** button. **Note:** For storage entities that contain an asterisk (*) within their names, the * is used as a wildcard character if you select `LIKE` from the **Operator** menu on the **Edit Filter** panel. For more information, see "Filtering the resources that appear on a report" on page 1326.
Available Columns / Included Columns	These panels display the columns that can be included or are already included in the generated report. • The **Available Columns** list displays the columns that you can include in the generated report that is not already included. • The **Included Columns** list displays the columns in the generated report. By default, all the columns for a report are listed in this list box and are displayed in the generated report. You can highlight the columns and use the up and down arrows to rearrange columns in the report. The order of the items in this list determines the order in which the columns display in the generated report. The columns with shaded-color names are unavailable columns. They cannot be selected, removed, or reordered. To include a column in a report: 1. In the **Available Columns** list, click the name of the column. To select more columns, press **Shift** and click the columns that you want to add. To select a range of columns, press **Ctrl** and click the first and last column. 2. Click the right-arrow button (>>) to add the columns to the **Included Columns** list. To remove a column in a report: 1. In the **Included Columns** list, click the name of the column. To select more columns, press **Shift** and click the columns that you want to add. To select a range of columns, press **Ctrl** and click the first and last column. 2. Click the left-arrow button (<<) to add the columns to the **Available Columns** list.

Selecting resources to display on a report:

Learn how to select the storage resources that displays in a report.

Use the Select Resource window to determine the objects (computers, file systems, logical volumes, storage subsystems, and so on) upon which you want to report. To select the resources on which you want to report:

1. Select the type of report you want to generate. For example, **Data Manager** > **Reporting** > **Capacity** > **Disk Capacity** > **By Disk**.
2. Click **Selection**. The Select Resources window opens.

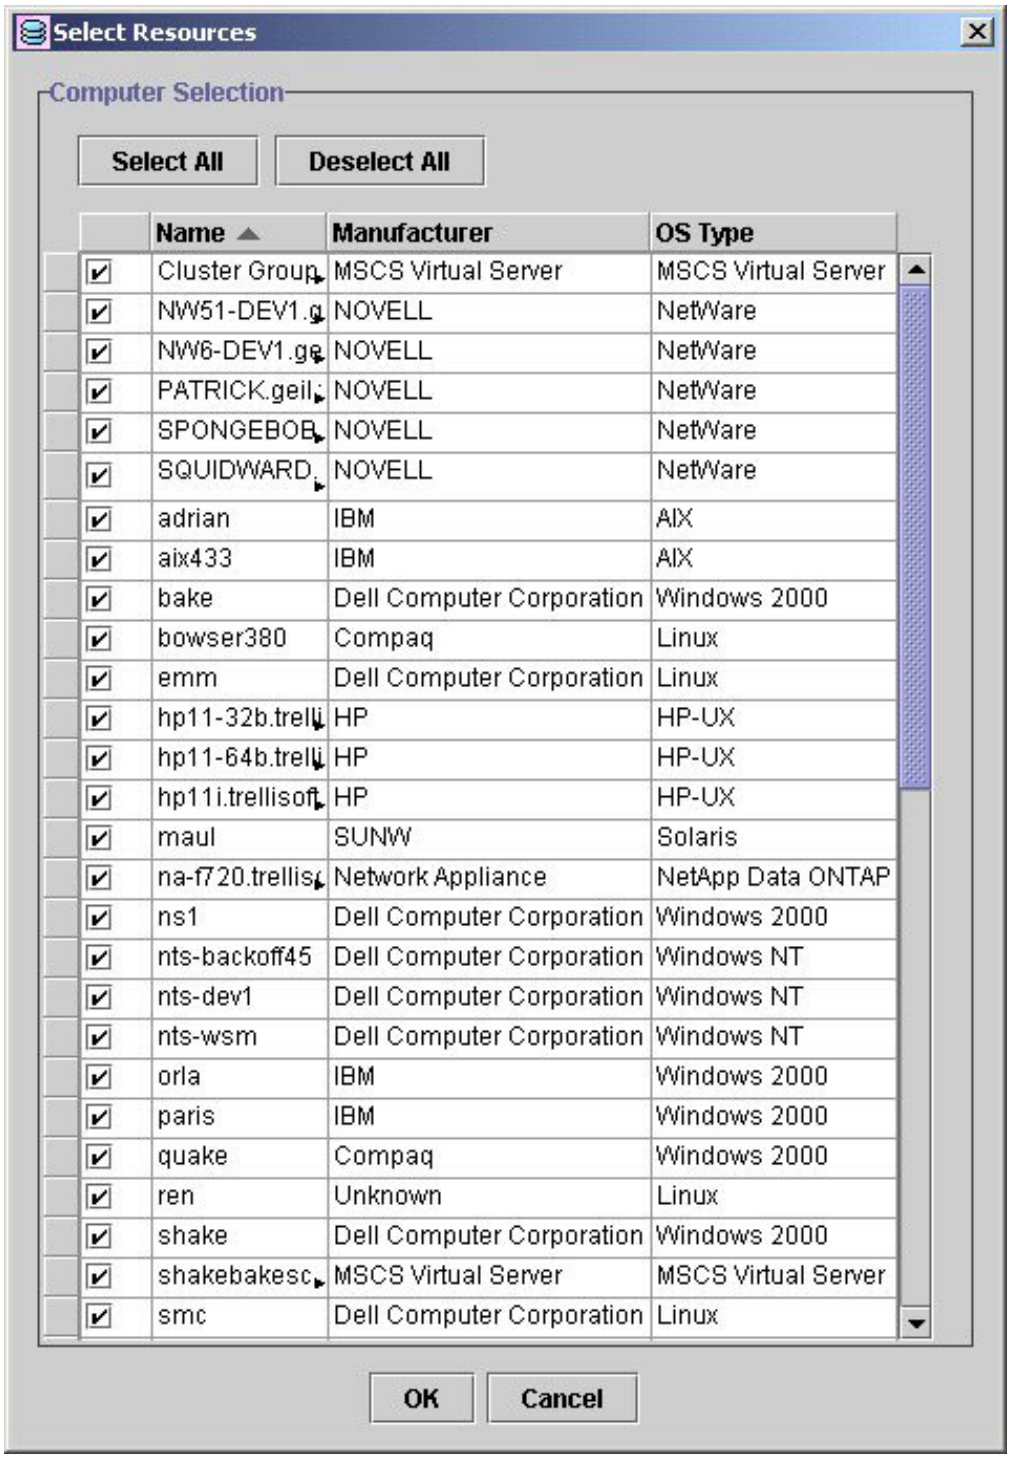

Figure 38. Selection Resources window

The objects that display in this window are determined by the report you are generating and the By-Clause you selected. For example, if you choose to generate the report **Data Manager** > **Reporting** > **Capacity** > **Disk Capacity** > **By Disk**, the Select Resources window displays a list of computers and subsystems that you can include in the report.

When the Select Resource window is initially opened, all objects in the list are selected.

- If you would like to generate the report for all the objects that are selected, click **OK**.
- To select/deselect individual objects in the selection list, select the check box for the object. When there is a check mark in the check box, then the object is selected for inclusion in the report.

You can use **Select All** or **Deselect All** to quickly select or deselect all of the objects in the list.

Note: If you generate a Network-wide report, the **Selection** button is disabled on the **Selection** tab because you chose to generate the report network-wide, for all objects.

Filtering the resources that appear on a report:

Learn how to apply conditions to the columns in a report to further filter the resources that appear in a report.

Use the Edit Filter window to apply conditions to the columns within a report to further select or restrict objects on which you want to report. Filters enable you to apply general rules to the report. For example, you can show only the file systems that contain over 80% used space, or show computers that match "w2s-*", or select all disks with percent used space >= to a specific value. To filter the objects that appear in a report:

1. Select the type of report you want to generate. For example, **Data Manager** > **Reporting** > **TPC-wide Storage Space** > **Disk Space** > **By Disk**.
2. Click **Filter...**. The Edit Filter window opens.

Note: For storage entities that contain an asterisk (*) within their names, the * is used as a wildcard character if you select **LIKE** from the **Operator** menu on this panel. For example, if you enter ***XYZ** in the **Value 1** field and select **LIKE** from the **Operator** menu, the resulting report includes any entities named "*XYZ", and also all entities whose names end with the letters "XYZ". To create a filter that includes only the storage entity named "*XYZ" in the resulting report, enter ***XYZ** in the **Value 1** field and select **IN** from the **Operator** menu.

The following table describes the interface elements that appear on the **Edit Filter** window:

Interface element	Description
Add	This button is used to specify multiple conditions for the filter. When you click **Add**, report columns will be displayed enabling you to define conditions for those columns.
Delete All	This button is used to delete all the conditions displayed on the table.
Case-sensitive string comparisons	This check box indicates that case-sensitive string comparisons should be done. Select this check box to ensure that values appearing in the generated report match the case of the values entered in the Value 1 and Value 2 fields.

Interface element	Description
Records must meet	This field specifies how the conditions in the filter are grouped together. • **All conditions**. All of the entered conditions (rows on the dialog box), must be true in order for the record to show up as a row in the report. • **At least one condition**. At least one of the entered condition must be true in order for the record to show up as a row in the report.
Column	This field is used to select specific columns on which you want to define conditions. The list on this field shows the columns available for the report. As you select a column from the list, it is displayed on the table, and the column name is removed from all other lists. Click to delete the column.
Operator	This field is used to select the condition's operator. Not all operators apply to all the column data-types.
Value 1/ Value 2	This field is used to specify fixed value(s) to be compared against the record's column value. Depending on the operator and data type, this field may be absent or display-only. If the field is display-only, an **Edit...** button will be enabled. If both this field and the **Edit...** button are enabled, this field allows shorthand value entry. If you click **Edit** , a window opens to provide more assistance in editing the value. A **Value 2** field will be present only for (NOT) BETWEEN. For certain columns that contain string text, you can: • put the string in quotes for the **LIKE** operator, or • use the **=** operator without quotes instead of LIKE (the default) For example, if you want to filter the Device Path (computer) column on the **Reporting → Storage Subsystems → Computer Views → By Filesystem/Logical Volume** report, you can search/filter on a name such as **"Disk 1"** or **"Disk 2"** with the **LIKE** operator. For the **=** operator, you can enter **Disk 1** or **Disk 2** with no quotes. **Note:** An asterisk is prepended to the names of space-efficient VDisks in the IBM Tivoli Storage Productivity Center database repository. When you use a filter to determine which volume names appear in a report, you must include the "*" as the first character in the filter value to ensure that those VDisks are included. For example, if you want to include VDisks whose name begins with "*vdisk1", you must enter ***vdisk1** in the **Value 1** field and select **LIKE** or **IN** from the **Operator** menu. If you select the **LIKE** operator, the *vdisk1 VDisk is included in the report, as well as other VDisks whose names end in "vdisk1".

Interface element	Description
Edit	This button is used to edit values pertaining to size, date, and for specific values used for IN, LIKE, etc. operators. Depending on the condition you are creating, a dialog box displays to help you to enter values that correspond to the condition.
Selection...	This button is used to determine the objects (computers, file systems, logical volumes, storage subsystems, etc.) upon which you want to report. The Select Resources window opens. If you clear the check box for a resource on the Select Resource window, a pushpin icon displays on the **Selection...** button. If you are generating a Network-wide report, the **Selection...** button will be disabled because you have chosen to generate the report network wide, for all objects.
OK	This button is used to save the filter. You are returned to the **Selection** tab. The **Filter** button will display a pushpin to indicate that a filter has been defined for the report.
Cancel	This button is used to return to the **Selection** tab without saving the filter.

Reporting groups

A reporting group is a user-defined set of storage assets, monitoring groups, and reporting groups upon which you want to report.

Use reporting groups to view storage information about:

- Resources that display in different monitoring groups. Because a resource can display in only one monitoring group at a time, you can use reporting groups to include resources from different monitoring groups and view reports on them.
- A subset of resources within a monitoring group. For example, you can create a reporting group for a specific subset of computers with the Payroll monitoring group and generate a report that contains information about those computers.
- Resources that display in other reporting groups.
- A subset of resources within other reporting groups.

Note: Unlike monitoring groups, a resource can be displayed in one or more reporting groups.

You can create the following types of reporting groups:

- Data Manager: **Computer**, **File system**. Within a computer-reporting group, you can include computer monitoring groups, specific computers, and other reporting groups. Within a file system reporting group, you can include file system monitoring groups, specific file systems, and other reporting groups.
- Data Manager for Databases: **Computer**, **Databases-Tablespaces**. Within a computer reporting group, you can include computer monitoring groups, specific computers, and other computer reporting groups. Within a databases-tablespaces reporting group, you can include databases-tablespace monitoring groups, specific databases or table spaces, and other databases-tablespaces reporting groups.

- Disk Manager: **Storage Subsystems**. Within a storage subsystem reporting group, you can include storage subsystem monitoring groups, specific storage subsystems, and other storage subsystem reporting groups.

When you generate a report, storage statistics about a reporting group displays in the corresponding "By Group" clause for that report. The following lists the reports that reflect reporting group information.

Computer reporting groups (Data Manager and Data Manager for Databases)
The following reports display computer reporting group information:

- **Data Manager** > **Reporting** > **Capacity reports**, **Data Manager for Databases** > **Reporting** > **Capacity reports**
- **Data Manager** > **Reporting** > **Usage reports**, **Data Manager for Databases** > **Reporting** > **Usage reports**
- **Data Manager** > **Reporting** > **Usage Violation reports**, **Data Manager for Databases** > **Reporting** > **Usage Violation reports**
- **Data Manager** > **Reporting** > **Backup reports**
- **IBM Tivoli Storage Productivity Center** > **Rollup Reports** > **Capacity**
- **IBM Tivoli Storage Productivity Center** > **My Reports** > **Batch Reports** > **Create Batch Report** > **Report page**, select one of the following items:
  - **Capacity** > **capacity_report_type** > **By Computer Group**
  - **Usage** > **usage_reporting_type** > **By Computer Group**
  - **Usage Violations** > **usage_violation_reporting_type** > **By Computer Group**
  - **Backup** > **backup_reporting_type** > **By Computer Group**
  - **Rollup Reports** > **Capacity** > **capacity_report_type** > **By Computer Group**
  - **Rollup Reports** > **Database Capacity** > **rdbms_type** > **By Computer Group**

File system reporting groups (Data Manager)
The following reports display file system reporting group information:

- **Data Manager** > **Reporting** > **Capacity reports**
- **Data Manager** > **Reporting** > **Usage reports**
- **Data Manager** > **Reporting** > **Usage Violation** > **Constraint Violations reports**
- **Data Manager** > **Reporting** > **Backup reports**
- **IBM Tivoli Storage Productivity Center** > **My Reports** > **Batch Reports** > **Create Batch Report** > **Report page**, select one of the following items:
  - **Capacity** > **capacity_report_type** > **By File System Group**
  - **Usage** > **usage_reporting_type** > **By File System Group**
  - **Usage Violations** > **Constraint Violations** > **By File System Group**
  - **Backup** > **backup_reporting_type** > **By File System Group**

Databases-Table spaces reporting groups (Data Manager for Databases)
The following reports display databases reporting group and table spaces reporting group information:

- **Data Manager for Databases** > **Reporting** > **Capacity reports**
- **Data Manager for Databases** > **Reporting** > **Usage reports**

- **Data Manager for Databases** > **Reporting** > **Usage Violations** > **Database-Table Space Quotas reports**Data Manager for Databases > Reporting > Usage Violations > Database-Tablespace Quotas reports

Storage subsystem reporting groups

The following reports display storage subsystem reporting group information:

- **Data Manager** > **Reporting** > **Storage Subsystems** > **Storage Subsystem Views** > **By Storage System Group**
- **Data Manager** > **Reporting** > **Capacity** > **Disk Capacity** > **By Storage Subsystem Group**
- **IBM Tivoli Storage Productivity Center** > **My Reports** > **Batch Reports** > **Create Batch Report** > **Report page**, select one of the following items:
  - **Storage Subsystems** > **Storage Subsystem Views** > **By Storage Subsystem Group**
  - **Capacity** > **Disk Capacity** > **By Storage Subsystem Group**

Creating a reporting group:

To create a reporting group:

1. Expand **Reporting** > **Groups** under:
   - Data Manager to create computer and file system reporting groups.
   - Data Manager for Databases to create computer and databases-tablespaces reporting groups
   - Disk Manager to create storage subsystem reporting groups
2. Right-click the type of group you want to create. A pop-up menu opens.
3. Select **Create <group type>** from the pop-up menu. The topic pane displays a window for defining the group.
4. Define information about the group. See the IBM Tivoli Storage Productivity Center online Help for detailed information about the windows for creating a group.
5. Click **File** > **Save** to save the group definition. The group is saved to the database repository. You can generate and view reports that are based on these groups.

Note: You can define filters to select specific computers, file systems, storage subsystems, table spaces, and databases for a reporting group. Tivoli Storage Productivity Center uses these filters to dynamically include resources to the group as they are discovered within your environment. When you define dynamic filters for reporting groups, the filters you enter might or might not be case-sensitive. The sensitivity depends on the database under which the Tivoli Storage Productivity Center database repository is stored. For example:

- The dynamic filters that you enter for reporting groups are case-sensitive when the database repository is stored on Oracle, Sybase, and UDB/DB2.
- The dynamic filters that you enter for reporting groups are not case-sensitive when the database repository is stored on Microsoft SQL Server.

Keep in mind these case-sensitive issues when you define filters for reporting groups. Prevent inaccurate reporting that is caused by unintended resources that are being included in reports.

Editing a reporting group:

To edit an existing reporting group:

1. Expand **Reporting** > **Groups** under:

- **Data Manager** to edit computer and file system reporting groups.
- **Data Manager for Databases** to edit computer and databases-tablespaces reporting groups
- **Disk Manager** to edit storage subsystem reporting groups

2. Expand the type of the reporting group that you want to edit. For example, expand **Disk Manager** > **Reporting** > **Groups** > **Storage Subsystem** to edit a storage subsystem reporting group.
3. Click the reporting group that you want to edit. The topic pane displays a window for editing the group.
4. Edit information about the group. See the IBM Tivoli Storage Productivity Center online Help for detailed information about the windows for creating a group.
5. Click **File** > **Save** to save the updated group definition.

Note: You can define filters to select specific computers, file systems, storage subsystems, table spaces, and databases for a reporting group. Tivoli Storage Productivity Center uses these filters to dynamically include resources to the group as they are discovered within your environment. When you define dynamic filters for reporting groups, the filters you enter might or might not be case-sensitive. This sensitivity depends on the database under which the Tivoli Storage Productivity Center database repository is stored. For example:

- The dynamic filters that you enter for reporting groups are case-sensitive when the database repository is stored on Oracle, Sybase, and UDB/DB2.
- The dynamic filters that you enter for reporting groups are not case-sensitive when the database repository is stored on Microsoft SQL Server.

Keep in mind these case-sensitive issues when you define filters for reporting groups. Prevent inaccurate reporting that is caused by unintended resources that are included in reports.

Data Manager reports

Use Data Manager reports to view both overview and detailed information about your storage resources, including computers, disks, file systems, storage subsystems, clusters, and users. These reports are constructed from the statistical information that is gathered by data collection jobs and accumulated in the database repository.

Note: You must run appropriate data collection jobs on the storage resources before you generate reports. If you do not run the data collection jobs before you generate reports, you might see inconsistent or inaccurate results.

Data Manager - Report categories:

Use the report categories within Data Manager to collect information about storage assets, availability, backup requirements, capacity, usage, usage violations, and monitored computer storage space.

The following table describes the categories of reports available in Data Manager.

Report Category	Use these reports to do the following	What job collects information for these reports?
Asset reports	View information about the agents, computers, disk controllers, hard disks, clusters, fabrics, storage subsystems, volumes, file systems, and general hardware inventory of storage assets within your environment. See "Asset reports" on page 1334 for more information about Asset reports.	Probes, discovery Data for System-wide Users and Monitored Directory reports are gathered by scans.
Availability reports	View information about the availability of your storage, both from a network point of view and from a computer up-time perspective. This allows you to see what percentage of the time the storage in your enterprise (or on particular server or group of servers), is off-network due to network problems or perhaps is unavailable due to system downtime. See "Availability Reports" on page 1358 for more information about Availability reports.	Ping data: pings Computer uptime — probes
TPC-wide Storage Space reports	View storage capacity at a file system, computer, or entire network level. For a specific computer, group of computers, or all the computers in your enterprise, you can view the following: • How much storage you have • How much storage is being used • What percentage of storage is free for expansion See "TPC-wide Storage Space Reports" on page 1360 for more information about Capacity reports.	Probes
Usage reports	View information about the usage and growth of the consumption of your storage. Use these reports to do the following: • View which servers and file systems are experiencing the heaviest (or lightest) load of data and storage access. • Identify wasted space by pinpointing files that are no longer needed or have not been accessed for the longest time. • Associate users with the amount of storage they are consuming on specific file systems, groups of file systems, specific computers, groups of computers, and throughout the entire network. • Select the files against which to run an IBM Tivoli Storage Manager archive/backup job. Tivoli Storage Manager protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage. See "Usage Reports" on page 1420 for more information about Usage reports.	Scans

Report Category	Use these reports to do the following	What job collects information for these reports?
Usage Violation reports	View information about the storage usage violations committed by users and user groups within your environment. Use these reports to do the following: • Create and enforce corporate storage usage policies. You can report on violations of quotas and constraints. • Select files in Constraint Violation reports against which to run an Tivoli Storage Manager archive or backup job. • Add files from Constraint Violation reports to an existing Tivoli Storage Manager archive or backup job. See "Usage Violation reports" on page 1427 for more information about Usage Violation reports.	Scans, Quotas, Constraints
Backup reports	Investigate the sizes of data that would be required to backup to achieve full backups, determine the amount of space required to perform an incremental backup of any or all of your data at any given point in time, and select files for archive/backup jobs. See "Backup reports" on page 1429 for more information about Usage Violation reports.	Scans
Monitored Computer Storage Space reports	Use these host-based reports to view the data that is collected by probes for computers, Network Attached Storage, and hypervisors. These reports contain detailed information about the consumption of file system space and enable you to relate that space back to external or internal storage without counting shared space multiple times. See "Monitored Computer Storage Space reports" on page 1431 for more information about Monitored Computer Storage Space reports.	Probes

Use By clauses:

Use the By clauses that are associated with Data Manager reports as a method for viewing the data in those reports from different perspectives.

As you drill down through Data Manager reports, the options repeat for each report type. Use the *By Clauses* options to generate variations of the reports with different perspectives and groupings of data.

The available By Clauses are described in the following table:

By Clause	Description
`By Directory`	Use this clause to select the directories that you want to include in a report.
`By Directory Group`	Use this clause to select the directory groups that you want to include in a report.

By Clause	Description
By Disk	Use this clause to select the disks that you want to include in a report.
By File system	Use this clause to select the file systems that you want to include in a report.
By File system Group	Use this clause to select the file system groups that you want to include in a report.
By Cluster	Use this clause to select the clusters or cluster resource groups that you want to include in a report. Cluster resource groups in a PowerHA SystemMirror for AIX or MSCS-clustered environment are listed as computers.
By Computer	Use this clause to select the computers that you want to include in a report.
By Computer Group	Use this clause to select the computer groups that you want to include in a report.
By Storage Subsystem	Use this clause to select the storage subsystems that you want to include in a report.
By Domain	Use this clause to select the Window Domains that you want to include in a report.
Network Wide	The reports and charts that you see are on all data, network wide. Use this view to generate a global view of the storage and storage usage for your whole enterprise.

Asset reports:

Use asset reports to view itemized statistics for your storage resources. Tivoli Storage Productivity Center provides a hierarchical view of your storage resources so that you can navigate to view your resources in greater detail.

The following section is a description of each of the report types that are under **Data Manager** > **Reporting** > **Asset**.

Asset Reporting Type	Description
By Cluster	View storage assets based on PowerHA SystemMirror for AIX or MSCS cluster membership. You can navigate to a cluster and view its nodes and cluster resource groups. For each cluster resource group, you can navigate to show its controllers, disks, file systems, logical volumes, and exports or shares. For a cluster node, you can view local resources.
By Computer	Navigate through storage assets in a hierarchical manner. You can navigate to a computer and show its controllers, disks, file systems, logical volumes, and directories. You can then navigate further to view more detailed information such as the disks that are under a specific controller or file systems that are under a disk.
By Hypervisor	Navigate through hypervisors in a hierarchical manner. You can use this report to navigate to and view details about the virtual machines, controllers, disks and file systems, or logical volumes.
By IBM Storwize V7000 Unified/IBM SONAS	Navigate through Storwize V7000 Unified or IBM SONAS clusters and directories in a hierarchical manner. You can use this report to navigate to and view details about the nodes, pools, Network Shared Disks, file systems, monitored directories, and exports that are associated with the clusters.

Asset Reporting Type	Description
By OS Type	View storage assets based upon the operating systems that are running on those storage servers. This report type provides a way to see what your operating system dependencies are throughout your organization. By navigating to any or each of these different supported operating system types, you can see how many of each server type you have and you can navigate to investigate the underlying storage.
By Storage Subsystem	View storage capacity at a computer, file system, storage system, volume, and disk level. These reports also show the relationships between the components of a storage system.
System Wide	View storage assets as a whole, across the entire system. Use these reports to view lists of all of your disks, computers, file systems, Storage Resource agents and Data agents that are installed in your environment.

By Cluster:
Use the "By Cluster" reports to navigate to a cluster to view its nodes and cluster resource groups.

- For a node, you can view its local resources, including controllers, disks, file systems, logical volumes, exports, shares, and directories.
- For a cluster resource group, you can view controllers, disks, file systems, logical volumes, exports, and shares.
- For a node, the number of processors that are reported are the logical processors. The amount of RAM that is reported is the available RAM.

1. Expand **Asset → By Cluster** node to display the nodes for each cluster that has been discovered and monitored by Data Manager.
2. Click a node to view cluster information, or further expand the cluster node to navigate to greater detail.

Cluster Information
Click the name of a cluster that appears under **Asset → By Cluster** to view information about that cluster such as:

- Name
- Domain
- IP address
- Discovered Time

Controller Information
As you navigate to Asset Reporting By Cluster for a cluster node or cluster resource group, you can expand the assets under a controller, or you can select a controller to view its detailed information.

The name of the controller's computer is shown, as well as the driver name, instance number, bus number, address, and type of driver (for example, SCSI, IDE, etc).

Note: Host Bus Adapter (HBA) controllers must be connected to the SAN and see volumes in order to be displayed under the Controllers node.

Disk Information
As you navigate to Asset Reporting By Cluster for a cluster node or cluster resource group, you can expand the assets under a disk, or you can select a disk to view its detailed information.

The disk detail information is shown in a tab dialog that is made of up three pages: **General**, **Paths**, and **Latest Probe**.

General page

Use the General page to view the computer name, the path name, SCSI Target ID, logical unit number, and the number of access paths. The disk information shown includes the hardware manufacturer of the disk, the model number of the disk, the firmware revision, the disk serial number, and the manufacture date along with the time when Data Manager discovered the disk. The page also includes a breakdown of how the disk is allocated, including the disk allocation point and the length of the allocation.

Latest Probe page

Use the Latest Probe page to view the information that was obtained by Data Manager about this disk during the most recent probe of that disk. This information on this page includes:

- Sector size and sector information
- the number of heads
- the number of cylinders
- the logical block size
- the capacity of the disk
- the RPM or revolutions per minute that the disk spins at
- the elapsed time the disk has been powered on
- spare blocks
- the current write protection setting
- the cache enabled setting
- whether failure is predicted for this disk
- the time and date of the last Probe
- the unallocated space
- the number of primary defects
- the number of grown defects
- the NonMedium errors
- the recorded operating temperature

Some of these values are manufacturer specific, so if they are unavailable, then you will see **N/A** for a value. Additional information about the breakdown of recovered errors, unrecovered errors, and the number of bytes broken down by Write, Read, and Verify is shown in this report. Also included are three SCSI disk bit settings that you can turn on/off and the change will actually be made to the disk settings the next time a probe is run on that disk. These include on/off settings for:

- Automatic read reallocation
- Automatic write reallocation
- Read continuous (no error correction)

Probe History page

Use the Probe History page to view the history of probes that have been run on the disk you are reporting on. Each row on this report represents one time that a probe has been run against the disk. Each row shows what time the probe ran, the type of record (which could be daily, weekly or monthly), bytes written, bytes read, number of grown defects.

Filesystem and Logical Volume Information

The Asset Reporting By Cluster Filesystem or Logical Volume information report contains two pages: **Filesystem** and **Logical Volume**.

Filesystem page

Use this page to view a pie chart that depicts used space and free space on the file system and includes the following information:

- Computer name
- File system type
- Use count
- Mount point
- Physical size of the file system
- Capacity of the file system
- The time and date of the last probe and scan
- When the file system was discovered on the system by Data Manager
- If the file system has been removed then a removed date/time will be shown

Logical Volume page

Use this page to view:

- The path name of the logical volume
- The capacity of the logical volume
- The use count
- The type of logical volume
- Whether it is a swap volume and mirrored
- A tree displaying the mapping of the logical volume on the disk area(s) that make up the logical volume

Exports or Shares Information

The Asset Reporting By Computer Exports or Shares information report includes the following information:

- the name of the computer the share is on
- the export name that it will be known by
- the path on the computer the share resides upon
- the protocol being used and the date and time that Data Manager discovered the export or share

Monitored Directory Information

To view a report on a directory that you are monitoring, expand the **Monitored Directories** node on the navigation tree and select the directory on which you want to report. This information is available for clustered nodes, but not cluster resource groups.

The Asset Reporting By Cluster Monitored Directory report includes the following information:

- computer name the directory resides on
- name of the directory
- owner
- OS Group
- total size of the directory
- number of files in the directory
- count of subdirectories in the directory
- average file size
- access time
- modification time and creation time of the directory
- last scan time and the time that Data Manager discovered the directory
- file attributes that are set
- list of the Directory Groups that the directory is a part of

By Computer:
Use the "By Computer" reports to navigate to a computer and view detailed information about its controllers, disks, file systems, logical volumes, and directories.

Computer Information

As you navigate to Asset Reporting By Computer, you can expand the assets under a computer, or you can select a computer to view the details for that computer. The detailed report shows the following information:

- Computer
- Host ID: the identifier of the host computer in hexadecimal format. Currently, IBM Tivoli Storage Productivity Center collects Host ID information for Oracle Solaris computers only. All other platforms display GUID information in the Host ID column.
- Group and Domain
- Network Address, IP address
- Time zone
- Hardware manufacturer and model numbers
- Processor type
- Count and speed
- Amount of RAM
- Operating System (OS) information such as:
  - Type of OS and version
  - The storage attributes of the server
  - The computer disk capacity, unallocated space, and free space

Controller Information

As you navigate to Asset Reporting By Computer, you can expand the assets under a controller, or you can select a controller to view the details for that controller. The name of the computer controller is shown, as well as the driver name, instance number, bus number, address, and type of driver (for example, SCSI, IDE, and so on).

Note: Host Bus Adapter (HBA) controllers must be connected to the SAN and see volumes in order to be displayed under the Controllers subnode.

Disk Information

As you navigate to Asset Reporting By Computer, you can expand the assets under a disk, or you can select a disk to view the details for that disk. The disk detail information is shown in a tab dialog that is made of up four tabs: **General**, **Paths**, **Latest Probe**, and **Probe History**.

General page

The General page includes the computer name, the path name, SCSI Target ID, logical unit number, the number of access paths.

The disk information shown includes:

- The hardware manufacturer of the disk
- The model number of the disk
- The disk firmware revision
- The disk serial number
- The disk manufacture date
- Indicators that specify whether the disk is encrypted or encryptable

- The time when Data Manager discovered the disk

The tab also includes a breakdown of how the disk is allocated, including the disk allocation point and the length of the allocation.

Paths The **Paths** page shows the following information:

- Host
- OS type
- Path
- Controller
- Instance
- Bus number
- SCSI target id
- Logical unit number

Latest Probe page

The **Latest Probe** page shows the information that was obtained by Data Manager about this disk during the most recent probe of that disk. This information on this tab includes:

- Sector size and sector information
- The number of heads
- The number of cylinders
- The logical block size
- The capacity of the disk
- The RPM or revolutions per minute that the disk spins at
- The elapsed time the disk has been powered on
- Spare blocks
- The current write protection setting
- The cache enabled setting
- Whether failure is predicted for this disk
- The time and date of the last Probe
- The unallocated space
- The number of primary defects
- The number of grown defects
- The NonMedium errors
- The recorded operating temperature

Some of these values are manufacturer-specific. If the values are unavailable, **N/A** is displayed.

Additional information about the breakdown of recovered errors, unrecovered errors, and the number of bytes broken down by Write, Read, and Verify is shown in this report.

Also included are three SCSI disk bit settings that you can turn on/off and the change will actually be made to the disk settings the next time a probe is run on that disk. These include on/off settings for:

- Automatic read reallocation
- Automatic write reallocation
- Read continuous (no error correction)

Probe History page

The **Probe History** page shows the history of probes that have been run on the disk you are reporting on. Each row on this report represents one time that a probe has been run against the disk. Each row shows what time the probe ran, the type of record (which could be daily, weekly, or monthly), bytes written, bytes read, number of grown defects.

Volume Group Information

Data Manager provides support for the following volume managers:

- Veritas Volume Manager (Oracle Solaris)
- Logical Volume Manager (AIX)

These volume managers allow you to group logical volumes and disks. As you navigate to Asset Reporting By Computer for an AIX or Oracle Solaris computer, you can expand the assets under a volume group to see the group disks and file systems or logical volumes.

You can also select a volume group to view detailed information about it. The group name is shown, as well as the type (LVM or Veritas), capacity, free space, number of volumes, number of disks, discovered time, and probe time.

File System / Logical Volume Information

The Asset Reporting By Computer File System / Logical Volume information report contains two pages: **File System** and **Logical Volume**.

File System page

This page shows a pie chart depicting used space and free space on the file system and includes the following information:

- Computer name
- File system type
- Use count
- Mount point
- Physical size of the file system
- Capacity of the file system
- The time and date of the last probe and scan
- When the file system was discovered on the system by Data Manager
- If the file system has been removed, then a removed date/time is shown

Also shown for UNIX-based systems are:

- The maximum file count
- Used inodes
- Free inodes. Inodes are areas reserved for index entries within the UNIX file systems. For Windows computers, these values are zero.
- The amount of used space and free space is shown

Logical Volume page

The information on this page includes:

- The path name of the logical volume
- The capacity of the logical volume
- The use count
- The type of logical volume
- Whether it is a swap volume and mirrored
- A tree displaying the mapping of the logical volume on the disk areas that make up the logical volume

Exports/Shares Information

The Asset Reporting By Computer Exports/Shares information report includes the following information:

- The name of the computer the share is on
- The export name that it will be known by

- The path on the computer the share resides upon
- The protocol being used and the date/time that Data Manager discovered the export/share

Monitored Directory Information

To view a report on a directory that you are monitoring, expand the **Monitored Directories** node on the navigation tree and select the directory on which you want to report.

Note: This information is not shown for cluster resource groups.
The Asset Reporting By Computer Monitored Directory report includes the following information:

- Computer name the directory resides on
- Name of the directory
- Owner
- OS Group
- Total size of the directory
- Number of files in the directory
- Count of subdirectories in the directory
- Average file size
- Access time
- Modification time and creation time of the directory
- Last scan time and the time that Data Manager discovered the directory
- File attributes that are set
- List of the Directory Groups that the directory is a part of

The Asset Reporting By Computer Monitored Directory report includes the following information:

- Computer name the directory resides on
- Name of the directory
- Owner
- OS Group
- Totalsize of the directory
- Number of files in the directory
- Count of subdirectories in the directory
- Average file size
- Access time
- Modification time and creation time of the directory
- Last scan time and the time that Data Manager discovered the directory
- File attributes that are set
- List of the Directory Groups that the directory is a part of

By IBM Storwize V7000 Unified/IBM SONAS:

Use "By IBM Storwize V7000 Unified/IBM SONAS" reports to navigate to the components of a Storwize V7000 Unified storage system or an IBM SONAS system. These components include the clusters, nodes, pools, Network Shared Disks, file systems, and exports that are associated with the storage system.

To view Storwize V7000 Unified or IBM SONAS asset reports, follow these steps:

1. Expand **Data Manager** > **Reporting** > **Asset** > **By IBM Storwize V7000 Unified/IBM SONAS**. The Storwize V7000 Unified storage systems or IBM SONAS systems that were discovered by Tivoli Storage Productivity Center are displayed.
2. Expand a storage system to view detailed information. The following tabbed pages are displayed:

Storage Subsystem Details
Use this page to view detailed information about the Storwize V7000 Unified storage system or the IBM SONAS system. The information available for the storage system includes the status, vendor, type and model, and the available and used space. You can also use this page to modify some information about the Storwize V7000 Unified storage system or the IBM SONAS system.

File Module Details
Use this page to view information about Storwize V7000 Unified.

You can view information about the components of the Storwize V7000 Unified storage system or the IBM SONAS system. When you click a component, report pages for that component are displayed.

Storwize V7000 Unified or IBM SONAS - Cluster Information:

This page shows general information about a selected cluster that is associated with a Storwize V7000 Unified storage system or an IBM SONAS system.

The following information is displayed on this page.

Table 115. Information about clusters associated with a Storwize V7000 Unified storage system or an IBM SONAS system

Field Name	Description
Cluster	The name of the cluster. For Storwize V7000 Unified or IBM SONAS clusters, the storage system name is included with the cluster name in the Computer or Device column. The format of the column is `storagesystem_name: cluster_name`, where storagesystem_name is the name of the Storwize V7000 Unified or IBM SONAS device that the cluster belongs to. For example, `tpcserver1.tpc.example.com: kq98n5d.acct` displays storage system tpcserver1.tpc.example.com that cluster kq98n5d.acct belongs to.
Cluster ID	The identifier of the cluster.
Type	The cluster can be one or more of the following types: **Interface** An interface cluster contains nodes that connect a system to an IP network for file-serving capabilities by using service protocols. **Management** A management cluster contains nodes that are used to configure, administer, and monitor a system. **Storage** A storage cluster contains nodes that provide the storage and file system for the data in a system.

Storwize V7000 Unified or IBM SONAS - Node Information:

This page shows general information about a selected node that is associated with a Storwize V7000 Unified storage system or an IBM SONAS system.

The following information is displayed on this page.

Table 116. Information about nodes associated with a Storwize V7000 Unified storage system or an IBM SONAS system

Field Name	Description
Node	The name of the node.
Description	The description of the node.
IP Address	The IP address for the node.
Role	The role of the node. The node can serve one or more of the following roles: **Interface** An interface node connects a system to an IP network for file-serving capabilities by using service protocols. **Management** A management node is used to configure, administer, and monitor a system. **Storage** A storage node provides the storage and file system for the data in a system.
Status	The status of the node. The status can be one of the following values: **Normal** The node is operating normally. **Critical** The node is either not operating or an error occurred. **Warning** At least one component of the node is not operating or an error occurred. **Unreachable** Tivoli Storage Productivity Center cannot reach the node.

Storwize V7000 Unified or IBM SONAS - Pool Information:

This page shows general information about a selected pool that is associated with a Storwize V7000 Unified storage system or an IBM SONAS system.

The following information is displayed on this page.

Table 117. Information about pools associated with a Storwize V7000 Unified storage system or an IBM SONAS system

Field Name	Description
Pool	The name of the pool.
Capacity	The capacity of the pool.
Available Space	The amount of unused space that is in the pool.

Storwize V7000 Unified or IBM SONAS - NSD Information:

This page shows general information about a selected Network Shared Disk (NSD) that is associated with a Storwize V7000 Unified storage system or an IBM SONAS system.

The following information is displayed on this page.

Table 118. Information about a selected NSD associated with a Storwize V7000 Unified storage system or an IBM SONAS system

Field Name	Description
NSD	The name of the NSD.
NSD ID	The identifier of the NSD.
Failover Group	The failure group that the NSD belongs to.
Type	The NSD type. The possible values are: • Metadata and Data • Metadata Only • Data Only
Pool	The pool that the NSD belongs to.
Status	The status of the NSD. The possible values are: **Normal** The NSD is operating normally. **Critical** The NSD is either not operating or an error occurred. **Warning** At least one component of the NSD is not operating or an error occurred. **Unreachable** Tivoli Storage Productivity Center cannot reach the NSD.
Disk Space	The total disk space for the NSD.
Available Space	The amount of unused disk space for the NSD.

Storwize V7000 Unified - NSD Probe History:

This page shows the probe history for a selected Network Shared Disk (NSD) that is associated with a Storwize V7000 Unified storage system.

The following information is displayed on this page.

Table 119. Information about a selected NSD associated with a Storwize V7000 Unified storage system

Field Name	Description
Time of Probe	The time that the probe ran.
Record Type	The record type that is shown in the report. Valid values are: • Daily • Weekly • Monthly
Bytes Written	The total number of bytes that were written since the NSD was formatted or manufactured.
Bytes Read	The total number of bytes that were read since the NSD was formatted or manufactured.
Grown Defects	The number of defects that were detected on the NSD after the NSD was used.
Recovered Write Errors	The number of write errors that were successful on a subsequent attempt.
Unrecovered Write Errors	The number of write errors that were not successful on a subsequent attempt.
Recovered Read Errors	The number of read errors that were successful on a subsequent attempt.
Unrecovered Read Errors	The number of read errors that were not successful on a subsequent attempt.
Recovered Verify Errors	The number of verify errors that were successful on a subsequent attempt.
Unrecovered Verify Errors	The number of verify errors that were not successful on a subsequent attempt.
Non-Medium Errors	The number of errors that are not related to the physical disk platter, such as a mechanical problem in the head movement or communication errors.
Available Disk Space	The amount of unused space on the NSD.
Primary Defects	The number of defects on the NSD when the NSD was new.
Temperature	The temperature of the NSD in degrees Celsius.

Storwize V7000 Unified - Correlated Storage Subsystem Volume:

This page shows information about the relationship between a selected Network Shared Disk (NSD) and the associated virtual disk or volume that is on the Storwize V7000 Unified storage system.

The following information is displayed on this page.

Table 120. Information about the relationship between a selected NSD and an associated virtual disk or volume on a Storwize V7000 Unified storage system

Field Name	Description
Virtual Disk Name	The name of the virtual disk.
Virtual Disk Space	The total storage capacity of a virtual disk.
Unavailable Virtual Disk Space	The amount of storage space for Redundant Array of Independent Disks (RAID) or mirror redundancy.
Type	The organization of the virtual disk. Values are: • Sequential • Striped
RAID Level	The RAID level of the virtual disk. For example, 0, 1, 5, and so on.
Is Encryptable	Indicates whether the virtual disk is encryptable. If the virtual disk is encryptable, the virtual disk uses full disk encryption, but does not belong to a defined System Storage DS8000 or System Storage DS5000 encryption group. Values are: • True (the virtual disk is encryptable) • False (the virtual disk is not encryptable)
Is Encrypted	Indicates whether the virtual disk is encrypted. If the virtual disk is encrypted, the volume uses full disk encryption and belongs to a defined DS8000 or DS5000 encryption group. Values are: • True (the virtual disk is encrypted) • False (the virtual disk is not encrypted)
Easy Tier	Indicates whether Easy Tier is on or off for the virtual disk.
Easy Tier Status	Indicates the Easy Tier functions that are active for the virtual disk copy. Values are: • Active (can move extents of this virtual disk copy for performance [automatic data placement]) • Inactive (no Easy Tier function is active) • Measured (statistics are being gathered for the virtual disk copy, but no extents are moved)
Tiers	The tier levels that are used by the virtual disk. Each tier is displayed as a row in a table.
Flash Copy	Indicates whether the virtual disk or copy is in a FlashCopy relationship and whether it is a FlashCopy source or a FlashCopy target. A virtual disk or copy that is not in a FlashCopy relationship is displayed with a None value, regardless of whether the virtual disk is on a system that supports FlashCopy.

Table 120. Information about the relationship between a selected NSD and an associated virtual disk or volume on a Storwize V7000 Unified storage system (continued)

Field Name	Description
Attributes	Indicates whether the disk is surfaced or not surfaced.

Storwize V7000 Unified or IBM SONAS - File System Information:

This page shows general information about a selected file system that is associated with a Storwize V7000 Unified storage system or an IBM SONAS system.

The following information is displayed on this page.

Table 121. Information about a selected file system associated with a Storwize V7000 Unified storage system or an IBM SONAS system

Field Name	Description
Computer	The name of the computer that contains the file system.
File System Type	The file system type, which is General Parallel File System (GPFS).
Mount Point	The name or mount point of the file system (for example, `/ibm/gpfs0`, or `/ibm/testfs`).
Physical Space	The physical size of the file system.
File System Space	The capacity of the file system.
Probe Last Run	The date and time that the last probe ran for the file system.
Scan Last Run	The time that the last scan ran for the file system.
Discovered Time	The date and time that the file system was discovered.
Removed Time	The date and time that the file system was removed, if applicable.
Maximum File Count	The maximum number of files allowed on the file system.
Used I-Nodes	The number of used i-nodes in the file system.
Free I-Nodes	The number of free i-nodes within a file system.
Used Space	The amount of used storage space within a file system. The amount of used and free space is also shown in a chart on the page.
Free Space	The amount of unused storage space in the file system.

Storwize V7000 Unified or IBM SONAS - File Set Information:

This page shows general information about a selected file set that is associated with a Storwize V7000 Unified storage system or an IBM SONAS system.

The following information is displayed on this page.

Table 122. Information about a selected file set associated with a Storwize V7000 Unified storage system or an IBM SONAS system

Field Name	Description
File Set	The name of the file set.
Path	The path for the file set. The path is displayed only if Linked is displayed in the **Status** field.
File System	The name of the file system that is associated with the file set.
Status	The status of the file system. Valid values are: **Linked** The file set is linked to the file system that is displayed in the **File System** field. **Unlinked** The file set is not linked to the file system that is displayed in the **File System** field.
Used Space	The amount of space that is used by the file set.
Used I-Nodes	The number of i-nodes that are used by the file set.

Storwize V7000 Unified or IBM SONAS - Export Information:

This page shows general information about a selected export that is associated with a Storwize V7000 Unified storage system or an IBM SONAS system.

The following information is displayed on this page.

Table 123. Information about a selected export associated with a Storwize V7000 Unified storage system or an IBM SONAS system

Field Name	Description
Export or Share	The name of the export.
Cluster	The name of the cluster that the export is associated with.
Path	The name of the path for the export.
Discovered	The date and time that the export was discovered.
Active	Indicates whether the export is active or inactive within the cluster. Valid values are: **Yes** The export is active. **No** The export is inactive. If an export is inactive, it is included in the list of exports, but the data that is associated with the export cannot be accessed.

Storwize V7000 Unified or IBM SONAS - Configured Protocol Information:

This page shows the protocols for a selected export that is associated with a Storwize V7000 Unified storage system or an IBM SONAS system.

Protocols are specified for an export when the export is created. The following protocols can be specified for an export:

- HTTP
- FTP
- Secure Copy Protocol (SCP)
- Network File System (NFS)
- Common Internet File System (CIFS)

For each protocol that is listed on this page, **Configured** or **Not Configured** is shown to designate whether the protocol as been specified for the export. In addition, the information in the following tables is provided for the NFS and the CIFS protocols when they are specified for the export.

Protocol Type NFS

Table 124. NFS protocols for a selected export associated with a Storwize V7000 Unified storage system or an IBM SONAS system

Field Name	Description
Client Name	The name of the client that can access the NFS protocol. If an asterisk (*) is shown, the settings apply to multiple clients.
Read Only	Specifies whether the export is read only. If the export is read only, files cannot be modified or created on this export. Values are: **Yes** Files cannot be modified or created. **No** Files can be modified or created.
Root Squash	Specifies whether root users on clients should be prevented from having privileged access to exported files. Values are: **Yes** Root squash is enabled. **No** Root squash is not enabled.
Secure	Specifies whether NFS access to the export is secure. Values are: **Yes** NFS is secure. **No** NFS is not secure.

Protocol Type CIFS

Table 125. CIFS protocols for a selected export associated with a Storwize V7000 Unified storage system or an IBM SONAS system

Field Name	Description
Read Only	Specifies whether the export is read only. If the export is read only, files cannot be modified or created on this export. Values are: **Yes** Files cannot be modified or created. **No** Files can be modified or created.
Browsable by Client	Specifies whether the export is hidden from the client. The export is not shown in the list of available exports. Values are: **Yes** The export is not hidden from the client. **No** The export is hidden from the client.
Hide Objects with No Read Access	Specifies whether the export is excluded from lists of files and directories when the user does not have permission to those files and directories. Values are: **Yes** The export is excluded. **No** The export is included.
Client/Group/SID	The client, group, or security identifier that can access the CIFS protocol.
Access	Specifies whether the client, group, or security identifier has access to the export. • Allowed • Denied
Value	Specifies the level of access that is granted to the client, group, or security identifier. Values are: • Full • Read • Write
Comment	A user-defined comment for the CIFS export. If a comment is not available, this field is blank.

Storwize V7000 Unified or IBM SONAS - Monitored Directory Information:

This page shows general information about a directory that is associated with a Storwize V7000 Unified storage system or an IBM SONAS system.

The following information is displayed on this page.

Table 126. Information about a directory associated with a Storwize V7000 Unified storage system or an IBM SONAS system

Field Name	Description
Computer	The name of the computer that contains the directory.

Table 126. Information about a directory associated with a Storwize V7000 Unified storage system or an IBM SONAS system (continued)

Field Name	Description
Directory	The name of the directory.
Owner	The owner of the directory.
OS Group	The operating system group that the directory belongs to.
Total Size	The total size of the directory.
File Count	The number of files in the directory.
Directory Count	The number of subdirectories in the directory.
Average File Size	The average size of the files in the directory.
Access Time	The date and time that the directory was last accessed.
Modification Time	The date and time that the directory was last modified.
Creation Time	The date and time that the directory was created.
Scan Time	The date and time that the directory was last scanned.
Discovered Time	The date and time that the directory was discovered.
Removed Time	The date and time that the directory was removed.
File attributes that are set	If any attributes were set for the files in the directory (for example, read only), the attributes are displayed. If no attributes were set, `No File Attributes Are Set` is displayed.
This Directory belongs to the following Directory Groups	A list of the directory groups that the directory belongs to.

By OS Type:
Use the "By OS Type Asset" report to view your storage assets based upon the operating systems under which they are running. This provides you with a an easy way to see what your OS dependencies are throughout your organization. You can navigate to any of the storage resources that appear under an operating system node to access all of the same sub-reports that as described for other asset reports.

By Storage Subsystem:

Browse your storage subsystem assets in a hierarchical manner. Data Manager provides reports for individual storage subsystems and any disk array subsystems whose SMI-S Providers are CTP certified by SNIA for SMI-S 1.0.2 or later.

For disk array systems, you can view information about:

- Disk groups
- Array sites
- Ranks
- Storage pools

- Disks
- Volumes

For virtualization systems, you can view information about:

- SAN Volume Controller clusters
- Storwize V7000
- Storwize V7000 Unified
- Managed disk groups
- Managed disks
- Virtual disks
- Virtual disk copies

Keep in mind the following considerations:

- A volume name in a report that is marked with an asterisk (*) indicates that the volume is a space-efficient volume. The * is not part of the actual volume name.
- The total capacity value for storage pools in TotalStorage Enterprise Storage Server subsystems includes only the formatted (usable) capacity for storage pools.
- For the following storage systems, you can navigate from storage subsystems and show managed disks, managed disk groups, and virtual disks:
  - SAN Volume Controller
  - Storwize V7000
  - Storwize V7000 Unified

  For Storwize V7000 and Storwize V7000 Unified, you can also navigate from the internal disks node to view the internal storage on the system.
  - Information that is displayed for managed disks and internal disks is the same as for storage subsystem disks. See "By Storage Subsystem" on page 1351 for a description of the available views.
  - Information that is displayed for virtual disks is the same as for storage subsystem volumes. See "By Storage Subsystem" on page 1351 for a description of the available views.
- For IBM FAStT, HP StorageWorks, and the XIV system, you can navigate to storage subsystems and show storage pools, disks, and volumes.

  Note: The total capacity value for storage pools in these storage subsystems includes raw (unconfigured) capacity—it does not include formatted/usable capacity.
- For Hitachi Data Systems TagmaStore subsystem, the relationship between a disk or volume is not displayed if the volume is using imported extents (virtual disks).

  Note: The total capacity value for storage pools in HDS subsystems includes the formatted/usable capacity for storage pools only.
- Some of the columns in the totals row for these reports can count storage multiple times if the report includes the following:
  - A storage virtualizer like a SAN Volume Controller
  - The disk arrays that host the backend storage to a storage virtualizer like a SAN Volume Controller

  In this situation, disks and logical unit numbers (LUNs) appear to belong to both a storage virtualizer like SAN Volume Controller and the disk arrays, and

are counted multiple times in some of the columns in the totals row. To avoid this situation, filter the report to display only the storage virtualizer or the disk arrays. If this situation is unavoidable, then some of the columns in the totals row can be adjusted as follows:

1. Use the **Physical Disk Space** column to determine the actual physical disk space on all arrays. This disk space is not counted multiple times.
2. Sum the **Disk Space** column for the disk arrays and subtract this result from the totals row to determine the total backend storage for a virtualizer.
3. To prevents counting the backend volumes multiple times, subtract the **Backend Volume Space** column from the **Volume Space** column in the totals row to determine the amount of volume space.

See "Report considerations and limitations" on page 1307 for additional information about reports.

Disk Groups

When you expand the **Disk Group** node, the navigation tree shows a list of disk groups for the storage subsystem. You can expand the name of a disk group to view the disks and volume spaces within that disk group. Navigate through both Volume Spaces and Disks subnodes to view:

- Disks and volumes associated with a volume space
- Volumes with segments on a disk
- Disks containing segments from a volume
- Detailed information about a disk

Array Sites

(IBM Tivoli Storage DS8000 only) When you expand the **Array Site** node, the navigation tree shows a list of array sites for the storage subsystem. Click **Array Site** to view a description of the information displayed about array sites. You can expand the name of an array site to view the ranks, storage pools, and disks associated with that array site.

Storage Pools

When you expand the **Storage Pool** node, the navigation tree shows a list of storage pools for the storage subsystem. You can expand the name of a storage pool to view the disks and volumes within that storage pool, including:

- Volumes with segments on a disk
- Disks containing segments from a volume
- Detailed information about a disk

Disks When you expand the **Disks** node, the navigation tree shows a list of disks for the storage subsystem. You can expand the name of disk to view the volumes associated with that disk. Click the name of a volume associated with a disk to view detailed information about that volume.

Volume Spaces

When you expand the **Volume Spaces** node, the navigation tree shows a list of volume spaces for the storage subsystem. You can expand the name of a volume space to view the disks and volumes within that volume space. Navigate through both the Disks and volumes subnodes to view:

- Volumes with segments on a disk
- Disks containing segments from a volume
- Detailed information about a disk
- Detailed information about a volume

Volumes

When you expand the volumes node, the navigation tree shows a list of volumes for the storage subsystem. Click the name of a volume to view detailed information about that volume.

Note:

- Volume Reporting is not available for IBM FAStT FlashCopy logical drive and repository logical drive.
- After creating a volume on a storage subsystem, you must run a probe on that storage subsystem to collect accurate information about disk allocation.

Managed Disk Groups

When you expand the Managed Disk Group node, the navigation tree shows a list of disk groups that are managed by the following storage systems:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You can expand the name of a disk group to view the physical disks and virtual disks that are associated with it.

Managed Disks

When you expand the Managed Disk node, the navigation tree shows a list of disks that are managed by the following storage systems:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You can expand the name of a disk to view the virtual disks and any available internal disks that are associated with it.

Virtual Disks and Virtual Disk Copies

There are three nodes in the navigation tree that are related to virtual disks (VDisks). These nodes might be included under the nodes for the SAN Volume Controllers in your environment:

- **Virtual Disks** (at the same level as **Managed Disk** and **Managed Disk Groups**): Expand this node to view a list of the logical disks that a cluster provides to the SAN. You can further expand each virtual disk (VDisk) to view its mirror copies (primary copies and additional copies).
- **Virtual Disks** (under **Managed Disk** and **Managed Disk Groups**): Expand this node to view the VDisks for which the primary copies are on the selected Managed Disk.
- **Virtual Disk Copies** (under **Managed Disk** and **Managed Disk Groups**): Expand this node to view all the virtual disk copies that are on the selected Managed Disk. The primary copies are not listed under this node because they are displayed under the **Virtual Disks** node.

  Tip:

  – The **Virtual Disk Copies** node is displayed in the navigation tree only if a monitored SAN Volume Controller supports virtual disk mirrors. If it supports virtual disk mirrors, all virtual disks will have a related primary virtual disk copy regardless of whether the virtual disk is mirrored. This is because all virtual disks on a SAN Volume

Controller that supports virtual disk mirrors are primary copies that can be updated to have mirrored disks or virtual disk copies.

– The icons for all virtual disk copies are shown as folders in the navigation tree because they represent logical extensions of a virtual disk.

- If the virtual disk is a virtual disk copy, the primary virtual disk copy is noted by -0 in the name of the logical volume, and all secondary copies are noted by a hyphen and a value larger than zero. If the virtual disk is not a virtual disk copy, only one logical volume is created and is noted by -0. The logical disk representation of a virtual disk should always exist, but if a probe is started for SAN Volume Controller 4.3.1 or later, the logical disk representation of the virtual disk expands to one or more logical volumes where virtual disk copies are represented by multiple logical volumes, and non-virtual disk copies are represented by a single logical volume.

Internal Disks

Internal storage information about the Storwize V7000 or Storwize V7000 Unified systems in your environment is displayed in the **Internal Disks** nodes in the navigation tree:

Internal Disks (at the same level as Virtual Disks, Managed Disk, and Managed Disk Groups)

Expand this node to view a list of the internal disks that a Storwize V7000 or Storwize V7000 Unified system provides to the SAN.

Internal Disks (under Managed Disk Groups > Managed Disk)

The internal disks for this node are only displayed if there are internal disks that are available on the managed disk.

Internal Disks (under Managed Disk)

The internal disks for this node are only displayed if there are internal disks available.

System-wide:

Use "System-wide" reports to do the following tasks.

- View your storage assets as a whole across the entire system. You can view lists of all of your disks, computers, storage subsystems, disk/volume groups, disks, file systems, logical volumes, storage subsystem volumes, exports or shares, monitored directories, users, and OS users. Additionally, you can view aggregated reports that are based on certain filters, such as grown defects per disk.
- View detailed information about Storage Resource agents and Data agents. Use these reports to monitor your agents and track their status.
- The **Selection** button is disabled for system-wide reports because all the monitored storage entities are automatically displayed for each report type. For example, all the computers that are monitored by IBM Tivoli Storage Productivity Center are displayed in the **System-wide** > **Computers** report.

Note: A volume name in a report that is marked with an asterisk (*) indicates that the volume is a space-efficient volume. The * is not part of the actual volume name. For more information about reports, see "Report considerations and limitations" on page 1307.

System-wide reports include options to go to storage resources for more detail, view charts of current statistics, and generate historical charts to show storage trends. The following icons are available on some system-wide reports.

Chart icon
Click to generate charts for the selected rows in the table. You can generate a chart that is based on the statistics that are currently displayed in the report. You can also generate a history chart that provides trending data for the selected storage subsystems. For more information about history charts, see Viewing historical trends.

Magnifying glass icon
Click to view detailed information about a specific file system or logical volume in a row.

Agents reports
Use these reports to track and monitor the Storage Resource agents and Data agents that are deployed in your environment. With this report, you can export a comma-separated file that contains the name and status of an agent. Specifically, use this report to run the following tasks.

- View information about the agents that are deployed in your environment
- Perform actions against one or more agents that are shown in the report

Each agent that is shown in the report has a magnifying glass icon next to it. Click this icon to display detailed information about an agent (which is identical to the information on the **Administrative Services** > **Data Sources** > **Data/Storage Resource agents** node in the navigation tree).

You can also select one or more rows on the report and right-click to display a menu. Use the options in this menu to run an action against multiple agents at the same time. The menu contains all the other operations available from the **Administrative Services** > **Data Sources** > **Data/Storage Resource agents** node in the navigation tree.

Computers reports

Use these reports to view all the computers in the system, which are sorted in order of the time they were last started, probed, or discovered. This information is valuable in environments where computers become more unreliable after they are up and running for long periods of time.

Storage Subsystem reports
Use these reports to view all your monitored storage subsystems and SAN Volume Controllers, sorted by their total capacity (By Disk Space) or alias (By Storage Subsystem).

Unmanaged Virtual Machines reports
Use these reports to view detailed information about the unmanaged virtual machines in your environment. Each row in the report gives you information for each unmanaged virtual machine.

Virtual Machines With No VMware reports
Use these reports to view detailed information about the virtual machines that have Storage Resource agents or Data agents that are installed but are not discovered through the VMWare data sources. Each row in the report gives you information for each virtual machine.

Disk/Volume Groups reports
Data Manager provides support for the following volume managers.

- Veritas Volume Manager (Oracle Solaris)
- Logical Volume Manager (AIX)

Through these volume managers, you can create groups of logical volumes and disks. Use these reports to view all your disk and volume groups, which are sorted by their computer (By Computer), storage capacity (By Disk/Volume Group Space) total free space (By Available Disk/Volume Group Space), time when they were last probed (By Probe Time), and time when they were discovered (By Discovered Time).

Disks reports

Use these reports to view all the disks in the system, which are sorted in order of how many recovered (By Recovered Errors) and unrecovered errors (By Unrecovered Errors) they have, time when they were probed (By Probe Time), and time when they were discovered (By Discovered Time).

Use the **Asset Reporting** > **System-wide** > **Disks** > **By Grown Defects** report to view the following information.

- All the disks in the system that are sorted in order of the number of grown defects that they have on them
- What disks in your system are experiencing problems and might be in danger of an imminent failure

The **Asset Reporting** > **System-wide** > **Disks** > **By Primary Defects** report returns no data if disk defects are not detected. Use this report to view the following information.

- All the disks in the system that are sorted in order of the number of primary defects they have on them
- What disks in your system are in good condition and which ones are more or less defective

File systems / Logical Volumes reports

Use these reports to view detailed information about file systems and logical volumes on your monitored computers. A logical volume is a partition on a disk that does not have a file system (raw partition). Use these reports to view the following information.

- Available file system space on monitored computers.
- All the file systems in the system that are sorted in order of the amount of free space they have on them.
- Which file systems have less free space and might be in danger of running out of free space.
- Which file systems have too much free space and where space consolidation might be necessary.
- A list of virtual disk files, which are not assigned to any of the virtual machines. This situation represents wasted storage capacity.
- The path from the file system on a virtual machine to the back-end storage volumes on a hypervisor or mapped to the hypervisor.
- The path from the file system on a virtual machine to the back-end storage volumes on an IBM SONAS or Storwize V7000 Unified system, or mapped to the IBM SONAS or Storwize V7000 Unified system.

Note: The values in the By Available File System Space report are based on data that is collected by probes. Probes collect space information as measured in blocks (1 block = 512 bytes). Tivoli Storage Productivity Center calculates file system space, including used space and free space,

which is based on the number of blocks it retrieves from a probe. It multiplies the number of blocks by 512 and displays the resulting value in the report.

Volumes reports

Use these reports to view the following information.

- Information about the storage subsystems that host volumes
- Storage capacity of the volumes hosted by storage subsystems
- Information about the storage subsystems that host volumes sorted by the date/time when they were probed

Exports or Shares report

Use this report to view all the Exports or Shares in your environment, which are sorted in the order that they were discovered.

Monitored Directories reports

Use these reports to view all the monitored directories in the system, which are sorted in order of the amount of storage space they are taking up (By Space Used), the time when they were last scanned (By Scan Time), the time when they were discovered (By Discovered Time), and when they were removed (Removed Directories).

Unmanaged Computers report

Use this report to view all of the new computers that are introduced into your system and that do not have Storage Resource agents or Data agents that are installed on them. It also lists other computers that you do not put under Data Manager control.

Users report

Use this report to view all the users in your environment that are sorted in the order that they were discovered.

OS User Groups report

Use this report to view all the OS User Groups in your environment that are sorted in the order when they were discovered.

Availability Reports:

Use Availability reports to monitor and report on the availability of your storage, both from a network point of view and from a computer uptime perspective.

Use these reports to view the percentage of time that the storage in your enterprise is off-network due to network problems or is unavailable due to system downtime. In a fast-paced business environment, 24x7 application availability is critical. These reports help you monitor the availability of the storage and data that back up your applications.

Reporting Type	Description
Ping	Use the ping report to monitor how available your storage is on the network. To do this monitor, regularly scheduled pings are sent out from the Data server to all of the monitored computers in your environment. The Tivoli Storage Productivity Center server checks not only for whether a Storage Resource agent is available on the network, but records the amount of time that the request was in transit. This information indicates the responsiveness of the various storage servers on the network as well. Use the **Report Filter Specification** section on the **Selection** panel to limit the time range of the information in a report. Check the **Limit To This Time Range** option and use the date selection fields to limit the time range for a report. A ping report shows the following information: • the name of a computer • percentage of time the computer is available on the network • number of pings that have been transmitted and the number that have been received • number of damaged requests and the average transit time for the request to go between the Tivoli Storage Productivity Center server and the computer that is being monitored • maximum transit time • minimum transit time • report start time and report end time for each computer in the report **Note:** You can sort on any of these columns by clicking on the column name. You can also generate charts and drill down on each computer for more detail.

Reporting Type	Description
Computer Uptime	Use the computer uptime report to track the availability of the storage based upon the percentage of the time the computer that hosts that storage is down or not running. The Storage Resource agent on the computer tracks this information and sends it to the server. A computer uptime report shows the following information: • name of a computer • percentage of time the computer is up and running • number of reboots that have been recorded • total amount of time the computer has been down and unavailable • average downtime that has been recorded • report start time and report end time for each computer in the report **Note:** Computer Uptime reports are not available for Storage Resource agents that are invoked using the non-daemon protocol for data collection. These reports do not contain data for computers on which non-daemon based Storage Resource agents are deployed. To collect computer uptime information using a Storage Resource agent, that agent must use a daemon service for runtime operation. You can sort on any of these columns by clicking on the column name. You can also generate charts and drill down on each computer for more detail.

TPC-wide Storage Space Reports:

Use these reports to view storage capacity of your resources at a file system, computer, or entire network level.

For a specific computer, group of computers, or all the computers in your enterprise, you can view the following information:

- How much storage you have
- How much storage is being used
- What percentage of storage is free for expansion

Reporting Type	Description
Disk Space	Use these reports to view and chart disk capacity, per disk, per computer, per computer group, per domain, and for the whole network. These reports include one row for each computer that shows the total amount of storage capacity on that computer and the associated computer and disk storage information. For detail on any computer in the report, click the magnifying glass next to the report row. You can also click the chart button within a report to view an easy-to-interpret bar chart, showing the total storage capacity of each computer that you chose to report on and chart. To get the exact amount of storage for a computer, position the cursor over the associated bar and a box displays the total amount for the selected computer. **Note:** To ensure the accuracy and consistency of the data that is displayed in Disk Capacity reports, run regularly scheduled probes against the hosts that import volumes and the storage subsystems upon which you want to report. Run probes against storage subsystems after you run them against any hosts. You can use the following By clauses to view disk space for the following storage entities: By Disk, By Storage Subsystem, By Cluster, By Computer, By Computer Group, By Domain, and Network-wide.
File System Space	Use these reports to view information for monitored file systems, including the total amount of storage capacity on those file systems and statistics for used space, free space, and total number of files. You can use the following By clauses to view file system space for the following storage entities: By File System, By File System Group, By Cluster, By Computer, By Computer Group, By Domain, Network-wide.
Consumed File System Space	Use these reports to view information about the amount of space that is consumed on a file system. You can use the following By clauses to view consumed file system space for the following storage entities: By File System, By File System Group, By Cluster, By Computer, By Computer Group, By Domain, Network-wide.
Available File System Space	Use File System Used Space reports to view information about the amount of free space on a file system. You can use the following By clauses to view available file system space for the following storage entities: By File System, By File System Group, By Cluster, By Computer, By Computer Group, By Domain, Network-wide.

TPC-Wide Storage Space - Disk Space report:

Use this report to view and chart disk storage capacity at the following levels: by disk, by computer, by cluster, by storage system, by computer monitoring or reporting group, by storage system monitoring or reporting group, by domain, and for the whole network.

Note: The names of some columns are changed. For your reference, the original names of those columns display within parentheses ( ) next to their new names.

By Disk

Chart icon
Click to generate charts about the capacity of disks that are displayed in the report.

Magnifying glass icon
Click to view a more detailed report about the capacity of a disk.

Device

Name of the device on which a disk is in. The following disks are displayed in this report.

- Computer disks that are internal to the computers monitored by IBM Tivoli Storage Productivity Center.
- Computer disks on virtual machines that are monitored by Tivoli Storage Productivity Center.
- Storage system disks on storage subsystems that are monitored by Tivoli Storage Productivity Center.
- Storage system disks that no longer exist after a probe of a storage system because information about that system was removed based on settings in the Removed Resource Retention panel.
- Back-end storage volumes that are not identified by Tivoli Storage Productivity Center as being assigned to a monitored storage system virtualizer.
- Computer disks that are not identified by Tivoli Storage Productivity Center as being assigned to a monitored storage system.

The following disks are not displayed in this report.

- Computer disks that are assigned to storage subsystems monitored by Tivoli Storage Productivity Center.
- Back-end storage volumes that are assigned to storage system virtualizers monitored by Tivoli Storage Productivity Center.

Path Name of a disk.

Disk Space (Capacity)

Total storage capacity for a disk. Keep in mind the following information about the TOTAL row for this column.

- The value that displays in TOTAL row for this column includes the disk space of spare disks within a storage subsystem.
- The value that displays in the TOTAL row for this column does not include computer disk space that is on virtual machines that are monitored by Tivoli Storage Productivity Center.

Note: When you use the stand-alone GUI, the Disk Space entries for Storwize V7000 virtualizers report the sum of backend capacity. Local disk space is not included. When you use the web-based GUI, local disk space is included in the calculation of the total device disk space of Storwize V7000 virtualizers.

Available Disk Space (Not Allocated Space)

Amount of space on the disk that is not allocated to any logical volume (or volumes for disks within storage subsystems). This value does not include RAID disk usage on storage subsystems that are identified by Tivoli Storage Productivity Center. The value that displays in the TOTAL row for this column includes the capacity of all disks within the storage system that is not allocated. The capacity of spare disks is included.

In Storage Subsystem

Indicates `Yes` when a disk is either of the following.

- An actual physical disk within a storage system that is monitored by Tivoli Storage Productivity Center.
- A back-end storage volume within a storage system virtualizer that is monitored by Tivoli Storage Productivity Center.

Manufacturer
Name of the manufacturer of a disk.

Note: The value `Unknown` is displayed in this column for IBM ESS storage subsystems.

Model
Model number of a disk.

Note: The value `Unknown` is displayed in this column for IBM ESS storage subsystems.

Serial Number
Serial number of a disk.

Primary Defects
Number of defects on the disk when it was new.

Note: Defect data is not collected for IBM TotalStorage Enterprise Storage Server, IBM FAStT, HP StorageWorks, Hitachi Data Systems storage subsystems.

Grown Defects
Number of defects that are detected on the disk since it was new.

Note: Defect data is not collected for TotalStorage Enterprise Storage Server, FAStT, HP StorageWorks, Hitachi Data Systems storage subsystems.

Fibre Attached
Indicates `Yes` when a disk is attached to a port that Tivoli Storage Productivity Center identifies as fibre.

Mapped to Storage Subsystem
Indicates `Yes` when a computer disk or a back-end storage volume can be mapped to a storage system.

Detectable
Indicates `Yes` when a disk is identified as being available to the server or the storage system after a successful probe.

Is Virtual
Indicates `Yes` when a disk is on a computer that Tivoli Storage Productivity Center identifies as a virtual machine.

Is Solid State
Indicates `Yes` when a disk is on a solid-state drive.

Is Encryptable
Indicates `Yes` when a disk can be encrypted.

Is Encrypted
Indicates `Yes` when a disk is encrypted.

Tier Indicates the tier level of a disk. Disks can have only one tier. For example, this column displays either the value `generic_ssd` or `generic_hdd`.

By Storage Subsystem

Chart icon
Click to generate charts about the disk capacity of storage subsystems that are displayed in the report.

Magnifying glass icon

Click to view a more detailed report about the capacity of disks that are associated with a storage system.

Name

Name or alias of a storage system that contains the disks or the name or alias of a storage system virtualizer (for example, SAN Volume Controller) that manages the disks. This report does not display storage subsystems that are discovered but not probed.

Disk Space (Disk Capacity)

Total raw (unformatted) disk capacity of a storage system. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column.

- Capacity of spare disks that are identified on TotalStorage Enterprise Storage Server, DS8000, and IBM System Storage DS6000 storage subsystems.
- Capacity of storage system disks that become missing after a Tivoli Storage Productivity Center storage system probe.
- The value that displays in the TOTAL row for this column does not include spare disks in its calculation.

Available Disk Space (Disk Space Not Allocated)

Total unformatted disk free space of a storage system. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column.

- Storage system disks that become missing after a storage system probe.
- RAID storage usage on storage subsystems that are identified by Tivoli Storage Productivity Center.
- External back-end storage for TagmaStore subsystems.
- Storage space that is used by volumes that are on external disks when Hitachi Data Systems (HDS) storage subsystems are configured as storage virtualizers.
- The value that displays in the TOTAL row for this column does not include spare disks in its calculation.

Number of Disks

Total number of disks in a storage system. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Storage system disks that become missing after a storage system probe.
- External back-end storage for TagmaStore subsystems.
- Disks in partner storage pools for NetApp storage devices.

Note: If you click the magnifying glass icon next to storage system, the resulting drill-down report generates a list of detectable storage system disks or back-end storage volumes only. Therefore, the total number of disks for each storage system drill-down report equals this value.

Number of Volumes (LUNs)

Total number of storage volumes in a storage system or the total number of virtual volumes that are virtual storage. This value does not include storage system volumes or virtual storage volumes that become missing after a storage system probe.

Volume Space (LUN Capacity)
Total capacity of all storage system volumes or virtual storage volumes within a storage system. This value does not include capacity information for storage system volumes or virtual storage volumes that become missing after a storage system probe.

Note: For space-efficient FlashCopy volumes, the volume capacity is the amount of storage space that is requested for these volumes, not the actual allocated amount. This condition can result in discrepancies in the overall storage capacity reported for a storage system that uses space-efficient FlashCopy volumes. This situation also applies to the Total Usable Volume Space and Total FlashCopy Target Capacity calculations.

Consumable Volume Space (Usable LUN Capacity)
Total amount of unique storage system volume space and virtual storage volume space on monitored storage system arrays. You can assign this space to servers within the network or they are already assigned to servers within the network.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes or virtual storage volumes that are used as FlashCopy targets.
- Storage system volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage system volumes or virtual storage volumes that are used as a Volume Configuration Management (VCM) database.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers. For example, SAN Volume Controller.

Note: For FAStT, the XIV system, HP StorageWorks, and Hitachi Data Systems storage subsystems, the value in this column is the same as the value in the Volume Capacity column.

FlashCopy Target Volume Space (Capacity)
Total amount of unique storage system volume space and virtual storage volume space on monitored storage system arrays that is identified as FlashCopy target space. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are not used as FlashCopy targets.

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Total amount of storage system volume space or virtual storage volume space on systems where the systems disks are identified by Tivoli Storage Productivity Center within a monitored storage system. The volumes include storage system volumes or virtual storage volumes that are identified as back-end storage volumes on monitored storage system

virtualizers. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that are assigned to a monitored storage system but are not identified by Tivoli Storage Productivity Center as a disk within the storage system. Some reasons why these disks are not identified as being in the storage system are included.
  - Tivoli Storage Productivity Center does not support the identification of disks for some storage system types.
  - The connection of the disk to the storage system is lost and is no longer detectable during a probe.
- Storage system volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage system volumes or virtual storage volumes that are used as a Volume Configuration Management (VCM) database.

Volume Space Not Correlated (LUN Capacity not Visible to Data Manager)
Total amount of unique storage system volume space and virtual storage volume space that is discovered on monitored storage system arrays that can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that are within a monitored storage system.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers.

Unavailable Disk Space (Disk Usage)
Amount of storage system volume or virtual storage volume space that is dedicated to redundancy. This value is dependent on the storage system SMI-S provider that returns a valid RAID value that is used to determine the disk usage. If the value is zero for a storage system, the disk usage cannot be calculated.

This amount does not include storage space information from storage system volumes or virtual storage volumes that become missing after a storage system probe.

Formatted Space (Capacity)
Total amount of formatted storage space that is associated with the storage system that is used or can be used for system storage volumes or virtual storage volumes. This value is dependent on the storage system SMI-S provider that returns a formatted value.

For some storage subsystems, this space is the total managed space, which includes space that cannot be used for storage volumes or virtual storage volumes. This situation can occur because of a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the total size of the managed disk groups that are created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

Formatted Space with No Volumes
Total amount of formatted storage space that is associated with the storage system that can be used for system storage volumes or virtual storage volumes. This value is dependent on the storage system SMI-S provider that returns a formatted value.

For some subsystems, this value represents the remaining managed space that includes space that cannot be used for storage volumes or virtual storage volumes. This situation can occur because of a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the space within the managed disk groups that does not include the size of the virtual disks that are created and the corresponding virtual disk extent usage.

Overall Unavailable Disk Space (Total Disk Usage)
(TotalStorage Enterprise Storage Server, DS6000, and DS8000 storage systems only) Total amount of unformatted space within disk groups or array sites.

Unformatted Disk Space
This value does not apply to all storage subsystems. For FAStT, TotalStorage Enterprise Storage Server, DS6000, and DS8000 storage systems, this value is the total raw disk space of any unformatted disk groups or array sites. This column displays `N/A` for the storage systems where this value is not applicable. For the XIV system, this value shows the physical (hard) space of the XIV system available for pool creation or expansion.

Cache
Amount of cache memory that is internal to the storage system.

Note: This column is blank for SAN Volume Controllers.

Manufacturer
Name of the manufacturer of a disk.

Model
Model number of a disk.

Serial Number
Serial number of a disk. This information is displayed for physical disks only. For MDisks, the value `Unknown` is displayed.

Firmware Revision
Microcode (OS) version.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage system volume space within the storage system that is mapped or assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes that are not mapped or are not assigned to host systems.

Unassigned Volume Space (Subsystem Unassigned LUN Capacity)
Total storage system volume space within a storage system that is not

mapped or not assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes that are mapped or assigned to host systems.

Unassigned and Available Formatted Disk Space (Subsystem Total Formatted Free Space)

Total storage system space that is not mapped or not assigned to host systems and the total amount of formatted storage space that is associated with the storage system that can be used for system storage volumes.

This total amount of formatted storage space is dependent on the storage system SMI-S provider that returns a formatted value. For some subsystems, this value represents the remaining managed space that might include space that cannot be used for storage volumes or virtual storage volumes. This situation can occur because of a number of reasons, such as space not being formatted until storage volumes are created.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column.

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are mapped or assigned to host systems.

zOS Volume Space (zOS LUN Capacity)

Total storage system space that is reserved for use on mainframe storage such as used by z/OS. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column.

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on open system server or non-mainframe storage.

Open System Volume Space (Open System LUN Capacity)

Total storage system space that is reserved for use on open system storage such as used by operating systems like Linux, AIX, HP/UX, Oracle Solaris, and Windows. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column.

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on mainframe storage.

Physical Disk Space (Capacity)

Total amount of physical disk space that is discovered on the monitored storage system arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column.

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

Backend Volume Space (Backend LUN Capacity)

Total amount of storage system volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage system virtualizer (for example, SAN Volume

Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are back-end disks on a storage system virtualizer that is not monitored.

Configured Real Space (GB)
Total amount of real available and formatted storage space in the storage pools that are associated with a storage subsystem.

Available Real Space (GB)
Amount of real storage space available to allocate to volumes in the storage pools that are associated with a storage subsystem.

Tiers Indicates the tiers that are used by the storage system. Each tier is delimited by a semi-colon (;). Possible values are: generic_hdd, generic_ssd

Information about tiers is available for SAN Volume Controller 6.1 and Storwize V7000 only. The value `N/A` is displayed in tier-related columns for other storage systems.

Tier Capacities
Total storage capacity for each tier that is used by the storage system. The capacity for each tier is delimited by a semi-colon (;).

Tier Free Capacities
Free storage capacity (available space) for each tier that is used by the storage system. The free capacity for each tier is delimited by a semi-colon (;).

By Storage Subsystem Group

Chart icon
Click to generate charts about the disk capacity of storage subsystems in the groups that are displayed in the report.

Magnifying glass icon
Click to view a more detailed report about the disk capacity of storage subsystems in a group.

Creator
ID of the user who created the storage system monitoring group or storage system reporting group.

Group Name of a group.

Disk Space (Disk Capacity)
Total raw (unformatted) disk capacity of the storage subsystems that are in a group. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Capacity of spare disks that are identified on TotalStorage Enterprise Storage Server, DS8000, and DS6000 storage systems.
- Capacity of storage system disks that become missing after a Tivoli Storage Productivity Center storage system probe.

Available Disk Space (Not Allocated Disk Space)
Total unformatted disk free space of the storage subsystems that are in a group. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column.

- Storage system disks that become missing after a storage system probe.
- RAID disk usage on storage subsystems that are identified by Tivoli Storage Productivity Center.
- External back-end storage for TagmaStore subsystems.
- Storage space that is consumed by volumes that are on external disks when Hitachi Data Systems (HDS) storage subsystems are configured as storage virtualizers.

Number of Disks

Total number of disks in the storage subsystems that are in a group. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Storage system disks that become missing after a storage system probe.
- External back-end storage for TagmaStore subsystems.
- Disks in partner storage pools for NetApp storage devices.

Note: If you click the magnifying glass icon next to storage system, the resulting drill-down report generates a list of detectable storage system disks or back-end storage volumes only. Therefore, the total number of disks for each storage system drill-down report equals this value.

Number of Volumes (LUNs)

Total number of storage volumes in a storage system or the total number of virtual volumes that are virtual storage. This value does not include storage system volumes or virtual storage volumes that become missing after a storage system probe.

Volume Space (LUN Capacity)

Total capacity of all storage system volumes or virtual storage volumes within the storage system. This value does not include capacity information of storage system volumes or virtual storage volumes that become missing after a storage system probe.

Note: For space-efficient FlashCopy volumes, the volume capacity is the amount of storage space that is requested for these volumes, not the actual allocated amount. This situation can result in discrepancies in the overall storage capacity reported for a storage system that uses space-efficient FlashCopy volumes. This situation also applies to the Total Usable Volume Space and Total FlashCopy Target Capacity calculations.

Consumable Volume Space (Usable LUN Capacity)

Total amount of unique storage system volume space and virtual storage volume space on monitored storage system arrays in a group that you can assign to servers within the network or are already assigned to servers within the network.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes or virtual storage volumes that are used as FlashCopy targets.
- Storage system volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage system volumes or virtual storage volumes that are used as a Volume Configuration Management (VCM) database.

- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers. For example, SAN Volume Controller.

Note: For FAStT, the XIV system, HP StorageWorks, and Hitachi Data Systems storage subsystems, the value in this column is the same as the value in the Volume Space column.

FlashCopy Target Volume Space (Capacity)
Total amount of unique storage system volume space and virtual storage volume space on monitored storage system arrays in a group that is identified as FlashCopy target space. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are not used as FlashCopy targets

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Total amount of storage system volume space or virtual storage volume space on systems where the systems disks are identified by Tivoli Storage Productivity Center as being in the monitored storage subsystems in a group. The volumes include storage system volumes or virtual storage volumes that are identified as back-end storage volumes on monitored storage system virtualizers.Tivoli Storage Productivity Center does not include the following scenarios in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that are assigned to a monitored storage system but are not identified by Tivoli Storage Productivity Center as a disk in the storage system. Some reasons why these disks are not identified as being in the storage system include.
  - Tivoli Storage Productivity Center does not support the identification of disks for some storage system types.
  - The connection of the disk to the storage system is lost and is no longer detectable during a probe.
- Storage system volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage system volumes or virtual storage volumes that are used as a Volume Configuration Management (VCM) database.

Volume Space Not Correlated (LUN Capacity not Visible to Data Manager)
Total amount of unique storage system volume space and virtual storage volume space that is discovered on monitored storage system arrays in a group that can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that are within a monitored storage system.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers.

Unavailable Disk Space
Amount of volume or virtual storage volume space of storage subsystems in a group that is dedicated to redundancy. This value is dependent on the storage system SMI-S provider that returns a valid RAID value that is used to determine the disk usage. If the value is zero for a storage system, the disk usage cannot be calculated.

This amount does not include storage space information from storage system volumes or virtual storage volumes that become missing after a storage system probe.

Formatted Space
Total amount of formatted storage space that is associated with the storage subsystems in a group that is used or can be used for system storage volumes or virtual storage volumes. This value is dependent on the storage system SMI-S provider that returns a formatted value.

For some storage subsystems, this space is the total managed space, which includes space that cannot be used for storage volumes or virtual storage volumes. This situation can occur because of a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the total size of the managed disk groups that are created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

Formatted Space with No Volumes
Total amount of formatted storage space that is associated with the storage subsystems in a group that can be used for system storage volumes or virtual storage volumes. This value is dependent on the storage system SMI-S provider that returns a formatted value.

For some subsystems, this value represents the remaining managed space that includes space that cannot be used for storage volumes or virtual storage volumes. This situation can occur because of a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the space within the managed disk groups that does not include the size of the virtual disks that are created and the corresponding virtual disk extent usage.

Unformatted Disk Space
This value does not apply to all storage subsystems. For FAStT, TotalStorage Enterprise Storage Server, DS6000, and DS8000 storage systems, this space is the total raw disk space of any unformatted disk groups or array sites. For the XIV system, the total physical (hard) space that is available for pool creation. This column displays `N/A` for the storage subsystems where this value is not applicable.

Cache
Amount of cache memory that is internal to the storage subsystems in the group.

Note: This column does not include values for SAN Volume Controllers.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total volume space of the storage subsystems within a group that is mapped or assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes that are not mapped or are not assigned to host systems.

Unassigned Volume Space (Subsystem Unassigned LUN Capacity)
Total volume space of the storage subsystems in a group that is not mapped or not assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes that are mapped or assigned to host systems.

zOS Volume Space (zOS LUN Capacity)
Total space of storage subsystems in a group that is reserved for use on mainframe storage such as used by z/OS. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column.

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on open system server or non-mainframe storage.

Open System Volume Space (Open System LUN Capacity)
Total space of the storage subsystems in a group that is reserved for use on open system storage such as used by operating systems such as Linux, AIX, HP/UX, Oracle Solaris, and Windows. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column.

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on mainframe storage.

Physical Disk Space (Capacity)
Total amount of physical disk space on the storage system arrays that are in a group. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column.

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

Backend Volume Space (Backend LUN Capacity)
Total amount of storage system volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on the monitored storage system virtualizers (for example, SAN Volume Controller) that are in a group. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column.

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are back-end disks on a storage system virtualizer that is not monitored.

By Cluster

Chart icon
Click to generate charts about the disk capacity of the clusters that are displayed in the report.

Magnifying glass icon
Click to view a more detailed report about the capacity of disks in a cluster.

Cluster
Name of a cluster.

Note: For Storwize V7000 Unified or IBM SONAS clusters, the storage system name is included with the cluster name in the Computer or Device column. The format of the column is `storagesystem_name: cluster_name`, where storagesystem_name is the name of the Storwize V7000 Unified or IBM SONAS device that the cluster belongs to. For example, `tpcserver1.tpc.example.com: kq98n5d.acct` displays storage system tpcserver1.tpc.example.com that cluster kq98n5d.acct belongs to.

Disk Space (Total Capacity)
Total amount of unique disk storage space for the computers in a cluster. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value in this column.

- Computer disks on a computer that is not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note:

- If the computer is a SAN File System client, this value includes the capacity of storage system volumes that are visible to the computers in the cluster and are owned by the SAN File System.
- The drill-down reports for each computer generate a list of detectable computer disks only. Therefore, the total capacity of disks for each computer drill-down report equals this value.

Available Disk Space (Total Not Allocated Space)
Total amount of unique disk storage space for a computer that is not allocated to any logical volume on the computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks on computers that are not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note: If the computer is a SAN File System client, then this value does not include the space not allocated on any storage system volumes visible to it that are owned by the SAN File System.

Owned Disk Space (Capacity)
Total amount of unique disk storage space for computers in a cluster that is owned by those computers. If Data Manager determines that a disk is configured for a different host, that disk's capacity is not counted in this

column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks on computers that are not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note:

- If the computer is a SAN File System client, this value does not include the capacity of any volumes that are visible to it that are owned by the SAN File System.
- If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads from the host ID. If this examination does not agree with the local host ID, the disk is classified as belonging to a different (Solaris) host.

Owned Available Disk Space (Owned Not Allocated Space)
Total amount of unique disk storage space for computers in a cluster that is not allocated to any logical volume within the computers and is owned by the computers. If Data Manager determines that a disk is configured for a different computer, the space that is not allocated for that disk is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks on computers that are not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

If the computer is a SAN File System client, then this value excludes the space not allocated on any storage system volumes visible to it that are owned by the SAN File System.

If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and read off the host-ID. If this examination does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)
Total amount of unique disk space on the computers in a cluster that is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks that are behind a fibre channel port.
- Computer disks on a computer that is not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Non-Fibre Channel Attached Available Disk Space (Total Other Attached Not Allocated Space)
Total amount of disk storage space on the computers in a cluster that is not allocated to any logical volume within the computers and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks that are behind a fibre channel port.
- Computer disks on a computer that is not yet probed.

- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Disk Space (Total Fibre Channel Attached Capacity)
Total amount of unique disk space on the computers in a cluster that is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks that are not behind a fibre channel port.
- Computer disks that are behind a fibre channel port that is not identified by a Storage Resource agent.
- Computer disks on a computer that is not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Not Allocated Space)
Total amount of disk storage space on the computers in a cluster that is not allocated to any logical volume within the computers and is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks that are not behind a fibre channel port.
- Computer disks that are behind a fibre channel port that is not identified by a Storage Resource agent.
- Computer disks on a computer that is not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

By Computer

Chart icon
Click to generate charts about the disk capacity of the computers that are displayed in the report.

Magnifying glass icon
Click to view a more detailed report about the disks on a computer.

Computer
Name of a computer.

Disk Space (Total Capacity)
Total amount of unique disk storage space for a computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks on a computer that is not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note:

- The value that displays in the TOTAL row for this column does not include disks on virtual machines where the Storage Resource agent is installed.
- If the computer is a SAN File System client, then this value also includes the capacity of any storage system volumes visible to this computer that are owned by the SAN File System.
- The drill-down reports for each computer generate a list of detectable computer disks only. Therefore, the total capacity of disks for each computer drill-down report equals this value.

Available Disk Space (Total Not Allocated Space)
Total amount of unique disk storage space for a computer that is not allocated to any logical volume on the computer. Tivoli Storage Productivity Center does not include the following disks from in its calculation of the value in this column.

- Computer disks on computers that are not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note: If the computer is a SAN File System client, then this value does not include the space not allocated on any storage system volumes visible to it that are owned by the SAN File System.

Owned Disk Space (Capacity)
Total amount of unique disk storage space for a computer that is owned by that computer. If Data Manager determines that a disk is configured for a different host, that disk's capacity is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks on computers that are not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

If the computer is a SAN File System client, this value does not include the capacity of any volumes that are visible to it that are owned by the SAN File System.

If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads from the host ID. If this examination does not agree with the local host ID, the disk is classified as belonging to a different (Solaris) host.

Owned Available Disk Space (Not Allocated Owned Space)
Total amount of unique disk storage space for a computer that is not allocated to any logical volume within the computer and is owned by the computer. If Data Manager determines that a disk is configured for a different computer, that disk space that is not allocated is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks on computers that are not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note:

- If the computer is a SAN File System client, then this value excludes the space not allocated on any storage system volumes visible to it that are owned by the SAN File System.
- If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and read off the host-ID. If this examination does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)
Total amount of unique disk space on a computer that is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks that are behind a fibre channel port.
- Computer disks on a computer that is not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Non-Fibre Channel Attached Available Disk Space (Total Other Attached Not Allocated Space)
Total amount of disk storage space for a computer that is not allocated to any logical volume within the computer and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks that are behind a fibre channel port.
- Computer disks on a computer that is not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Disk Space (Total Fibre Channel Attached Capacity)
Total amount of unique disk space on a computer that is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks that are not behind a fibre channel port.
- Computer disks that are behind a fibre channel port that is not identified by a Storage Resource agent.
- Computer disks on a computer that is not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Not Allocated Space)
Total amount of disk storage space on a computer that is not allocated to any logical volume within the computer and is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks that are not behind a fibre channel port.
- Computer disks that are behind a fibre channel port that is not identified by a Storage Resource agent.
- Computer disks on a computer that is not yet probed.

- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

OS Type
Operating system that runs on a computer.

Network Address
Network address of a computer.

IP Address
IP address of a computer.

Time Zone
Time zone in which a computer is running.

CPU Architecture
Architecture of the processor (for example, IA32).

By Computer Group

Chart icon
Click to generate charts about the disk capacity of computers in the groups that are displayed in the report.

Magnifying glass icon
Click to view a more detailed report about the disk capacity of computers in a group.

Creator
Creator of a computer monitoring group or computer reporting group.

Group
Name of a group.

Disk Space (Total Capacity)
Total amount of unique disk storage space for the computers in a group. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks that are on a virtual machine that is monitored by Tivoli Storage Productivity Center.
- Computer disks on a computer that is not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Available Disk Space (Total Not Allocated Space)
Total amount of unique disk storage space for computers in a group that is not allocated to any logical volume on the computers. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks on computers that are not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note: If the computer is a SAN File System client, then this value does not include the space not allocated on any storage system volumes visible to it that are owned by the SAN File System.

Owned Disk Space (Capacity)

Total amount of unique disk storage space for computers in a group that is owned by those computers. If Data Manager determines that a disk is configured for a different host, that disk's capacity is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks on a cluster of virtual machines that is monitored by Tivoli Storage Productivity Center.
- Computer disks on computers that are not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note:

- If the computer is a SAN File System client, this value does not include the capacity of any volumes that are visible to it that are owned by the SAN File System.
- If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads from the host ID. If this examination does not agree with the local host ID, the disk is classified as belonging to a different (Solaris) host.

Owned Available Disk Space (Owned Not Allocated Space)

Total amount of unique disk storage space for the computers in a group that is not allocated to any logical volume within the computer group and is owned by a computer within the group. If Data Manager determines that a disk is configured for a different computer, the disk space that is not allocated is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks on computers that are not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note:

- If the computer is a SAN File System client, then this value excludes the not allocated space on any storage system volumes visible to it that are owned by the SAN File System.
- If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and read off the host-ID. If this examination does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)

Total amount of unique disk space on the computers in a group and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks that are behind a fibre channel port.
- Computer disks that are discovered on virtual machine servers where the Storage Resource agent is installed.
- Computer disks on a computer that is not yet probed.

- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Non-Fibre Channel Attached Available Disk Space (Total Other Attached Not Allocated Space)

Total amount of disk storage space on the computers in a group that is not allocated to any logical volume within the computers and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks that are behind a fibre channel port.
- Computer disks on a computer that is not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Disk Space (Total Fibre Channel Attached Capacity)

Total amount of unique disk space on the computers in a group that is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks that are not behind a fibre channel port.
- Computer disks that are behind a fibre channel port that is not identified by a Storage Resource agent.
- Computer disks that are discovered on virtual machine servers where the Storage Resource agent is installed.
- Computer disks on a computer that is not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Not Allocated Space)

Total amount of disk storage space on the computers in a group that is not allocated to any logical volume within the computers and is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks that are not behind a fibre channel port.
- Computer disks that are behind a fibre channel port that is not identified by a Storage Resource agent.
- Computer disks on a computer that is not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

By Domain

Chart icon

Click to generate charts about the disk capacity of the computers in a domain.

Magnifying glass icon

Click to view a more detailed report about the disk capacity of computers in a domain.

Domain
Name of a domain.

Disk Space (Capacity)
Total amount of unique disk storage space for the computers in a domain. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks on a computer that is not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Available Disk Space (Not Allocated Space)
Total amount of unique disk storage space for computers in a domain that is not allocated to any logical volume on the computers. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column.

- Computer disks on computers that are not yet probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Network-wide

Chart icon
Click to generate charts about the disk capacity across a network.

Magnifying glass icon
Click to view a more detailed report about the disk capacity of computers in a network.

Disk Space (Capacity)
Total amount of internal physical disk space on monitored storage system arrays and disk space on monitored systems that are not storage system volumes. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column.

- Storage system disks or computer disks that become missing after a probe.
- Back-end storage volumes that are not identified by Tivoli Storage Productivity Center as being assigned to a monitored storage system virtualizer.
- Computer disks that are identified by Tivoli Storage Productivity Center as being assigned to a monitored storage system.
- Computer disks that are within a virtual machine server where the Storage Resource agent is installed.

Consumed File System Space (File System Used Space)
Total amount of unique file system storage space that is used or consumed by the systems where the Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems on discovered hypervisors that are assigned to discovered virtual machines where the Storage Resource agent is installed.

- File systems on discovered virtual machines where the Storage Resource agent is installed. The hypervisor for the virtual machine is not yet discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Available File System Space (File System Free Space)
Total amount of unique file system storage space that is not used or available to the systems where the Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Raw Volume Space
Total amount of storage system volume space that is the following.

- Assigned or mapped to hosts.
- Storage system volume space on a system that it is monitored by Tivoli Storage Productivity Center.
- Logical volume space on hosts that are not within a storage system that is monitored by Tivoli Storage Productivity Center.

Unavailable Disk Space (Disk Usage)
Total amount of storage system volume space that is identified as the following.

- For RAID or mirror redundancy on a storage system.
- As storage system volume space used for replication targets.
- As logical volume space on hosts that is for RAID or mirror redundancy on those hosts.

Available Disk Space (Not Allocated Space)
Total amount of disk space that does not include the following from the network capacity.

- Network Raw Volume Space
- Network Disk Usage
- Network Unknown Volume Space

Unknown Storage Subsystem Volume Space (Unknown LUN Capacity)
Total amount of storage system volume space that is not mapped or assigned to a host and is not identified on a host that is probed by a Storage Resource agent.

TPC-Wide Storage Space - File System Space report:

Use these reports to view detailed information about the storage capacity of file systems on systems in your environment. You can view this information in different ways by using the following By-Clauses: by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

Note: The names of some columns are changed. For your reference, the original names of those columns are within parentheses ( ) next to their new names.

By File system

Chart icon
Click this button to generate charts about the storage space for a file system.

Computer
Name of the computer on which a file system is in.

Mount Point
Name or mount point (UNIX/Linux) of the file system (for example, `c:\`, `d:\`, `/opt`, `/export/home`, and so on). The file systems that are displayed in this report include the following information.

- File systems on virtual machines that are monitored by IBM Tivoli Storage Productivity Center.
- File systems on hypervisors that are monitored by Tivoli Storage Productivity Center.
- File systems that become missing after a probe.

File System Space (Capacity)
Amount of file system storage space for a file system.

Note:

- The value that is displayed in the TOTAL row for this column does not include file systems on virtual machines where the Storage Resource agent is installed.
- The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**. The difference represents the space that is wasted because of the system overhead.

Detectable
Indicates "Yes" when a file system is identified as being available to the server after a successful probe.

Percent Used Space
Percentage of file system storage space that is used within a file system.

Consumed File System Space (Used Space)
Amount of used storage space within a file system. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead). Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

The value in the TOTAL row for this column represents the difference between the TOTAL values in the **File System Space** and **Available File System Space** columns. If the calculated value for the TOTAL row is negative, the column displays 0.

The TOTAL row does not add the used space from VMWare computers. See the *VMWare capacity reports* topic in the *Planning* chapter of the *Tivoli Storage Productivity Center Installation and Configuration Guide* for information about how file system and disk capacity are calculated in this report for VMWare computers. This information is also available in the *Planning* section of the *Tivoli Storage Productivity Center Information Center*.

The value in the TOTAL row for this column does not include the following file systems.

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent is installed.
- File systems on discovered virtual machine servers where the Storage Resource agent is installed, but the hypervisor for the virtual machine servers is not discovered and the total used space exceeds the capacity of the total file system capacity.

Available File System Space (Free Space)
Amount of unused storage space within a file system.

File Count
Total number of files within a file system. This value does not include files on file systems that are not scanned.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS™) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a file system. This value does not include directories on file systems that are not scanned.

Percent Used Inodes
Percentage of used Inodes within a file system.

Used Inodes
Number of used Inodes within a file system.

Free Inodes
Number of free Inodes within a file system.

Volume Group
Volume group to which a logical volume belongs. If the volume is not part of a volume group this column displays “N/A”.

Logical Volume Space (Logical Volume Allocation)
Logical volume space that is used by the logical volume that contains the file system.

Disk Detectable
Indicates “Yes” when a file system's logical volume is identified as having an underlying host disk.

By File System Group

Chart icon
Click this button to generate charts about the capacity of file systems in file system monitoring groups or file system reporting groups.

Magnifying glass icon
Click this button to view more detailed information about the capacity of file systems in a file system monitoring group or file system reporting group.

Creator
ID of the user who created a group.

Group Name of a file system monitoring group or file system reporting group.

File System Space (Capacity)
Amount of file system storage space for the file systems in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems on virtual machines where the Storage Resource agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference that represents the space that is wasted because of the system overhead.

Percent Used Space
Percentage of total file system storage space that is used within a group.

Consumed File System Space (Used Space)
Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems on hypervisors that are assigned to virtual machine servers where the Storage Resource agent is installed if both the virtual machine file systems and hypervisor file systems are within the group.
- File systems on virtual machine servers where the Storage Resource agents are deployed. However, the hypervisor for the virtual machine servers is not discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are not mounted.

File Count
Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column.

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column.

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a group.

Used Inodes
Number of used Inodes within a group.

Free Inodes
Number of free Inodes within a group.

By Cluster

Chart icon
Click this button to generate charts about the used file system space for a cluster.

Magnifying glass icon
Click this button to view more detailed information about the used storage space on file systems that are in a cluster.

Cluster
Name of cluster.

File System Space (Capacity)
Total amount of unique file system storage space in a cluster. Tivoli Storage

Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference that represents the space that is wasted because of the system overhead.

Percent Used Space
Percentage of total file system storage space that is used within a cluster.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the file systems in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

File Count
Total number of files within a cluster. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column.

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by an agent.

When computers are monitored where the Windows operating system is installed, the number of files that are indicated in this column might be less than the number of files that are reported by the operating system.

Directory Count
Total number of directories within a cluster. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column.

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by an agent.

When computers are monitored where the Windows operating system is installed, the number of folders that is indicated in this column might be higher than the number of folders that is reported by the operating system.

Percent Used Inodes
Percentage of used Inodes within a cluster.

Used Inodes
Number of used Inodes within a cluster.

Free Inodes
Number of free Inodes within a cluster.

By Computer

Chart icon
Click this button to generate a chart about the used file system storage space of the computers that are displayed in the report.

Magnifying glass icon
Click this button to view a more detailed report about the used file system storage space on a computer.

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computer where the Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note:

- The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference that represents the space that is wasted because of the system overhead.
- The value that displays in the TOTAL row for this column does not include file systems on virtual machines where the Storage Resource agent is installed.
- The drill-down reports for each computer generate a list of detectable file systems only. Therefore, the total capacity of file systems for each computer drill-down report equals this value.

Percent Used Space
Percentage of total file system storage space that is used within a computer.

Consumed File System Space (Used Space)
Amount of unique file system storage space that is used or consumed by the computer where the Storage Resource agent is installed. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead). Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

The value that displays in the TOTAL row for this column represents the difference between the TOTAL values in the **File System Space** and **Available File System Space** columns. If the calculated value for the TOTAL row is negative, the column displays 0.

The TOTAL row does not add the used space from VMWare computers. See the *VMWare capacity reports* topic in the *Planning* chapter of the *Tivoli Storage Productivity Center Installation and Configuration Guide* for information about how file system and disk capacity are calculated in this report for VMWare computers. This topic is also available in the *Planning* section of the *Tivoli Storage Productivity Center Information Center*.

The value that displays in the TOTAL row for this column does not include the following file systems.

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent is installed.
- File systems on discovered virtual machine servers where the Storage Resource agent is installed. However, the hypervisor for the virtual machine servers is not discovered and the total used space exceeds the capacity of the total file system capacity.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to a computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

File Count
Total number of files within a computer. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column.

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the **File Count** column is 200.

Directory Count
Total number of directories within a computer. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column.

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a computer.

Used Inodes
Number of used Inodes within a computer.

Free Inodes
Number of free Inodes within a computer.

By Computer Group

Chart icon
Click this button to generate charts about the used file system storage space for a computer monitoring group or computer reporting group.

Magnifying glass icon
Click this button to view more detailed information about the file systems in a computer monitoring group or computer reporting group.

Creator
ID of the user who created a group.

Group Name of a computer monitoring group or computer reporting group.

File System Space (Capacity)
Amount of file system storage space for the computers in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems on virtual machines where the Storage Resource agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference that represents the space wasted because of the system overhead.

Percent Used Space
Percentage of total file system storage space that is used within a group.

Consumed File System Space (Used Space)
Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems on hypervisors that are assigned to virtual machine servers where the Storage Resource agent is installed if both the virtual machine file systems and hypervisor file systems are within the group.
- File systems on virtual machine servers where the Storage Resource agent is installed. However, the hypervisor for the virtual machine servers is not discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are not mounted.

File Count
Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column.

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: For groups that contain Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the **File Count** column is 200.

Directory Count
Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column.

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a group.

Used Inodes
Number of used Inodes within a group.

Free Inodes
Number of free Inodes within a group.

By Domain

Chart icon
Click this button to generate charts about the file system used space for a domain.

Magnifying glass icon
Click this button to view more detailed information about the file system used space of computers that are in a domain.

Domain
Name of a domain.

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference that represents the space wasted because of the system overhead.

Percent Used Space
Percentage of total file system storage space that is used within a domain.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server
- File systems that are not mounted.

File Count
Total number of files within a domain. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column.

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a domain. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column.

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a domain.

Used Inodes
Number of used Inodes within a domain.

Free Inodes
Number of free Inodes within a domain.

Network-wide

Chart icon
Click this button to generate a chart about the file system used space in a network.

Magnifying glass icon
Click this button to view more detailed information about the file system used space in a network.

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on

the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference that represents the space wasted because of the system overhead.

Percent Used Space
Percentage of total file system storage space that is used within a network.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server
- File systems that are not mounted.

File Count
Total number of files within a network. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column.

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of

directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a network. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column.

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a network.

Used Inodes
Number of used Inodes within a network.

Free Inodes
Number of free Inodes within a network.

TPC-Wide Storage Space - Consumed File System Space report:

Use these reports to view detailed information about the used storage space in file systems on computers in your environment. You can view this information in different ways by using the following By clauses: by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

Note: The names of some columns are changed. For your reference, the original names of those columns are within parentheses ( ) next to their new names.

By File System

Chart icon
Click this button to generate charts about the used space information for a file system.

Computer
Name of the computer on which a file system is in.

Mount Point
Name or mount point (UNIX/Linux) of the file system (for example, `c:\`, `d:\`, `/opt`, `/export/home`, and so on). The file systems that are displayed in this report include the following information.

- File systems on virtual machines that are monitored by IBM Tivoli Storage Productivity Center.
- File systems on hypervisors that are monitored by Tivoli Storage Productivity Center.
- File systems that become missing after a probe.

Consumed File System Space (Used Space)
Amount of used storage space within a file system. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

The value in the TOTAL row for this column does not include the following file systems.

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent is installed.
- File systems on discovered virtual machine servers where the Storage Resource agent is installed. However, the hypervisor for the virtual machine servers is not discovered and the total used space exceeds the capacity of the total file system capacity.

Detectable
Indicates "Yes" when a file system is identified as being available to the server after a successful probe.

Percent Used Space
Percentage of file system storage space that is used within a file system.

Available File System Space (Free Space)
Amount of unused storage space within a file system.

File System Space (Capacity)
Amount of file system storage space for a file system.

Note:

- The value in the TOTAL row for this column does not include file systems on virtual machines where the Storage Resource agent is installed.
- The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**. The difference represents the space that is wasted because of the system overhead.

File Count
Total number of files within a file system. This value does not include files on file systems that are not scanned.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a file system. This value does not include directories on file systems that are not scanned.

Percent Used Inodes
Percentage of used Inodes within a file system.

Used Inodes
Number of used Inodes within a file system.

Free Inodes
Number of free Inodes within a file system.

Volume Group
Volume group to which a logical volume belongs. If the volume is not part of a volume group this column displays “N/A”.

Logical Volume Space (Allocation)
Logical volume space that is used by the logical volume that contains the file system.

Disk Detectable
Indicates “Yes” when a file system's logical volume is identified as having an underlying host disk.

By File System Group

Chart icon
Click this button to generate charts about the used file system storage space in a file system monitoring group or file system reporting group.

Magnifying glass icon
Click this button to view more detailed information about the used file system storage space in a file system monitoring group or file system reporting group.

Creator
ID of the user who created a group.

Group Name of a file system monitoring group or file system reporting group.

Consumed File System Space (Used Space)
Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems on hypervisors that are assigned to virtual machine servers where the Storage Resource agent is installed if both the virtual machine file systems and hypervisor file systems are within the group.
- File systems on virtual machine servers where the Storage Resource agent is installed. However, the hypervisor for the virtual machine servers is not discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Percent Used Space
Percentage of total file system storage space that is used within a group.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.

- File systems that are not mounted.

File System Space (Capacity)
Amount of file system storage space for the file systems in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems on virtual machines where the Storage Resource agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**. The difference represents the space that is wasted because of the system overhead.

File Count
Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column.

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column.

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a group.

Used Inodes
Number of used Inodes within a group.

Free Inodes
Number of free Inodes within a group.

By Cluster

Chart icon
Click this button to generate charts about the used file system space for a cluster.

Magnifying glass icon
Click this button to view more detailed information about the used storage space on file systems that are in a cluster.

Cluster
Name of cluster.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the file systems in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Percent Used Space
Percentage of total file system storage space that is used within a cluster.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

File System Space (Capacity)
Total amount of unique file system storage space in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**. The difference represents the space that is wasted because of the system overhead.

File Count
Total number of files within a cluster. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column.

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

When computers are monitored where the Windows operating system is installed, the number of files that are indicated in this column might be less than the number of files that are reported by the operating system.

Directory Count
Total number of directories within a cluster. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column.

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

When computers are monitored where the Windows operating system is installed, the number of folders that are indicated in this column might be higher than the number of folders that is reported by the operating system.

Percent Used Inodes
Percentage of used Inodes within a cluster.

Used Inodes
Number of used Inodes within a cluster.

Free Inodes
Number of free Inodes within a cluster.

By Computer

Chart icon
Click this button to generate charts about the used file system storage space of the computers that are displayed in the report.

Magnifying glass icon
Click this button to view a more detailed report about the used file system storage space on a computer.

Computer
Name of the computer on which a file system is in.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computer where the Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Percent Used Space
Percentage of total file system storage space that is used within a computer.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to a computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computer where the Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note:

- The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**. The difference represents the space that is wasted because of the system overhead.
- The value in the TOTAL row for this column does not include file systems on virtual machines where the Storage Resource agent is installed.
- The drill-down reports for each computer generate a list of detectable file systems only. Therefore, the total capacity of file systems for each computer drill-down report equals this value.

File Count
Total number of files within a computer. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column.

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of

directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a computer. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column.

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a computer.

Used Inodes
Number of used Inodes within a computer.

Free Inodes
Number of free Inodes within a computer.

By Computer Group

Chart icon
Click this button to generate charts about the used file system storage space for a computer monitoring group or computer reporting group.

Magnifying glass icon
Click this button to view more detailed information about the file systems in a computer monitoring group or computer reporting group.

Creator
ID of the user who created a group.

Group Name of a computer monitoring group or computer reporting group.

Consumed File System Space (Used Space)
Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems on hypervisors that are assigned to virtual machine servers where the Storage Resource agent is installed if both the virtual machine file systems and hypervisor file systems are within the group.
- File systems on virtual machine servers where the Storage Resource agent is installed. However, the hypervisor for the virtual machine servers is not discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Percent Used Space
Percentage of total file system storage space that is used within a group.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are not mounted.

File System Space (Capacity)
Amount of file system storage space for the computers in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems on virtual machines where the Storage Resource agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**. The difference represents the space that is wasted because of the system overhead.

File Count
Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column.

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: For groups that contain Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column.

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a group.

Used Inodes
Number of used Inodes within a group.

Free Inodes
Number of free Inodes within a group.

By Domain

Chart icon
Click this button to generate charts about the file system used space for a domain.

Magnifying glass icon
Click this button to view more detailed information about the file system used space of computers that are in a domain.

Consumed File System Space (Domain)
Name of a domain.

Used Space
Total amount of unique file system storage space that is used or consumed by the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Percent Used Space
Percentage of total file system storage space that is used within a domain.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server
- File systems that are not mounted.

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**. The difference represents the space that is wasted because of the system overhead.

File Count
Total number of files within a domain. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column.

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a domain. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column.

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a domain.

Used Inodes
Number of used Inodes within a domain.

Free Inodes
Number of free Inodes within a domain.

Network-wide

Chart icon
Click this button to generate charts about the file system used space in a network.

Magnifying glass icon
Click this button to view more detailed information about the file system used space in a network.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Percent Used Space
Percentage of total file system storage space that is used within a network.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server
- File systems that are not mounted.

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**. The difference represents the space that is wasted because of the system overhead.

File Count
Total number of files within a network. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column.

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a network. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column.

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a network.

Used Inodes
Number of used Inodes within a network.

Free Inodes
Number of free Inodes within a network.

TPC-Wide Storage Space - Available File System Space report:

Use these reports to view detailed information about the available storage space in file systems on systems in your environment. You can view this information in different ways by using the following By clauses: by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

Note: The names of some columns are changed. For your reference, the original names of those columns are within parentheses ( ) next to their new names.

By File System

Chart icon
Click this button to generate a chart about the used space information for a file system.

Computer
Name of the computer on which a file system is in.

Mount Point
Name or mount point (UNIX/Linux) of a file system (for example, `c:\`, `d:\`, `/opt`, `/export/home`, and so on). The file systems that are displayed in this report include the following information.

- File systems on virtual machines that are monitored by IBM Tivoli Storage Productivity Center.
- File systems on hypervisors that are monitored by Tivoli Storage Productivity Center.
- File systems that become missing after a probe.

Available File System Space (Free Space)
Amount of unused storage space within a file system.

Detectable
Indicates “Yes” when a file system is identified as being available to the server after a successful probe.

Percent Available File System Space (Percent Free Space)
Percentage of file system storage space that is unused within a file system.

Consumed File System Space (Used Space)
Amount of used storage space within a file system. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

The value in the TOTAL row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent is installed.
- File systems on discovered virtual machine servers where the Storage Resource agent is installed. However, the hypervisor for the virtual machine servers is not discovered yet and the total used space exceeds the capacity of the total file system capacity.

File System Space (Capacity)
Amount of file system storage space for a file system.

Note:

- The value in the TOTAL row for this column does not include file systems on virtual machines where the Storage Resource agent is installed.
- The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**. The difference represents the space that is wasted because of the system overhead.

File Count
Total number of files within a file system. This value does not include files on file systems that are not scanned.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a file system. This value does not include directories on file systems that are not scanned.

Percent Used Inodes
Percentage of used Inodes within a file system.

Used Inodes
Number of used Inodes within a file system.

Free Inodes
Number of free Inodes within a file system.

Volume Group
Volume group to which a logical volume belongs. If the volume is not part of a volume group this column displays “N/A”.

Logical Volume Space (Allocation)
Logical volume space that is used by the logical volume that contains the file system.

Disk Detectable
Indicates "Yes" when a file system's logical volume is identified as having an underlying host disk.

By File System Group

Chart icon
Click this button to generate charts about the used file system storage space in a file system monitoring group or file system reporting group.

Magnifying glass icon
Click this button to view more detailed information about the used file system storage space in a file system monitoring group or file system reporting group.

Creator
ID of the user who created a group.

Group Name of a file system monitoring group or file system reporting group.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are not mounted.

Percent Available File System Space (Percent Free Space)
Percentage of total file system storage space that is not used or available within a group.

Consumed File System Space (Used Space)
Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Storage Resource agent is installed if both the virtual machine file systems and hypervisor file systems are within the group.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems. The total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

File System Space (Capacity)
Amount of file system storage space for the file systems in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Storage Resource agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**. The difference represents the space that is wasted because of the system overhead.

File Count
Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a group.

Used Inodes
Number of used Inodes within a group.

Free Inodes
Number of free Inodes within a group.

By Cluster

Chart icon
Click this button to generate a chart about the used file system space for a cluster.

Magnifying glass icon
Click this button to view more detailed information about the used storage space on file systems that are in a cluster.

Cluster
Name of cluster.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Percent Available File System Space (Percent Free Space)
Percentage of total file system storage space that is available within a cluster.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the file systems in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

File System Space (Capacity)
Total amount of unique file system storage space in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**. The difference represents the space that is wasted because of the system overhead.

File Count
Total number of files within a cluster. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by an agent.

When computers are monitored where the Windows operating system is installed, the number of files that are indicated in this column might be less than the number of files that are reported by the operating system.

Directory Count
Total number of directories within a cluster. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by an agent.

When computers are monitored where the Windows operating system is installed, the number of folders that is indicated in this column might be higher than the number of folders that is reported by the operating system.

Percent Used Inodes
Percentage of used Inodes within a cluster.

Used Inodes
Number of used Inodes within a cluster.

Free Inodes
Number of free Inodes within a cluster.

By Computer

Chart icon
Click this button to generate a chart about the used file system storage space of the computers that are displayed in the report.

Magnifying glass icon
Click this button to view a more detailed report about the used file system storage space on a computer.

Computer
Name of the computer on which a file system is in.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to a computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Percent Available File System Space (Percent Free Space)
Percentage of total file system storage space that is available within a computer.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computer where the Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.

- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computer where the Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note:

- The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference represents the space that is wasted because of the system overhead.
- The value in the TOTAL row for this column does not include file systems on virtual machines where the Storage Resource agent is installed.
- The drill-down reports for each computer generate a list of detectable file systems only. Therefore, the total capacity of file systems for each computer drill-down report equals this value.

File Count
Total number of files within a computer. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a computer. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.

- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a computer.

Used Inodes
Number of used Inodes within a computer.

Free Inodes
Number of free Inodes within a computer.

By Computer Group

Chart icon
Click this button to generate charts about the used file system storage space for a computer monitoring group or computer reporting group.

Magnifying glass icon
Click this button to view more detailed information about the file systems in a computer monitoring group or computer reporting group.

Creator
ID of the user who created a group.

Group Name of a computer monitoring group or computer reporting group.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are not mounted.

Percent Available File System Space (Percent Free Space)
Percentage of total file system storage space that is available within a group.

Consumed File System Space (Used Space)
Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Storage Resource agent is installed if both the virtual machine file systems and hypervisor file systems are within the group.
- File systems on virtual machine servers where the Storage Resource agent is installed. However, the hypervisor for the virtual machine servers is not discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

File System Space (Capacity)
Amount of file system storage space for the computers in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Storage Resource agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference represents the space that is wasted because of the system overhead.

File Count
Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: For groups that contain Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a group.

Used Inodes
Number of used Inodes within a group.

Free Inodes
Number of free Inodes within a group.

By Domain

Chart icon
Click this button to generate charts about the file system used space for a domain.

Magnifying glass icon
Click this button to view more detailed information about the file system used space of computers that are in a domain.

Domain
Name of a domain.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server
- File systems that are not mounted.

Percent Available File System Space (Percent Free Space)
Percentage of total file system storage space that is available within a domain.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference represents the space that is wasted because of the system overhead.

File Count
Total number of files within a domain. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a domain. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a domain.

Used Inodes
Number of used Inodes within a domain.

Free Inodes
Number of free Inodes within a domain.

Network-wide

Chart icon
Click this button to generate a chart about the file system used space in a network.

Magnifying glass icon
Click this button to view more detailed information about the file system used space in a network.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server
- File systems that are not mounted.

Percent Available File System Space (Percent Free Space)
Percentage of total file system storage space that is available within a network.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This difference occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference represents the space that is wasted because of the system overhead.

File Count
Total number of files within a network. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that are not scanned by a Storage Resource agent.

Note: On Windows computers, AIX systems, and systems that use IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a network. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.

- Directories on file systems that are not scanned by a Storage Resource agent.

Percent Used Inodes
Percentage of used Inodes within a network.

Used Inodes
Number of used Inodes within a network.

Free Inodes
Number of free Inodes within a network.

Usage Reports:

Use Usage reports to view the usage and growth of the consumption of your storage. You can view these reports that are based on specific file systems and computers, groups of file systems and computers, or throughout your entire enterprise. These reports help you project storage consumption for the future and maximize the current storage assets that you have in place. This result is achieved by eliminating wasted space and making the most of the space you have.

Use these reports to do the following actions.

- View which servers and file systems are experiencing the heaviest (or lightest) load of data and storage access.
- Identify wasted space by pinpointing files that are no longer needed or not accessed in a long time.
- Associate users with the amount of storage they are using on specific file systems, groups of file systems, specific computers, groups of computers, and throughout the entire network.
- Select the files against which to run an IBM Tivoli Storage Manager archive/backup job. Tivoli Storage Manager protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on auxiliary storage.

Access Load reports: Use Access Load reports to view which machines and file systems within your environment are experiencing the heaviest (or lightest) load of data and storage access.

Access Time reports
Use Access Time reports to view and chart the amount of data and the number of files that have been accessed during the last day, the last week, the last month, the last year, and over. You can view this information at the directory level, the file system level, the computer level, at the domain level, and for the entire network.

The following report shows a single pie chart that depicts all of the data across the entire enterprise, giving you an easy way to see how much of your corporate data is accessed every day, every week, and even data that has not been accessed in over a year.

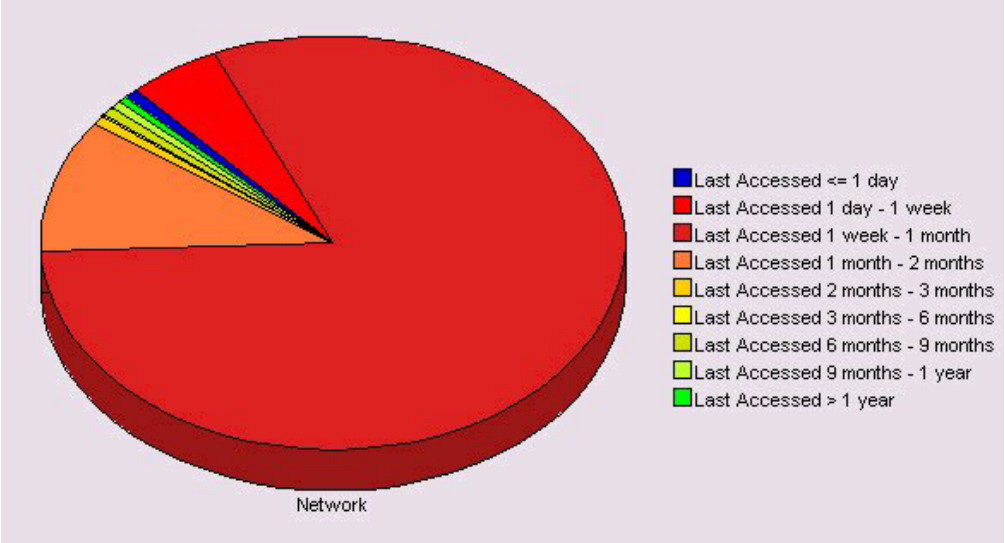

Figure 39. Usage Reporting—Access Load—Access Time—Chart Network Wide

Modification Time reports

Use Modification Time reports to view and chart the amount of data and the number of files that have been modified during the last day, the last week, the last month, the last year, and over. You can view this information at the directory level, the file system level, the computer level, at the domain level, and for the entire network.

The following report shows a single pie chart that depicts all of the data across the entire enterprise, allowing you to see how much of your corporate data is modified every day, every week, and even data that has not been modified in over a year.

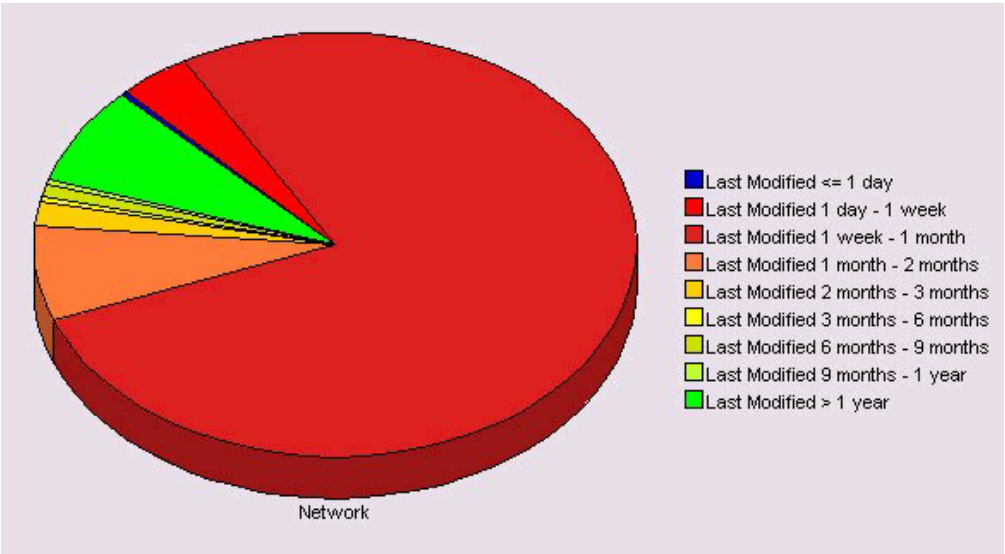

Figure 40. Usage Reporting—Access Load—Modification Time—Chart Network Wide

Creation Time reports

Use Creation Time reports to view and chart the amount of data and the number of files that have been created during the last day, the last week,

the last month, the last year, and over. You can view this information at the directory level, the file system level, the computer level, at the domain level, and for the entire network.

Note: Because scans do not collect creation time information for the files stored on UNIX machines, creation time reports do not contain any data for the non-Windows machines in your environment.

Files reports: Use Files Reports to report on the files found on the storage in your environment during theData Manager scan process. During the scan process, Data Manager gathers statistics on your files and attributes about the files. Select the files against which to run an IBM Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 547 for information on how to create archive and backup jobs from the files shown in a Files report. The information gathered can be reported on in a variety of ways using this facility:

Largest Files reports

Use these reports to view information about the largest files in your environment. Largest Files reports enable you to do the following:

- Generate reports that contain detailed information about the largest files found in your environment. You can view reports on the largest files in the following ways: by directory, by directory group, by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

  Note: The number of files that are gathered and stored in the repository is dependent upon the number that you specify in the profile that you use during scanning. The default profile gathers information on the 20 largest files. You can select a profile for the report by clicking the **Profile** list box.
- Select files against which to run anTivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 547 for information on how to create archive/backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive/backup job. See "Archive and backup functions" on page 547 for information on how to add files from reports to an archive/backup job.

Duplicate Files reports

Use these reports to view information about the files found in your environment during a scan that contain duplicate file names. Duplicate Files reports enable you to do the following:

- Generate reports on the files found in your environment during a scan that contain duplicate file names. Use this report to view a list of files that might no longer be needed and could be wasting storage space. You can use the data generated in a duplicate files report to view a list of files that might no longer be needed and could be wasting storage space. You can view reports on duplicate files in the following ways: by directory, by directory group, by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

  Note: The number of files that are gathered and stored in the repository is dependent upon the number that you specify in the profile that you

use during scanning. The default profile gathers information on the 20 largest files. You can select a profile for the report by clicking the **Profile** list box.

- Select files against which to run an Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 547 for information on how to create archive and backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 547 for information on how to add files from reports to an archive or backup job.

Before generating a Duplicate Files report for environments that contain a large amount of files (i.e., environments that might contain over 100,000 duplicate files), we recommend you tune theData Manager repository. To do this, run the appropriate RDBMS utility to collect statistics on the repository tables. Run the following utilities depending on the RDBMS where the repository is stored:

- UDB: runstats
- Oracle: analyze
- Microsoft SQL Server: update stats
- Sybase: update stats

Run the appropriate utility against all the tables in the repository schema you specified during the installation process.

To generate the data for a duplicate files report, you can either create a constraint that will save filenames for *n* number of violation files, or run a scan that includes one or all of the following existing profiles:

- TPCUser.Largest Files
- TPCUser.Largest Orphans
- TPCUser.Most at Risk
- TPCUser.Most Obsolete

There is not a specific Duplicate Files profile that collects information on all duplicate files in a file system. The filenames that are triggered by the previously listed profiles or constraints will be saved to the repository and the duplicate files report is generated from the data gathered from the scan or constraint.

For example, create a scan using the Largest Files profile to collect information on the 100 largest files. The 100 largest files are saved in the t_stat_file table in the repository. If five duplicate files were saved as a result of being flagged by the Largest File Profile, then those files will be reported as duplicates in the Duplicate file Report.

Most Obsolete Files reports

Use these reports to view information about the files found in your environment that have not been accessed in the longest period of time. Most Obsolete Files reports enable you to do the following:

- Generate reports on the files found in your environment that have not been accessed in the longest period of time. It is very likely that many of these files are no longer needed, at least online, and could be wasting space. You can view reports on the largest files in the following ways: by directory, by directory group, by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

Note: The number of files that are gathered and stored in the repository is dependent upon the number that you specify in the profile that you use during scanning. The default profile gathers information on the 20 largest files. You can select a profile for the report by clicking the **Profile** list box.

- Select files against which to run an Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 547 for information on how to create archive/backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 547 for information on how to add files from reports to an archive or backup job.

Orphan Files reports

Use these reports to view information about files that are owned by users that are no longer active in the environment. Orphan files reporting enables you to do the following:

- Generate reports on files that are owned by users that are no longer active in the environment. It is very likely that many of these files are no longer needed and could be wasting space. You can view reports on orphan files in the following ways: by directory, by directory group, by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

  You can also generate reports on the oldest orphan files that have not been accessed for the longest periods of time, as well as largest orphan files.

  Note: The number of files that are gathered and stored in the repository is dependent upon the number that you specify in the profile that you use during scanning. The default profile gathers information on the 20 largest files. You can select a profile for the report by clicking the **Profile** list box.
- Select files against which to run an Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 547 for information on how to create archive and backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 547 for information on how to add files from reports to an archive or backup job.

File Size Distribution reports

Use these reports to view information about the distribution of file sizes across storage resources. Use File Size Distribution reports to generate reports that show the distribution of file sizes on a specific directory, a file system, a computer, the domain, or across your entire network.

The following report shows a single pie chart for each computer in the report.

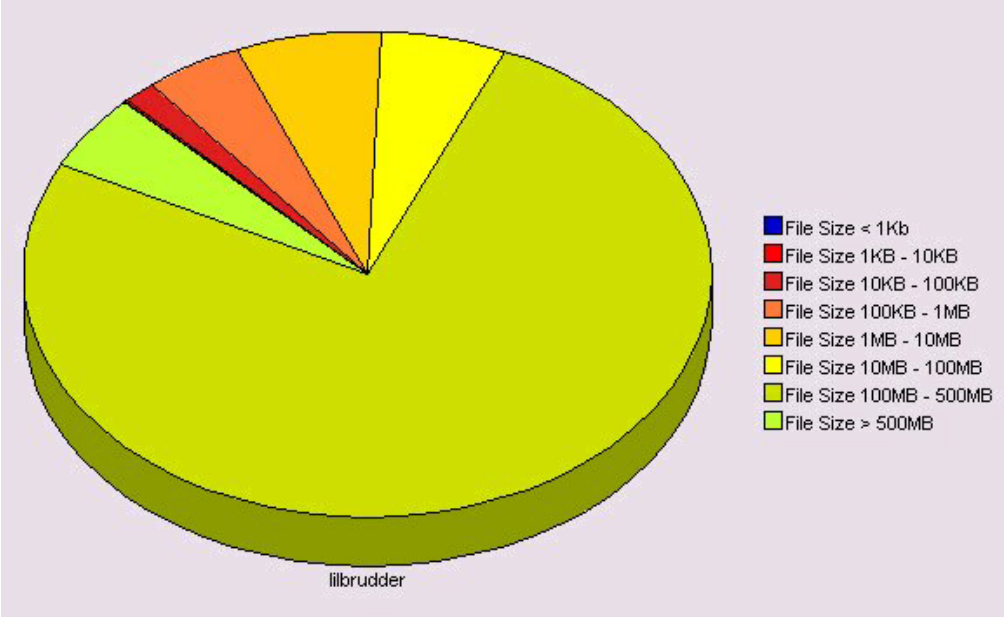

Figure 41. Usage Reporting—Files—File Size Distribution—File Count Chart Tab

Each pie chart depicts the number of files that fall into each file size bucket that are described in the legend. For example, the blue portion of the pie represents the number of files on each computer that are smaller than 1024 bytes long. This gives you an easy way to get a quick grasp of what are the size distribution of your data files.

File Summary reports

Use these reports to view general information about the files in your environment. Use File Summary reports generate overview information for files, also at the directory level, the file system level, the computer level, the domain level or for the whole network.

This report shows you the following (for each entity at the level you are reporting by):

- The total size of the files
- The number of files
- The number of directories
- The average file size

File Types reports

Use these reports to view information about the storage usage of different file types in your environment. File Types reports enable you to generate reports on the file types found in your environment during a scan. The scan that collects data for these reports must use the TPCUser.Summary By File Type profile. For example, when you run a scan using the TPCUser.Summary By File Type profile,Data Manager will collect storage usage information about such file types as .exe, .zip, .sys, .pdf, .doc, .dll, .wav, .mp3, .avi, etc. Each of these file types is represented by its own row in the generated reports.

Use these reports to do the following:

- Relate the space used by applications to the total capacity and used space. For example, you can use these reports to view the total amount of storage consumed by Adobe Acrobat files and Lotus Notes mail databases.
- View which applications are consuming the most space across a given set of storage resources

You can view these reports in the following ways:

- **Group by Resource**. These reports organize data by storage resource (for example, file system, computer, cluster, etc.) and space used.
- **Group by File Type**. These reports organizes data by file type and space used. They show where the most spaced used by each type of file resides.

Archive/Backup jobs from reports:

Use the Archive/Backup function from Reporting to define an archive or backup job to run against the files that you select from specified reports.

Tivoli Storage Manager protects your organization's data from hardware failures and other let you back up specified files. These files are not excluded in the include-exclude list and meet the requirement for serialization in the backup copy group of the management class assigned to each file. An incremental backup enables you to back up new or changed files or directories from a client domain or from specified directories or files. These directories or files are not excluded in the include-exclude list and meet the requirements for frequency, mode, and serialization as defined by a backup copy group of the management class assigned to each file. See *IBM Tivoli Storage Manager Client User's Guide* for more information about Tivoli Storage Manager archives and backups.

This function enables you to select a specific file or group of files fromData Manager reports that you want to archive or back up using the Tivoli Storage Manager tool.

Archive/backup is run from**Data Manager -> Reporting**. You cannot create archive and backup jobs from the Archive/Backup node - rather you create these jobs by selecting files from the previously listed reports or the reports that you create. Each archive and backup that you create from reports appears as a node under the Archive/Backup node. The reports that contain files against which you can run an archive or backup job include:

- **Data Manager -> Reporting -> Usage -> Files -> Largest Files**
- **Data Manager -> Reporting -> Usage -> Files -> Most Obsolete Files**
- **Data Manager -> Reporting -> Usage -> Files -> Orphan Files**
- **Data Manager -> Reporting -> Usage -> Files -> Duplicate Files**
- **Data Manager -> Reporting -> Backup -> Most at Risk Files ->**
- **Data Manager -> Reporting -> Usage Violations -> Constraint Violations**

You can also run an archive or backup job from a report you create.

Note:

See *Creating Archive/Backup jobs* in *Reporting* for information about how to create an archive or backup job from the files shown in reports.

Directories reports:
Use Directories reports to report upon the directories found on the storage in your environment during the scan process. During the scan process, Data Manager gathers a vast number of statistics on your directories and a number of attributes about the directories. Before you can view information about directories in these reports, you must do the following:

1. Use the **Monitoring → Groups → Directory** function to create directory groups that contain the directories upon which you want to report.
2. Include the directory groups you created in scans.
3. Schedule and run the scans that contain the directory groups (and directories) upon which you want to report.

Largest Directories reports

Use Largest Directories reports to view reports on the *n* largest directories found in your environment:

- for specific directories
- on specific computers or groups of computers
- on specific file systems or groups of file systems
- in specific domains
- on the entire network

The number of largest (or the value of *n* in finding the *n* largest directories), is determined by the profile that you use for doing the reporting. When you define a profile in the **Monitoring → Profiles** node, you specify the number of largest directories upon which to collect information. When you generate a Largest Directories report, you can select what profile to use.

Monitored Directory Space reports

Use Monitored Directory reports to report on the monitored directories that you have included in a group via the **Monitoring → Groups → Directory** facility and included in a scan. During the scan process, Data Manager gathers a number of statistics on your directories and a number of attributes about the directories.

Users reports: Use Users Reports to view information about the users that are found as file owners in your environment. These reports provide an easy and convenient way to view and aggregate storage resource consumption by owner across directories, file systems, computers, domains, and across your entire network.

OS User Groups reports: OS User Groups reports give you the exact same capabilities as the reports under Users reports. The only difference is that you are doing your analysis of OS User Groups as opposed to actual single users.

Usage Violation reports:

Use Usage Violation reports to view information about quota and constraint violations.

With Usage Violation reports, you can do these actions:

- Create and enforce corporate storage usage policies that are based on reports that show quota and constraint violations.

  Use quotas to control how much storage a user, or a group of users, can use on a file system or group of file systems, a computer or group of computers, and

throughout the entire network. Use Constraints to set policies about which types of files, owners, and file sizes are allowed on a file system or group of file systems, a computer or group of computers, and throughout the entire network.

- Select files in Constraint Violation reports against which to run an IBM Tivoli Storage Manager archive/backup job. See "Archive and backup functions" on page 547 for information on how to create archive/backup jobs from the files that are shown in a Files report.
- Add files from Constraint Violation reports to an existing Tivoli Storage Manager archive or backup job. See "Create an archive/backup job" on page 548 for information on how to add files from reports to an archive/backup job.

Constraint Violations reports: Constraint Violations reporting enables you to do the following:

- View the files that have violated a constraint. The report shows one row for each Constraint. The row shows the detail of the constraint and the violating space totals for the files that violate the constraint. You can view reports on the files that violate constraints in the following ways: by file system, by file system group, by computer, by computer group, by domain, and by constraint.
- Select files against which to run an Tivoli Storage Manager archive or backup job. See the "Archive and backup functions" on page 547 section for information on how to create archive and backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive or backup job. See the "Archive and backup functions" on page 547 section for information on how to add files from reports to an archive or backup job.

Quota Violations reports: Quota Violations reports display information on the quotas that you define. Monitor the space that is being used by the various entities that have quotas that are defined upon them.

Quotas can be defined on either of two categories of users: on a specific user or group of users that you define within either Users or on OS User Groups. The two categories are described in the following list:

User You can define a quota on a user or group of users. The users are the owners of the files that Data Manager finds as it scans your systems.

OS User Group
You can define a quota on an OS user group or group of OS user groups. OS user groups are defined outside the tool and are found by Data Manager as it scans your systems.

Within each of those categories, you can define quotas at any one of following three levels:

Network
You can define quotas on users or groups of users at the entire network level. You can control how much storage any single user or group of users is using throughout your enterprise.

Computer
You can define quotas on users or groups of users at the computer level. You can control how much a single or group of users is using on a single computer or on any group of computers.

File system
You can define quotas on users or groups of users at the file system level.

You can control how much a single or group of users is using on a single file system, or on any group of file systems.

All Quotas reports

Use All Quotas reports to report on all quotas that are violated, including the quotas at the network level, computer level, and file system level. With the **Selection** tab, you can select which users are to do the quota reporting.

The **Objects** tab displays information for each user and computer and file system:

- How much space is being used
- There
- What the Quota threshold is
- What the percentage of the Quota that is being used

This report gives you a simple way to analyze how much space is being used relative to the Quotas that you defined on your systems and users.

Network-wide Quotas reports

Use Network-wide Quota reports to view all the network-wide quotas that are violated by users and OS user groups in your environment. Quota information that is displayed in these reports includes the amount of space that is being used, the threshold value, and the percentage over the threshold value. You can view these reports that are sorted by user, by OS user group, and by the creator of the quota.

Computer Quotas reports

Use Computer Quota reports to view all the computer quotas that are violated by users and OS user groups in your environment. Quota information that is displayed in these reports includes the amount of space that is being used, the threshold value, and the percentage over the threshold value. You can view these reports that are sorted by user, by OS user group, and by the creator of the quota.

File System Quotas reports

Use File System Quota reports to view all the computer quotas that are violated by users and OS user groups in your environment. Quota information that is displayed in these reports includes the amount of space that is being used, the threshold value, and the percentage over the threshold value. You can view these reports that are sorted by user, by OS user group, and by the creator of the quota.

Backup reports:

Use Backup reports to identify the most at risk files in your enterprise that are not backed up properly. Critical data that is not backed up might expose you to costly losses of information.

Backup reports also enable you to:

- Investigate the sizes of data that would be required to backup to achieve full backups of any subset of your enterprise (including the whole enterprise).
- Determine the amount of space required to perform an incremental backup of any or all of your data at any given point in time.

- Select files in Most at Risk Files reports against which to run an IBM Tivoli Storage Manager archive/backup job.
- Add files from Most at Risk Files reports to an existing Tivoli Storage Manager archive/backup job.

Most At Risk Files reports: Most At Risk Files reporting enables you to do the following:

- View the files that have the earliest modification time, but have not been backed up since they were modified. This can alert you to a dangerous situation where you may have vulnerable data that could be lost and not recovered due to lack of backups.
- Select files against which to run an IBM Tivoli Storage Manager archive or backup job. See the "Archive and backup functions" on page 547 section for information on how to create archive or backup jobs from the files shown in a Files report.
- Add files to an existingTivoli Storage Manager archive or backup job. See the "Archive and backup functions" on page 547 section for information on how to add files from reports to an archive or backup job.

Modified Files not Backed Up reports: Use Modified Files not Backed Up reports to view the numbers of files that have the earliest modification time, but have not been backed up since they were modified. This can alert you to a dangerous situation where you may have vulnerable data that could be lost and not recovered due to lack of backups.

Backup Storage Requirements reports: Use Backup Storage Requirements reports to determine, at a network, domain computer, file system, or directory level, how much storage will be required to do backups. You can determine how much storage will be required to do both full backups and to do incremental backups.

Reporting Type	Description
Full Backup Size	Use Full Backup Size reports to view how much space it will take to do full backups of specific computers, file systems, directories, groups of those entities, domains, and the entire network.
Incremental Backup Size	(Windows only) Use this report to view the total size of all the files that have not been backed up since the last scan. You can generate this report by Directory, by Directory Group, by File System, by File System Group, by Computer, by Computer Group, by Domain and Network Wide. This allows you to see for any of these levels, how much data would need to be backed up for any of these desired time intervals. This is useful for backup sizing and can be used to determine appropriate backup groups and sets for timing your backups.

Reporting Type	Description
Incremental Range Sizes	Use this report to determine how much storage space you would need to back up data that has been modified in the last day, week, month, two months, three months, six months, nine months, or year. For example, if you want to back up the files that have been modified in the last 3 months only, generate this report and view the values in the "Last Modified <= 3 months" column to determine 1) how many files changed and 2) how much storage those changed files are consuming. You can generate this report by Directory, by Directory Group, by File System, by File System Group, by Cluster, by Computer, by Computer Group, by Domain and Network Wide. This allows you to see for any of these levels, how much data would need to be backed up for any of these desired time intervals. This is useful for backup sizing and can be used to determine appropriate backup groups and sets for timing your backups.

Monitored Computer Storage Space reports:

Use these host-based reports to view the data that is collected by probes for computers, network attached storage, and hypervisors. These reports contain detailed information about the consumption of file system space and enable you to relate that space back to external or internal storage without counting shared space multiple times.

Monitored computer storage space reports are organized into the following categories:

Disk Storage
Use these reports to view information about disks that reside on computers monitored by IBM Tivoli Storage Productivity Center. This includes information about storage subsystem disks, known disks with serial numbers, or unknown disks without serial numbers.

Non-Disk Storage
Use these reports to view information about storage space that does not reside on computer disks that are monitored by Tivoli Storage Productivity Center. This type of storage includes the following:

- Storage that resides on a remote mount or a network attached server.
- Storage that is temporarily allocated to processor memory.

You must perform the following actions before viewing these reports:

- Run a probe against the storage entities upon which you want to report.
- If you want to generate By Computer Group, By Filesystem Group, and By Storage Subsystem Group reports, perform the following tasks:
  - Create computer monitoring group and file system monitoring groups using the **Data Manager** > **Monitoring** > **Groups** node and then probe them.
  - Create storage subsystem monitoring group using the **Disk Manager** > **Monitoring** > **Groups** node and probe them.
  - Create a computer reporting group or file system reporting group using the **Data Manager** > **Reporting** > **Groups** and **Disk Manager** > **Reporting** > **Groups** node. Probe the entities in the group.

– Create a subsystem reporting group using the **Disk Manager** > **Reporting** > **Groups** node. Probe the storage subsystems in the group.

Related Topics

- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Monitored Computer Storage Space reports - Storage entities:

Monitored Computer Storage Space reports display information about storage entities that you probe. Use this topic to view descriptions of those entities.

You can use Monitored Computer Storage Space reports to view information about the following storage entities:

Computer
Computer that is monitored by a Storage Resource agent. This computer can be a single host, a virtual machine, or a hypervisor. A hypervisor is a machine that hosts one or more virtual machines.

Disk on a storage subsystem
Computer disk that originates from a storage subsystem. If the disk's storage subsystem is not monitored by IBM Tivoli Storage Productivity Center, that storage subsystem is represented by a portion of the serial number of the disk. If the storage subsystem is monitored by Tivoli Storage Productivity Center, that storage subsystem is represented by the name collected by a probe.

Disk not on a storage subsystem with a serial number
Computer disk that Tivoli Storage Productivity Center is unable to determine if it originates from a storage subsystem. Tivoli Storage Productivity Center is able to extract the serial number of that disk.

Disk not on a storage subsystem without a serial number
A valid computer disk for which Tivoli Storage Productivity Center is unable to do the following:

- Determine if the disk originates from a storage subsystem.
- Extract the serial number of the disk.

Disk Group
A collection of disk drives on a computer.

File System
Storage entity that creates, removes, and modifies files on a computer. File systems can be shared among monitored computers and span portions of different disks. If a file system spans disk entities, then there will be a single logical volume beneath it that describes how the file system is spanned.

(Parent) Logical Volume
Logical representation of a volume that spans many logical volumes. If there is no file system created on this logical volume, information about the parent logical volume is displayed in the reports.

(Child) Logical Volume
Logical representation of a volume that spans one or more portions of a single or multiple disks. This entity helps determine how the parent logical

volumes and file systems are allocated on the disks. If there is no file system created on this logical volume, or there is no parent logical volume for this logical volume, the child logical volume is displayed in the reports.

Non-Disk Storage
Storage space that does not reside on a computer disk. This includes storage residing on a remote mount or a network attached server and storage that is temporarily allocated to processor memory.

Volume Group
A collection of logical volumes on a computer.

Virtual Machine (VM) Disk File
Storage entity that represents a computer disk on a computer that is a virtual machine.

Virtual Machine File System (VMFS)
Storage entity that creates, removes, and modifies VM Disk files on a computer that is a hypervisor hosting one or more virtual machines.

Monitored Computer Storage Space reports - Common storage space values:

Monitored Computer Storage Space reports provide detailed storage space information about the computers probed by IBM Tivoli Storage Productivity Center. Use this topic to view descriptions of the common storage values that appear in these reports.

You can view the following storage space information in Monitored Computer Storage Space reports:

Space Total amount of storage space for the entity that a row represents:

- For file system rows, this is the total logical volume space where a file system resides.
- For logical volume rows, this is the total logical volume space.
- For computer disk rows, this is the total disk space of a computer disk.
- For computer space rows, this is the total storage subsystem volume space on a computer.
- For storage subsystem space rows, this is the total storage subsystem volume space on all monitored computers.

Consumed Space
Amount of consumed storage space for the entity that a row represents:

- For file system rows, this is the space that is consumed by data that can be accessed by the end user of a computer.
- For logical volume rows, this is the space that can be consumed by a file system, but is not yet consumed by a file system.
- For computer disk rows, this is the space that is consumed by the file systems and logical volumes on a disk.
- For computer space rows, this is the storage subsystem volume space that is consumed on a monitored computer.
- For storage subsystem space rows, this is the storage subsystem volume space that is consumed on all monitored computers.

Available Space
Amount of available storage space for the entity that a row represents:

- For file system rows, this is the space that is available for consumption.

- For logical volume rows, this value is zero (0) because no file system has been created on the logical volume that would help determine the amount of available space.
- For computer disk rows, this is the space that is available on the file systems on a computer disk and the space on a computer disk not yet allocated to a logical volume.
- For computer space rows, this is the storage subsystem volume space that is available for consumption on a computer.
- For storage subsystem space rows, this is the storage subsystem volume space that is available for consumption on all monitored computers.

Unavailable Space

Amount of unavailable storage space for the entity that a row represents:

- For file system rows, this space is unavailable for consumption because it is reserved by the software RAID level on the underlying logical volume and is required for data that is used by the operating system to maintain a file system.
- For logical volume rows, this space is unavailable because it is reserved for the software RAID level of a logical volume.
- For computer disk rows, this is the space that is unavailable on the file systems and the logical volumes on a computer disk.
- For computer space rows, this is the storage subsystem volume space that is unavailable on a computer.
- For storage subsystem space rows, this is the storage subsystem volume space that is unavailable on all monitored computers.

Monitored Computer Storage Space - Disk Storage reports:

Use these reports to view information about disks residing within computers that are monitored by IBM Tivoli Storage Productivity Center. This includes information about storage subsystem disks, known disks with serial numbers, and unknown disks without serial numbers.

Disk Storage reports are organized into the following categories:

All Disk Storage

Use these reports to view information about all the disks residing on computers that are monitored by Tivoli Storage Productivity Center.

Disk Storage On Storage Subsystems

Use these reports to view information about disks that Tivoli Storage Productivity Center has identified as storage subsystem volumes.

Disk Storage Not On Storage Subsystems

Use these reports to view information about disks that Tivoli Storage Productivity Center has not identified as storage subsystem volumes.

Related Topics

- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Disk Storage - All Disk Storage reports:

Use these reports to view information about all the disks residing on computers that are monitored by IBM Tivoli Storage Productivity Center.

See All Disk Storage reports - Column Descriptions for descriptions of all the columns that appear in All Disk Storage reports.

Use By clauses to filter the entities that appear in a report. For example, if you want to view an All Disk Storage report that displays specific computers, select **By Computer** and use the **Selection** and **Filter** options to determine which computers appear in that report.

You can customize the information in each report by relating the monitored entities to other storage. The following options for relating entities to storage in these reports includes the following:

- Use the **Relate To:** options in the By Computer, By Computer Group, and By Disk/Volume Group reports to relate storage entities to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space
  - Computer Disk Space
  - Distributed Computer Disk Space
- Use the **Relate To:** options in the By File System/Logical Volume and By File System Groups reports to relate storage entities to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space

Related Topics

- All Disk Storage reports - Column Descriptions
- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

All Disk Storage reports - Column Descriptions:

Use these reports to view information about all the disks residing on computers that are monitored by IBM Tivoli Storage Productivity Center. Use this topic to view descriptions of the columns that appear in these reports.

The following buttons and columns appear in the By clauses of All Disk Storage reports:

Note: Not all of the following columns appear in each By clause.

Selection...
Click this button to select the specific entities that you want to include in the report. For example, in the **By Filesystem/Logical Volume** report, click this button to select the specific file systems or logical volumes upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict the entities that you want to include in the report.

Relate <entity> To:
Select how you want to relate monitored storage entities in the report to other storage.

Magnifying glass icon
This button appears on all **By <entity> Group** reports. Click this button to view a more detailed report about the entities in a group.

Available Disk Space
Amount of space on a disk that is available for use. For computer groups, this is the amount of space on disks that is available for use by the computers in a group. For disk and volume groups, this is the external disk space available for use within those groups.

This space includes the file system space that is available on the disk and the space that is not allocated to a logical volume. For example, unpartitioned space.

The XIV system GUI and the XIV system command-line interface (CLI) report sizes in GB of disks, volumes, pools, and so on, based on 10^9. Tivoli Storage Productivity Center reports sizes in GB based on 2^30. For example, a volume showing a size of 17 GB in the XIV system GUI is reported as 16 GB (XIV GB size * 1000 * 1000 * 1000 / 1024 /1024 / 1024 rounded to the nearest GB) in Tivoli Storage Productivity Center.

Available Disk/Volume Group Space
Amount of available space in a disk group or volume group.

Available File System/Logical Volume Space
Amount of space on the file system or logical volume that is available for use. For computers groups, this is the amount of space on file systems or logical volumes that is available for use by the computers within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that can be consumed. For disk groups and volume groups, this is the external logical volume or file system space available for use within those groups.

Computer
Name of a computer where a disk, file system, logical volume, or disk group is located.

Computer OS Type
The operating system running on a computer.

Consumed Disk Space
Amount of space on a disk that is consumed. For computer groups, this is the amount of disk consumed by the computers within those groups. For disk and volume groups, this is the external disk space consumed within those groups.

This space includes the file system space that is consumed on a disk. If a logical volume on this disk space has no file system, then this space includes the consumed portion of the logical volume that has been allocated to the disk.

Consumed Disk/Volume Group Space
Amount of consumed disk or volume group space.

Consumed File System or Logical Volume Space
Amount of space on the file system or logical volume consumed by the computer on a disk. For computer groups, this is the amount of space on

file systems or logical volumes that is consumed by the computers within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that is consumed. For disk groups and volume groups, this is the external logical volume or file system space consumed within those groups.

Correlated
Indicates **Yes** if Tivoli Storage Productivity Center is able to relate the host storage directly to a storage subsystem. You can relate this correlated storage back to the actual storage subsystem storage for a more detailed view of how the host storage is allocated.

If a computer disk, file system, or logical volume appears in the Disk Storage On Storage Subsystem reports but is not correlated, it might be the result of the following reasons:

- The storage subsystem was not probed. Run a probe to collect data about the storage subsystem then generate this report.
- The storage subsystem storage was unassigned from the host storage, but the host has not been probed. Run a probe to collect data about the host then generate this report.
- The storage subsystem storage was not correctly identified by a probe of the host. If this is the problem, check the probe log of the host for messages relating to SCSI command errors and the probe results of the storage subsystem.

Creator
ID of the user who created a group.

Disk Path
Path name of a disk as seen by the computer.

Disk Space
Total amount of space allocated to a computer disk. For computer groups, this is the total amount of space allocated to the disks on the computers within those groups. For disk and volume groups, this is the external disk space within the disk group or volume group.

Note:

- This space is distributed across the hypervisor disk space from where the virtual machine disk's file resides. The portion of space identified to a particular external hypervisor disk is reported only. This is true for all space values in this report.
- The XIV system GUI and the XIV system command-line interface (CLI) report sizes in GB of disks, volumes, pools, and so on, based on 10^9. Tivoli Storage Productivity Center reports sizes in GB based on 2^30. For example, a volume showing a size of 17 GB in the XIV system GUI is reported as 16 GB (XIV GB size * 1000 * 1000 * 1000 / 1024 /1024 / 1024 rounded to the nearest GB) in Tivoli Storage Productivity Center.

Disk/Volume Group Name
Name of a disk group or volume group where a logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Disk/Volume Group Space
Total amount of disk group or volume group space. This is space inclusive of all storage space and applies to all space values related to disk groups or volume groups.

Disk/Volume Group Type
Type of a disk group.

Fibre Attached
Indicates **Yes** when Tivoli Storage Productivity Center determines that a disk is behind a fibre channel adapter port.

File System Type
Type of a file system, such as NTFS.

File System/Logical Volume Path
Name of a file system or logical volume path.

Group Name of a group.

Grown Defects
Number of defects detected on the disk since it was new.

Hypervisor Name
Name of the hypervisor where a virtual machine was created. This column displays **N/A** if not applicable to a disk.

Logical Volume Space
Total amount of space allocated to the logical volume on a disk. For computer groups, this is the total amount of space allocated to the logical volumes on disks within those groups. For file system groups, this is the total amount of space allocated to the logical volumes within those groups. For disk groups and volume groups, this is the external logical volume space within those groups.

This space includes the file system space on the disk reported if the entity displayed is a file system. The actual size of the file system on the disk reported is in addition to the consumed and available space.

Note: When relating a monitored entity to distributed file system/logical volume space, the space is distributed across the disks where a logical volume or file system resides. The portion identified to a disk is reported only.

(Disk) Manufacturer
The manufacturer of a disk.

Mapped to Storage Subsystem
Indicates **Yes** if Tivoli Storage Productivity Center has identified a disk as being located on a storage volume from a storage subsystem. The storage subsystem might no be monitored by Tivoli Storage Productivity Center.

(Disk) Model
Model number of a disk.

Number of Disks
Number of disks in a computer group. For disk groups and volume groups, this is the number of external disks within those groups.

Number of File Systems/Logical Volumes
Number of logical file systems or logical volumes within a group. For disk groups and volume groups, this is the number of external logical file systems or logical volumes within those groups.

The logical volumes underlying the file systems are not counted for the value in the column.

Overallocated
Indicates **Yes** if Tivoli Storage Productivity Center is unable to uniquely identify disk storage across more than one host. Check the probe log of the host for messages relating to SCSI command errors.

Path Type
Path type of the entity being reported: file system or a logical volume.

Primary Defects
Number of defects on the disk when it was new.

Shared
Indicates **Yes** if a file system or logical volume is shared by more than one computer.

Serial Number
Serial number of a disk.

Unavailable Disk Space
Amount of space on the disk not available to be consumed. For the computer groups, this is the amount of space on disks that is not available to be consumed by the computers in a group. For disk and volume groups, this is the external disk space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Unavailable Logical Volume Space
Amount of space on a logical volume that is not available for use by a computer. For computer groups, this is the amount of space on logical volumes that is not available for use by the computers within those groups. For file system groups, this is the amount of space on the logical volumes within those groups that is not available for use. For disk groups and volume groups, this is the external logical volume space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

VM Disk File
Virtual machine disk's file on a monitored hypervisor that represents the disk on the virtual machine. This column displays **N/A** if not applicable to a disk.

VMFS Disk
Hypervisor disk where the virtual machine disk resides. This column displays **N/A** if not applicable to a disk.

Note: If a virtual machine disk's file spans multiple hypervisor disks, duplicate entries for the virtual machine disk are generated. If the report is distributed, the space is distributed across the hypervisor disks that are not external.

VFMS Mount Point
Hypervisor file system's mount point where the virtual machine disk's file resides. This column displays **N/A** if not applicable to a disk.

Related Topics

- All Disk Storage reports overview

- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Disk Storage - Disk Storage on Storage Subsystems reports:

Use these reports to view information about computer disks that IBM Tivoli Storage Productivity Center has identified as storage subsystem volumes.

See Disks On Storage Subsystems - Column Descriptions for descriptions of all the columns that appear in Disk Storage on Storage Subsystem reports.

Use By clauses to filter the entities that appear in a report. For example, if you want to view a report that displays specific computers, select **By Storage Subsystem** and use the **Selection** and **Filter** options to determine which storage subsystems appear in the generated report.

You can customize the information in each By clause by relating the monitored entities to other storage. The following options for relating storage entities (storage subsystems, computers, groups, file systems, logical volumes, disk groups, volume groups) to other storage in these reports includes the following:

- Use the **Relate Storage Subsystems To:** options in the By Storage Subsystem and By Storage Subsystem Group reports to relate storage subsystems to the following storage:
  - Computer Space
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space
  - Computer Disk Space
  - Distributed Computer Disk Space
- Use the **Relate Computers To:** and **Relate Disk/Volume Group To:** options in the By Computer, By Computer Group, and By Disk/Volume Group reports to relate storage computers, computer groups, disk groups, or volume groups to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space
  - Computer Disk Space
  - Distributed Computer Disk Space
  - Storage Subsystem Space (not available for disk groups and volume groups)
- Use the **Relate FileSystems/Logical Volumes To:** options in the By File System/Logical Volume and By Filesystem Group reports to relate file systems or logical volumes to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space

Considerations

Keep in mind the following considerations when generating Disk Storage on Storage Subsystems reports:

By Storage Subsystem (Relate Storage Subsystems To: File System/Logical Volume Space, Distributed File System/Logical Volume Space)

Duplicate entries appear in By Storage Subsystem reports for file systems

or logical volumes that span multiple storage subsystems. The storage space shown in the report is adjusted in the totals. The storage space values include space that is not storage subsystem storage if a file system or logical volume spans disks that are both storage subsystem disks and other storage disks.

When relating storage subsystems or computers to distributed file system/logical volume space, the reports displays duplicate entries for file systems or logical volumes that span multiple disks, but the displayed storage space unique. The values for storage space exclude space that is not storage subsystem storage if a file system or logical volume spans disks that are both storage subsystem disks and other storage disks.

By Storage Subsystem (Relate Storage Subsystems To: Computer Disk Space, Distributed Computer Disk Space)

Duplicate entries appear in these reports for virtual machines whose disks spans multiple hypervisor disks. The storage space shown in the report is adjusted in the totals. The storage space values include space that is not storage subsystem storage if the file system on a virtual machine spans hypervisor disks that are external or if the file system spans other storage disks.

When relating storage subsystems or computers to distributed computer disk space, this report displays duplicate entries for virtual machines whose disks spans multiple hypervisor disks, but the storage space shown in the report is unique. The storage space values exclude space that is not storage subsystem storage if the virtual machine disk's file system spans hypervisor disks that are external or the virtual machine disk's file system spans other storage disks.

By XIV system

The XIV system GUI and the XIV system command-line interface (CLI) report sizes in GB of disks, volumes, pools, and so on, based on 10^9. Tivoli Storage Productivity Center reports sizes in GB based on 2^30. For example, a volume showing a size of 17 GB in the XIV system GUI is reported as 16 GB (XIV GB size * 1000 * 1000 * 1000 / 1024 /1024 / 1024 rounded to the nearest GB) in Tivoli Storage Productivity Center.

Related Topics

Keep in mind the following considerations when generating Disk Storage on Storage Subsystems reports:

- Disks On Storage Subsystems - Column Descriptions
- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Disks On Storage Subsystems - Column Descriptions:

Use these reports to view information about disks on monitored computers that IBM Tivoli Storage Productivity Center has identified as storage subsystem volumes.

The following buttons and columns appear in the By clauses of Disks On Storage Subsystems reports:

Note: Not all of the following columns appear in each By clause.

Selection...
Click this button to select the specific entities that you want to include in the report. For example, in the **By Storage Subsystem** report, click this button to select the specific storage subsystems upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict the entities that you want to include in the report.

Relate <entity> To:
Select how you want to relate monitored storage entities in the report to other storage. See Disk Storage - Disk Storage Not On Storage Subsystems reports for more information.

Magnifying glass icon
This button appears on all **By <entity> Group** reports. Click this button to view a more detailed report about the entities in a group.

Available Disk Space
Amount of disk space on a storage subsystem that is available for a computer to consume. For computer groups, this is the amount of space on disks that is available for use by the computers within those groups. For storage subsystem groups, this is the amount of space on disks that is available for use by the storage subsystems within those groups. For disk and volume groups, this is the external disk space available for use within those groups.

This space includes the file system space that is available on the disk and the space that has not been allocated to a logical volume.

Available Disk/Volume Group Space
Amount of available space in a disk group or volume group.

Available File System/Logical Volume Space
Amount of space on the file system or logical volume that is available for use by the computer on a storage subsystem. For computers groups and storage subsystem groups, this is the amount of space on file systems or logical volumes that is available for use on storage subsystems by the entities within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that can be consumed. For disk groups and volume groups, this is the external logical volume or file system space available for use within those groups.

Computer
Name of a computer where a disk, file system, logical volume, or disk group is located.

Computer OS Type
The operating system running on a computer.

Consumed Disk Space
Amount of disk space on a storage subsystem that is consumed by a computer. For storage subsystem groups, this is the amount of consumed disk space within those groups. For computer groups, this is the amount of disk space consumed by the computers those groups. For disk and volume groups, this is the external disk space consumed within those groups.

This space includes the file system space that is consumed on a disk. If a logical volume on this disk space has no file system, then this space includes the consumed portion of the logical volume that has been allocated to the disk.

Consumed Disk/Volume Group Space
Amount of consumed disk or volume group space.

Consumed File System/Logical Volume Space
Amount of space on the file system or logical volume that is consumed by the computer on a storage subsystem. For computer groups and storage subsystem groups, this is the amount of space on file systems or logical volumes that is consumed by the entities within those groups. For file system groups, this is the amount of storage subsystem space on file systems or logical volumes within those groups that is consumed. For disk groups and volume groups, this is the external logical volume or file system space consumed within those groups.

Correlated
Indicates **Yes** if Tivoli Storage Productivity Center is able to relate the host storage directly to a storage subsystem. You can relate this correlated storage back to the actual storage subsystem storage for a more detailed view of how the host storage is allocated.

If a computer disk, file system, or logical volume appears in the Disk Storage On Storage Subsystem reports but is not correlated, it might be the result of the following reasons:

- The storage subsystem was not probed. Run a probe to collect data about the storage subsystem then generate this report.
- The storage subsystem storage was unassigned from the host storage, but the host has not been probed. Run a probe to collect data about the host then generate this report.
- The storage subsystem storage was not correctly identified by a probe of the host. If this is the problem, check the probe log of the host for messages relating to SCSI command errors and the probe results of the storage subsystem.

Creator
ID of the user who created a group.

Disk Path
Path name of a disk as seen by the computer.

Disk Space
Total disk space on a storage subsystem seen by a computer. For computer groups, this is the total disk space within those groups that is associated with storage subsystems. For storage subsystem groups, this is the total disk space on the storage subsystems within those groups. For disk and volume groups, this is the external disk space within the disk group or volume group.

Note: This space is distributed across the hypervisor disk space from where the virtual machine disk's file resides. The portion of space identified to a particular external hypervisor disk is reported only. This is true for all space values in this report.

Disk/Volume Group Name
Name of a disk group or volume group where a logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Disk/Volume Group Space
Total amount of disk group or volume group space. This is space inclusive of all storage space and applies to all space values related to disk groups or volume groups.

Disk/Volume Group Type
Type of a disk group. For example, a volume group might be LVM or Veritas. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Fibre Attached
Indicates **Yes** when Tivoli Storage Productivity Center determines that a disk is behind a fibre channel adapter port.

File System Type
Type of a file system, such as NTFS.

File System/Logical Volume Path
Name of a file system or logical volume path.

Group Name of a group.

Grown Defects
Number of defects detected on the disk since it was new.

Hypervisor Name
Name of the hypervisor where a virtual machine was created. This column displays **N/A** if not applicable to a disk.

Logical Volume Space
Total amount of space allocated to the logical volume on a disk. For computer groups, this is the total logical volume space in computer groups that is associated with storage subsystems. For storage subsystem groups, this is the total logical volume space on the storage subsystems within those groups. For file system groups, this is the total amount of space allocated to the logical volumes within those groups. For disk groups and volume groups, this is the external logical volume space within those groups.

This space includes the file system space on the disk reported if the entity displayed is a file system. The actual size of the file system on the disk reported is in addition to the consumed and available space.

Note: When relating a monitored entity to distributed file system/logical volume space, the space is distributed across the disks where a logical volume or file system resides. The portion identified to a disk is reported only.

(Disk) Manufacturer
The manufacturer of a disk.

Mapped to Storage Subsystem
Indicates **Yes** if Tivoli Storage Productivity Center has identified a disk as being located on a storage volume from a storage subsystem. The storage subsystem might not be monitored by Tivoli Storage Productivity Center.

(Disk) Model
Model number of a disk.

Number of Disks
Number of disks assigned to a storage subsystem or group. For disk groups and volume groups, this is the number of external disks within those groups. For NetApp storage devices, this number does not include disks in partner storage pools.

Number of File Systems/Logical Volumes
Number of logical file systems or logical volumes within a group. For disk groups and volume groups, this is the number of external logical file systems or logical volumes within those groups.

The logical volumes underlying the file systems are not counted for the value in the column.

Overallocated
Indicates **Yes** if Tivoli Storage Productivity Center is unable to uniquely identify disk storage across more than one host. Check the probe log of the host for messages relating to SCSI command errors.

Path Type
Path type of the entity being reported: file system or a logical volume.

Primary Defects
Number of defects on the disk when it was new.

Shared
Indicates **Yes** if a file system or logical volume is shared by more than one computer.

Serial Number
Serial number of a disk.

Storage Subsystem Manufacturer
The manufacturer of a storage subsystem.

Storage Subsystem Type
The type of storage subsystem.

Unavailable Disk Space
Amount of disk space on the storage subsystem that is not available for a computer to consume. For the computer groups, this is the amount of space on disks that is not available to be consumed by the computers within those groups. For the storage subsystem groups, this is the amount of space on disks that is not available to be consumed within those groups. For disk and volume groups, this is the external disk space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Unavailable Logical Volume Space
Amount of space on the logical volume that is not available for use by the computer on a storage subsystem. For computer groups and storage subsystem groups, this is the amount of space on logical volumes that is not available for use by the entities within those groups. For file system groups, this is the amount of space on the logical volumes within those groups that is not available for use. For disk groups and volume groups, this is the external logical volume space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

VM Disk File
Virtual machine disk's file on a monitored hypervisor that represents the disk on the virtual machine. This column displays **N/A** if not applicable to a disk.

VMFS Disk
Hypervisor disk where the virtual machine disk resides. This column displays **N/A** if not applicable to a disk.

Note: If a virtual machine disk's file spans multiple hypervisor disks, duplicate entries for the virtual machine disk are generated. If the report is distributed, the space is distributed across the hypervisor disks that are not external.

VFMS Mount Point
Hypervisor file system's mount point where the virtual machine disk's file resides. This column displays **N/A** if not applicable to a disk.

Related Topics

- Disks On Storage Subsystems reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Disk Storage - Disk Storage Not On Storage Subsystems reports:

Use these reports to view information about computer disks that IBM Tivoli Storage Productivity Center has not identified as storage subsystem volumes.

These reports are organized into the following categories:

Disks With Serial Numbers
These reports display information about non-storage subsystem disks that Tivoli Storage Productivity Center has identified with a valid serial number.

Disks Without Serial Numbers
These reports display information about non-storage subsystem disks that Tivoli Storage Productivity Center has not identified with serial numbers that uniquely identify those disks. This might occur on internal disks within a hypervisor and virtual machines that have storage created from internal disks within a hypervisor. If this problem occurs for other types of storage, check the probe log of the host for messages relating to SCSI command errors.

See Disks Not On Storage Subsystems - Column Descriptions for descriptions of all the columns that appear in Disk Storage Not On Storage Subsystems reports.

Use By clauses to filter the entities that appear in a report. For example, if you want to view a Disk Storage Not On Storage Subsystems report that displays specific computers, select **By Computer** and use the **Selection** and **Filter** options to determine which computers appear in that report.

You can customize the information in each By clause by relating the monitored entities to other storage. The following options for relating storage entities (computers, computer groups, file systems, file system groups, logical volumes, disk groups, volume groups) to other storage in these reports includes the following:

- Use the **Relate Computers To:** and **Relate Disk/Volume Group To:** options in the By Computer, By Computer Group, and By Disk/Volume Group reports to relate computers, computer groups, disk groups, and volume groups to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space
  - Computer Disk Space
  - Distributed Computer Disk Space
- Use the **Relate FileSystems/Logical Volumes To:** options in the By File System/Logical Volume and By Filesystem Group reports to relate file systems or logical volumes to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space

Related Topics

- Disks Not On Storage Subsystems - Column Descriptions
- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Disks Not On Storage Subsystems - Column Descriptions:

Use Disks Not On Storage Subsystems reports to view information about disk storage that is not associated with storage subsystems. IBM Tivoli Storage Productivity Center has not identified the disks in this report as storage subsystem volumes. Use this topic to view descriptions of the columns that appear in these reports.

The following buttons and columns appear in the By clauses of Disks Not On Storage Subsystems reports:

Note: Not all of the following columns appear in each By clause.

Selection...
Click this button to select the specific entities that you want to include in the report. For example, in the **By Filesystem/Logical Volume** report, click this button to select the specific file systems or logical volumes upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict the entities that you want to include in the report.

Relate <entity> To:
Select how you want to relate monitored storage entities in the report to other storage. See Disk Storage - Disk Storage Not On Storage Subsystems reports for more information.

Magnifying glass icon

This button appears on all **By <entity> Group** reports. Click this button to view a more detailed report about the entities in a group.

Available Disk Space

Amount of space on a disk that is available for use by a computer. For computer groups, this is the amount of space on disks that is available for use by the computers in a group. For disk and volume groups, this is the external disk space available for use within those groups.

This space includes the file system space that is available on the disk and the space that has not been allocated to a logical volume.

Available Disk/Volume Group Space

Amount of available space in a disk group or volume group.

Available File System/Logical Volume Space

Amount of space on the file system or logical volume that is available for use on a disk. For computers groups, this is the amount of space on file systems or logical volumes that is available for use by the computers within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that can be consumed. For disk groups and volume groups, this is the external logical volume or file system space available for use within those groups.

Computer

Name of a computer where a disk, file system, logical volume, or disk group is located.

Computer OS Type

The operating system running on a computer.

Consumed Disk Space

Amount of space on a disk that is consumed. For computer groups, this is the amount of disk consumed by the computers in those groups. For disk and volume groups, this is the external disk space consumed within those groups.

This space includes the file system space that is consumed on a disk. If a logical volume on this disk space has no file system, then this space includes the consumed portion of the logical volume that has been allocated to the disk.

Consumed Disk/Volume Group Space

Amount of consumed disk or volume group space.

Consumed File System/Logical Volume Space

Amount of space on the file system or logical volume consumed by the computer on a disk. For computer groups, this is the amount of space on file systems or logical volumes that is consumed by the computers within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that is consumed. For disk groups and volume groups, this is the external logical volume or file system space consumed within those groups.

Creator

ID of the user who created a group.

Disk Path

Path name of a disk as seen by the computer.

Disk Space
Total amount of space allocated to a computer disk. For computer groups, this is the total amount of space allocated to the disks on the computers within those groups. For disk and volume groups, this is the external disk space within the disk group or volume group.

Note: This space is distributed across the hypervisor disk space from where the virtual machine disk's file resides. The portion of space identified to a particular external hypervisor disk is reported only. This is true for all space values in this report.

Disk/Volume Group Name
Name of a disk group or volume group where a logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Disk/Volume Group Space
Total amount of disk group or volume group space. This is space inclusive of all storage space and applies to all space values related to disk groups or volume groups.

Disk/Volume Group Type
Type of a disk group. For example, a volume group might be LVM or Veritas. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Fibre Attached
Indicates **Yes** when Tivoli Storage Productivity Center determines that a disk is behind a fibre channel adapter port.

File System Type
Type of a file system, such as NTFS.

File System/Logical Volume Path
Name of a file system or logical volume path.

Group Name of a group.

Grown Defects
Number of defects detected on the disk since it was new.

Hypervisor Name
Name of the hypervisor where a virtual machine was created. This column displays **N/A** if not applicable to a disk.

Logical Volume Space
Total amount of space allocated to the logical volume on a disk. For computer groups, this is the total amount of space allocated to the logical volumes on disks within those groups. For file system groups, this is the total amount of space allocated to the logical volumes within those groups. For disk groups and volume groups, this is the external logical volume space within those groups.

This space includes the file system space on the disk reported if the entity displayed is a file system. The actual size of the file system on the disk reported is in addition to the consumed and available space.

Note: When relating a monitored entity to distributed file system/logical volume space, the space is distributed across the disks where a logical volume or file system resides. The portion identified to a disk is reported only.

(Disk) Manufacturer
The manufacturer of a disk.

(Disk) Model
Model number of a disk.

Number of Disks
Number of disks in a computer group. For disk groups and volume groups, this is the number of external disks within those groups.

Number of File Systems/Logical Volumes
Number of logical file systems or logical volumes within a group. For disk groups and volume groups, this is the number of external logical file systems or logical volumes within those groups.

The logical volumes underlying the file systems are not counted for the value in the column.

Overallocated
Indicates **Yes** if Tivoli Storage Productivity Center is unable to uniquely identify disk storage across more than one host. Check the probe log of the host for messages relating to SCSI command errors.

Path Type
Path type of the entity being reported: file system or a logical volume.

Primary Defects
Number of defects on the disk when it was new.

Shared
Indicates **Yes** if a file system or logical volume is shared by more than one computer.

Serial Number
Serial number of a disk.

Unavailable Disk Space
Amount of space on the disk not available to be consumed. For the computer groups, this is the amount of space on disks that is not available to be consumed by the computers in a group. For disk and volume groups, this is the external disk space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Unavailable Logical Volume Space
Amount of space on a logical volume that is not available for use by a computer. For computer groups, this is the amount of space on logical volumes that is not available for use by the computers within those groups. For file system groups, this is the amount of space on the logical volumes within those groups that is not available for use. For disk groups and volume groups, this is the external logical volume space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

VM Disk File
Virtual machine disk's file on a monitored hypervisor that represents the disk on the virtual machine. This column displays **N/A** if not applicable to a disk.

VMFS Disk
Hypervisor disk where the virtual machine disk resides. This column displays **N/A** if not applicable to a disk.

Note: If a virtual machine disk's file spans multiple hypervisor disks, duplicate entries for the virtual machine disk are generated. If the report is distributed, the space is distributed across the hypervisor disks that are not external.

VFMS Mount Point
Hypervisor file system's mount point where the virtual machine disk's file resides. This column displays **N/A** if not applicable to a disk.

Related Topics

- Disks Not On Storage Subsystems reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Disk Storage - Non-Disk Storage reports:

Use these reports to view information about storage space that does not reside on disks that are monitored by IBM Tivoli Storage Productivity Center. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory.

Use By clauses to filter the entities that appear in a report. For example, if you want to view a report that displays specific computers, select **By Computer** and use the **Selection** and **Filter** options to determine which computers appear in the generated report.

The following By clauses are available for Non-Disk Storage reports:

- By Computer
- By Computer Group
- By Filesystem/Logical Volume
- By Filesystem Group
- By Disk/Volume Group

You can customize the information in each By clause by relating the monitored entities to other storage. The following options for relating storage entities (computers, computer groups, file systems, file system groups, logical volumes, disk groups, volume groups) to other storage in these reports includes the following:

- File System or Logical Volume Space
- Distributed File System or Logical Volume Space

Related Topics

- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Non-Disk Storage - By Computer report, By Filesystem/Logical Volume report:

Use these reports to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. Select **By Computer** if you want to determine which computers are displayed in the report. Select **By Filesystem/Logical Volume** if you want to determine which file systems or logical volumes appear in the report.

Selection...
For the **By Computer** reports, click this button to select the specific computers upon which you want to report. For the **By Filesystem/Logical Volume** reports, click this button to select the specific file systems or logical volumes upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict computers or file system/logical volumes upon which you want to report.

Relate <entity> To:
Select how you want to relate computers or file systems in the report to other storage entities. See]Disk Storage - Non-Disk Storage reports for more information.

Computer
Name of a computer.

File System/Logical Volume Path
Name of a file system or logical volume path.

Note: Each monitored file system or logical volume path is displayed only once in this report because if the storage space is not distributed (there are no underlying disks), then there is no need to report distribution.

Logical Volume Space
Total amount of space allocated to the logical volume on a disk. This space includes the file system space if the entity displayed is a file system. The actual size of the file system is in addition to the consumed and available space.

Note: When relating computers to distributed file system/logical volume space, the space is distributed across the logical volumes where the file system or parent logical volume resides. The portion identified to a particular file system or parent logical volume is reported only. If a file system does not have an underlying logical volume, the logical volume space is the size of that file system.

Consumed File System/Logical Volume Space
Amount of space on a file system or logical volume that is consumed by a computer.

Available File System/Logical Volume Space
Amount of space on a file system or logical volume that is available for use by a computer.

Unavailable Logical Volume Space
Amount of space on the logical volume that is not available for use by a computer. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Computer OS Type
Operating system running on a computer.

Disk/Volume Group Name
Name of a disk or volume group where the logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Disk/Volume Group Type
Type of a disk or volume group where the logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

File System Type
Type of a file system, such as NTFS.

Path Type
Path type of an entity: file system or a logical volume.

Shared
Indicates **Yes** if a file system or logical volume is shared by more than one computer.

Non-Disk Storage - By Computer Group report, By Filesystem Group report:

Use these reports to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. Select **By Computer Group** if you want to determine which computer groups are displayed in the report. Select **By Filesystem Group** if you want to determine which file system groups appear in the report.

Selection...
For the **By Computer Group** reports, click this button to select the specific computer groups upon which you want to report. For the **By Filesystem Group** reports, click this button to select the specific file system group upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict groups upon which you want to report.

Relate <entity> To:
Select how you want to relate groups in the report to other storage entities. See Disk Storage - Non-Disk Storage reports for more information.

A unique row appears for each group in this report.

Magnifying glass icon
Click this button to view a more detailed report about the computers or file systems in a group.

Creator
ID of the user who created the group.

Group Name of a group.

Logical Volume Space
Total amount of space allocated to the computers or file systems in a group. This space includes the file system space if the entity displayed is a file system. The actual size of the file system is in addition to the consumed and available space.

Note: When relating groups to distributed file system/logical volume space, this space is distributed across the disks from where the logical volume or file system reside. This is true for all space values in the Distributed File System/Logical Volume Space report.

Consumed File System/Logical Volume Space
Amount of space on file systems or logical volumes within a group that is consumed.

Available File System/Logical Volume Space
Amount of space on file systems or logical volumes within a group that can be consumed.

Unavailable Logical Volume Space
Amount of space on the logical volumes in a group that is not available for use. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Number of File Systems/Logical Volumes
Number of file systems or logical volumes in a group.

Non-Disk Storage - By Computer report, By Filesystem/Logical Volume report:

Use these reports to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. Select **By Computer** if you want to determine which computers are displayed in the report. Select **By Filesystem/Logical Volume** if you want to determine which file systems or logical volumes appear in the report.

Selection...
For the **By Computer** reports, click this button to select the specific computers upon which you want to report. For the **By Filesystem/Logical Volume** reports, click this button to select the specific file systems or logical volumes upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict computers or file system/logical volumes upon which you want to report.

Relate <entity> To:
Select how you want to relate computers or file systems in the report to other storage entities. See]Disk Storage - Non-Disk Storage reports for more information.

Computer
Name of a computer.

File System/Logical Volume Path
Name of a file system or logical volume path.

Note: Each monitored file system or logical volume path is displayed only once in this report because if the storage space is not distributed (there are no underlying disks), then there is no need to report distribution.

Logical Volume Space
Total amount of space allocated to the logical volume on a disk. This space includes the file system space if the entity displayed is a file system. The actual size of the file system is in addition to the consumed and available space.

Note: When relating computers to distributed file system/logical volume space, the space is distributed across the logical volumes where the file system or parent logical volume resides. The portion identified to a particular file system or parent logical volume is reported only. If a file system does not have an underlying logical volume, the logical volume space is the size of that file system.

Consumed File System/Logical Volume Space
Amount of space on a file system or logical volume that is consumed by a computer.

Available File System/Logical Volume Space
Amount of space on a file system or logical volume that is available for use by a computer.

Unavailable Logical Volume Space
Amount of space on the logical volume that is not available for use by a computer. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Computer OS Type
Operating system running on a computer.

Disk/Volume Group Name
Name of a disk or volume group where the logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Disk/Volume Group Type
Type of a disk or volume group where the logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

File System Type
Type of a file system, such as NTFS.

Path Type
Path type of an entity: file system or a logical volume.

Shared
Indicates **Yes** if a file system or logical volume is shared by more than one computer.

Non-Disk Storage - By Computer Group report, By Filesystem Group report:

Use these reports to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. Select **By Computer Group** if you want to determine which computer groups are displayed in the report. Select **By Filesystem Group** if you want to determine which file system groups appear in the report.

Selection...
For the **By Computer Group** reports, click this button to select the specific computer groups upon which you want to report. For the **By Filesystem Group** reports, click this button to select the specific file system group upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict groups upon which you want to report.

Relate <entity> To:
Select how you want to relate groups in the report to other storage entities. See Disk Storage - Non-Disk Storage reports for more information.

A unique row appears for each group in this report.

Magnifying glass icon
Click this button to view a more detailed report about the computers or file systems in a group.

Creator
ID of the user who created the group.

Group Name of a group.

Logical Volume Space
Total amount of space allocated to the computers or file systems in a group. This space includes the file system space if the entity displayed is a file system. The actual size of the file system is in addition to the consumed and available space.

Note: When relating groups to distributed file system/logical volume space, this space is distributed across the disks from where the logical volume or file system reside. This is true for all space values in the Distributed File System/Logical Volume Space report.

Consumed File System/Logical Volume Space
Amount of space on file systems or logical volumes within a group that is consumed.

Available File System/Logical Volume Space
Amount of space on file systems or logical volumes within a group that can be consumed.

Unavailable Logical Volume Space
Amount of space on the logical volumes in a group that is not available for use. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Number of File Systems/Logical Volumes
Number of file systems or logical volumes in a group.

Disk Storage on Storage Subsystems reports - By Disk/Volume Group:

Use this report to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. This report enables you to determine which disk groups or volume groups are displayed.

Selection...
Click this button to select the specific disk groups or volume groups upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict disk groups or volume groups upon which you want to report.

Relate Disk/Volume Group To:
Select how you want to relate the non-disk storage in disk groups or volume groups to other storage entities. See Disk Storage - Non-Disk Storage reports for more information.

A unique row appears for each disk group or volume group.

Magnifying glass icon
Click this button to view a more detailed report about the disk or file system within the disk group or volume group.

Disk/Volume Group Name
Name of a group.

Disk/Volume Group Type
Type of a group. For example, a volume group might be LVM or Veritas.

Computer
Name of the computer where a disk group or volume group resides.

Computer OS Type
Operating system running on a computer.

Disk/Volume Group Space
Total amount of disk group or volume group space. This is space inclusive of all storage space and applies to all space values related to disk groups or volume groups.

Consumed Disk/Volume Group Space
Amount of consumed disk group or volume group space.

Available Disk/Volume Group Space
Amount of available disk group or volume group space.

Logical Volume Space
External logical volume space within a disk group or volume group.

Consumed File System/Logical Volume Space
External logical volume or file system space that is consumed within a disk group or volume group.

Available File System/Logical Volume Space
External logical volume or file system space that is available for use within a disk group or volume group.

Unavailable Logical Volume Space
External logical volume space that is not available for use within a disk group or volume group. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Number of File Systems/Logical Volumes
Number of external logical file systems or logical volumes within a disk or volume group. The logical volumes underlying the file systems are not counted.

Additional reporting support and information:

Learn about the additional storage devices upon which you can report within Data Manager.

IBM FAStT Subsystem support: The IBM FAStT subsystem reporting feature further extends the subsystem reporting capabilities of Data Manager. You can use the Reporting facility to view information about the capacity, allocation, and volumes of an IBM FAStT subsystem. You can view IBM FAStT subsystem information by using the following menu options to generate reports:

- **Reporting → Asset → By Storage Subsystem →** <IBM FAStT subsystem name>

- **Reporting → Asset → System-wide → Storage Subsystems**
- **Reporting → Asset → System-wide → Volumes**
- **Reporting → Asset → System-wide → Disks**

Note: Data Manager shows an incorrect value for the **Storage Pool Free Space** column in IBM FAStT subsystem reports (for example, **Asset → By Storage Subsystem → <FAStT_subsystem_name>**, and **Asset → System-wide → Storage Subsystems**). The value for **Storage Pool Free Space** in these reports reflects the space from one Free Capacity node rather than all the free space within the defined arrays on a FAStT subsystem.

Hitachi Data Systems Subsystem support: Use the Reporting facility to report on Hitachi Data Systems (HDS) storage subsystems. You can view reports that contain information about the capacity, allocation, and volumes of an HDS subsystem.

Any storage pools or volumes created using the HDS **HiCommand** software will not be reported through the HDS CIMOM provider. **HiCommand** is the native interface for HDS subsystems. If you configure an HDS subsystem through **HiCommand** , IBM Tivoli Storage Productivity Center will not be able to gather information about storage pools and volumes from that subsystem. The HDS CIM agent provider will return all storage pools created through the **HiCommand** interface as a single pool containing all of the **HiCommand** created pools. This single pool is always reported as Raid 5.

Note: Data Manager does not report on the raw capacity of disks internal to the HDS subsystem. Information on the internal disks of the HDS subsystem is not available through the CIMOM interface provided by the Hitachi HiCommand Device Manager.

You can view HDS storage subsystem information by using the following menu options to generate reports:

- **Reporting → Asset → By Storage Subsystem →** <HDS subsystem name>
- **Reporting → Asset → System-wide → Storage Subsystems**
- **Reporting → Asset → System-wide → Volumes**
- **Reporting → Asset → System-wide → Disks**
- **Reporting → TPC-wide Storage Space → Disk Space**

HP StorageWorks Subsystem support: Use the Reporting facility to report on HP StorageWorks storage subsystems. You can view reports that contain information about the capacity, allocation, and volumes of an HP StorageWorks subsystem.

You can view HP StorageWorks storage subsystem information by using the following menu options to generate reports:

- **Reporting → Asset → By Storage Subsystem →** <HP StorageWorks subsystem name>
- **Reporting → Asset → System-wide → Storage Subsystems**
- **Reporting → Asset → System-wide → Volumes**
- **Reporting → Asset → System-wide → Disks**

IBM Tivoli Storage SAN File Systems support: You can also perform scans and probes to gather information about IBM Tivoli Storage SAN File Systems and generate reports based on that information using the Reporting facility. You can

view information gathered by probes of the SAN File System by using the following menu options to generate reports:

- **Reporting → Asset → By OS Type → IBM SAN File System**
- **Reporting → Asset → By Computer**
- **Reporting → TPC-wide Storage Space → File System Space**
- **Reporting → TPC-wide Storage Space → Consumed File System Space**
- **Reporting → TPC-wide Storage Space → Available File System Space**

Note: Due to the complexity of the SAN File System, the values displayed for total capacity, used space and free space in the Capacity and Asset reports might not be accurate. The values displayed are the values returned by the probe or scane and they may not represent actual values for your SAN File System. Use the SAN File System admin console to obtain the actual values.
You can view information gathered by scans of the SAN File System by using the following menu options to generate reports:

- **Reporting → Usage**
- **Reporting → Usage Violations**
- **Reporting → Backup**

Refer to the *Creating Directory Groups* section in the *Monitoring* chapter for more information about how to configure SAN File System support within the Reporting facility.

Report Array, Loop, and Adapter Pair for IBM TotalStorage Enterprise Storage Server Disk Groups support: Disk Manager provides report information on Array, Loop, and Adapter Pairs for TotalStorage Enterprise Storage Server Disk Groups. This information is available when managing an TotalStorage Enterprise Storage Server through IBM CIM Agent (ICAT) for TotalStorage Enterprise Storage Server version 1.2.0.13. Contact IBM Customer Support to obtain IBM CIM Agent for TotalStorage Enterprise Storage Server version 1.2.0.13.

You can view Array, Loop, and Adapter Pair information by using the menu option **Reporting → Asset Reporting → By Storage Subsystem** to generate reports.

Conversion of volume sizes by IBM Tivoli Storage Productivity Center:

Tivoli Storage Productivity Center uses 1024 bytes to represent a kilobyte. However, some element managers use 1000 bytes to represent a kilobyte. When reporting volume sizes, Tivoli Storage Productivity Center converts base-1000 values to base-1024 values.

The following table shows the byte values that are associated with base-1024 and base-1000 systems:

Unit	Abbreviation	Bytes in base-1024	Bytes in base-1000
Kilobyte	KB in base-1024 kB in base-1000 see note	2^{10} or 1024	10^3 or 1000
Megabyte	MB	1024^2 or 1,048,576	10^6 or 1,000,000
Gigabyte	GB	1024^3 or 1,073,741,824	10^9 or 1,000,000,000
Terabyte	TB	1024^4 or 1,099,511,627,716	10^{12}or 1,000,000,000,000

Unit	Abbreviation	Bytes in base-1024	Bytes in base-1000
Note: In SI Metric Style, the abbreviation of the prefix kilo is lowercase k. IBM distinguishes between lowercase k (meaning 1000 or 10^3) and uppercase K (meaning 1024 or 2^{10}). Therefore, KB means 1024 bytes, and kB means 1000 bytes.			

You can verify the volume size reported by Tivoli Storage Productivity Center against the size reported by element managers that use a base-1000 system. For example, suppose an element manager reports the size of a volume as 62,000 MB. Use the following calculations to verify the total number of base-1024 GB as reported by Tivoli Storage Productivity Center:

1. Multiply the volume size reported by the element manager by the number of bytes in a base-1000 MB:

   `62,000 MB × 1,000,000 bytes per MB = 62,000,000,000 bytes`
2. Divide the result by the number of bytes in a base-1024 MB:

   `62,000,000,000 bytes / 1,048,576 bytes per MB = 59,127.81 MB`
3. Divide the quotient by the number of MB in a base-1024 GB:

   `59,127.81 MB / 1024 MB per GB = 57.74 GB`

Tivoli Storage Productivity Center reports the volume size as 57.74 GB.

VMware capacity reports:

You must probe both the ESX Server and the Storage Resource agent on the virtual machines before you can generate accurate reports for disk and file system capacity.

The TOTAL row in the capacity report shows the capacity for the file system or the disk capacity. The total includes virtualized disks, virtual machines, non-virtualized disks, and non-virtualized machines.

For example, you have an ESX Server that has a capacity of 100 GB and 60 GB is allocated to the virtual machine. The virtual machine uses 5 GB of space. Both the ESX Server (H1) and the virtual machine (VC1) are probed. You also have a physical machine (PC1) that is probed. The capacity is shown in the report, as follows:

Column heading	Capacity	Used space (calculated as Capacity minus Free Space)	Free space
TOTAL	130 GB	25 GB	105 GB
H1	100 GB	60 GB	40 GB
VC1	60 GB	5 GB	55 GB
PC1	30 GB	20 GB	10 GB

If you probed the physical machine (PC1) and the virtual machine (VC1) but did not probe the ESX Server (hypervisor), the capacity is shown as follows:

Column heading	Capacity	Used Space (calculated as Capacity minus Free Space; negative values are shown as 0)	Free Space
TOTAL	30 GB	0 GB	65 GB
PC1	30 GB	20 GB	10 GB
VC1	60 GB	5 GB	55 GB

If you probed the hypervisor (H1) and the physical machine (PC1) but did not probe the virtual machine (VC1), the capacity is shown as follows:

Column heading	Capacity	Used Space (calculated as Capacity minus Free Space)	Free Space
TOTAL	130 GB	80 GB	50 GB
H1	100 GB	60 GB	40 GB
PC1	30 GB	20 GB	10 GB

If you have an ESX Server that has devices that are directly attached to the virtual machine, the file system capacity created on the virtual machine is not added in the TOTAL row.

For example, the file system capacity shown in the table is based on the following factors:

- The file system (FS1) created on the hypervisor has a capacity of 100 GB and uses 60 GB of space.
- The file system (FS1) created from internal storage on the hypervisor (H1) has a capacity of 60 GB and uses 5 GB of space.
- The file system (FS2) on the device that is directly attached to the virtual machine (VC1) has a capacity of 5 GB and uses 4 GB of space.
- The file system created on the physical machine (PC1) has a capacity of 30 GB and uses 20 GB of space.

Column heading	Capacity	Used Space (calculated as Capacity minus Free Space)	Free Space
TOTAL	130 GB	24 GB	106 GB
H1 FS1	100 GB	60 GB	40 GB
VC1 FS1	60 GB	5 GB	55 GB
VC1 FS2	5 GB	4 GB	1 GB
PC1 FS1	30 GB	20 GB	10 GB

Data Manager for Databases reports

Use Data Manager for Databases for reports to view both overview and detailed information about your RDBMS storage resources, including instances, tablespaces, databases, devices, containers, data files, fragments, tables, control files, redo logs,

archive log directories, users. These reports are constructed from the statistical information gathered by data collection jobs and accumulated in the database repository.

Using Data Manager for Databases reports:

Learn how to use Data Manager for Databases reports to view information about the instances, tablespaces, databases, devices, containers, data files, fragments, tables, control files, redo logs, archive log directories, and users within your environment.

Before you begin:

Learn about information you should know before generating and viewing Data Manager for Databases reports.

Before you can use reports to view detailed information about tables in an RDBMS, you must have 1) defined a table group that contains the tables about which you want to view storage information and 2) scheduled and run a scan against that table group. See "Working with groups of resources" on page 337 for information about using groups. See "Collecting storage usage and trending information" on page 306 for information about scans and how to create them.

Data Manager for Databases - Report Categories:

Use the report categories within Data Manager for Databases reports to collect information about RDBMS storage assets, capacity, usage, and usage violations.

The following table describes the categories of reports available in Data Manager for Databases.

Report Category	Use these reports to do the following tasks	What job collects information for these reports?
Asset Reports	View the information gathered by probes. While running probes, Data Manager for Databases itemizes the information for each type of database. • Oracle: instances, table spaces, tables, data files, control files, redo logs, and archive log directories • SQL Server: instances, databases, tables, and data files • Sybase: instances, devices, databases, tables, and fragments. • UDB: instances, databases, table spaces, tables, containers, and users	Probes

Report Category	Use these reports to do the following tasks	What job collects information for these reports?
Capacity Reports	View the storage capacity of different entities within an RDBMS. Specifically, use Capacity reports to view: • **Instance storage capacity within all DBMSs.** If you are monitoring the instances for different DBMSs, you can use the **All DBMSs** node to view a summary of the storage capacity for all those instances in a single view. • **Oracle databases** – How much storage you have for instances – How much storage is being used for data files – What percentage of storage for a data file is free for expansion • **SQL Server databases** – How much storage you have for instances – How much storage is being used for data files – What percentage of storage for a data file is free for expansion • **Sybase databases** – How much storage you have for instances – How much storage is being used for devices – What percentage of storage for a device is free for expansion – How much storage is being used for fragments – What percentage of storage for a fragment is free for expansion • **UDB databases** – How much allocated storage you have for a UDB Instance and the objects it contains – How much storage is being used for a UDB Instance and the objects it contains – What percentage of UDB storage is free for expansion The Capacity reporting function provides you with the ability to view the storage capacity, used space, and free space for your RDBMS objects in a single view.	Probes

Report Category	Use these reports to do the following tasks	What job collects information for these reports?
Usage Reports	Monitor and report on the usage and growth of the consumption of your DBMS storage: • **Storage usage within all DBMSs.** If you are monitoring the multiple DBMSs, use this node to view common storage usage information for databases, table spaces, tables, and users across all monitored DBMSs. • **Oracle databases** – View detailed information for table spaces including their total size, used space, empty used space, # of tables and indexes, and so on. – Identify the tables consuming the most space within your environment. – Identify the segments with the most extents and most empty used space. – Associate users with the amount of storage they are consuming on specific databases, table spaces, groups of table spaces, specific computers, groups of computers, and throughout the entire network. • **SQL Server databases** – View detailed information for databases including their total size, used space, empty used space, # of tables and indexes, and so on. – Identify the tables consuming the most space within your environment. – Associate users with the amount of storage they are consuming on specific databases, groups of databases, specific computers, groups of computers, and throughout the entire network. • **Sybase databases** – View detailed information for databases including their total size, used space, empty used space, # of tables and indexes, and so on. – Identify the tables consuming the most space within your environment. – Associate users with the amount of storage they are consuming on specific databases, groups of databases, specific computers, groups of computers, and throughout the entire network. • **UDB databases** – View detailed information for table spaces including their total size, used space, empty used space, # of tables and indexes, and so on. – Identify the tables consuming the most space within your environment. You can view this information at the table space, database, and computer level, as well as at a network-wide level. – Associate users with the amount of storage they are consuming on specific databases, table spaces, groups of table spaces, specific computers, groups of computers, and throughout the entire network.	Scans
Usage Violation Reports	Create and enforce corporate storage usage policies. You can report on violations of quotas that are set at the network-wide, instance, table space, and database levels.	Scans, Quotas

Reporting on RDBMS Objects:

Learn about the RDBMS objects upon which you can report using Data Manager for Databases reports.

The following table lists the RDBMS objects whose storage usage is monitored by Data Manager for Databases.

Oracle	SQL Server	Sybase	DB2/UDB
Instances	**Instances**	**Instances**	**Instances**
Databases	**Databases**	**Devices**	**Databases**
Tablespaces	**Datafiles**	**Databases**	**Tablespaces**
Datafiles	**Tables**	**Fragments**	**Containers**
Tables	**Indexes**	**Tables**	**Tables**
Indexes	**Log Files**	**Indexes**	**Indexes**
Extents			
Segments			
Redo Log files			
Control Files			
Archive Log Directories			

Note:

1. Each of the major reporting categories in the Reporting facility can be expanded in the navigation tree to reveal additional sub-categories of reporting. Some of the nodes within the reporting tree will appear only if you set up access for the RDBMS to which that node applies. For example, if you set up for monitoring Oracle and SQL Server instances only, nodes representing the other RDBMSs supported by Data Manager for Databases (such as Sybase and UDB) will not appear within the navigation tree. If you have set up access for only one RDBMS, the **All DBMSs** node will also not be available in some nodes.
2. Data Manager for Databases does not currently support the monitoring of clustered database applications.

Using By clauses:

Use the By clauses associated with Data Manager for Databases reports as a method for viewing the data in those reports from different perspectives.

As you drill down through Data Manager for Databases's reports, you will notice a set of options that repeats for each report type. These *By Clauses* allow you to generate variations of the reports with different perspectives and groupings of the data. Some of the available By Clauses are described in the following table:

By Clause	Description
By DBMS Type	Use this clause to select the DBMS that you want to report upon.
By Instance	Use this clause to select the instances that you want to report upon and the report that will be generated based upon those instances.
By Device (Sybase only)	Use this clause to select the devices that you want to report upon and the report that will be generated based upon those devices.
By Database	Use this clause to select the database that you want to report upon and the report that will be generated based upon those databases.
By Tablespace	Use this clause to select the tablespaces that you want to report upon and the report that will be generated based upon those tablespaces.
By Tablespace Group	Use this clause to select the tablespace groups that you want to report upon and the report that will be generated based upon those tablespace groups.

By Clause	Description
By Fragments (Sybase only)	Use this clause to select the fragments that you want to report upon and the report that will be generated based upon those fragments.
By Container (UDB only)	Use this clause to select the containers that you want to report upon and the report that will be generated based upon those containers.
By Computer	Use this clause to select the computers that you want to report upon and the report that will be generated based upon those computers.
By Computer Group	Use this clause to select the computer groups that you want to report upon and the report that will be generated based upon those computer groups.
By User	Use this clause to select the users that you want to report upon and the report that will be generated based upon those users.
Network Wide	The reports and charts that you see will be generated on all data, network wide. This allows you to get a global view of the storage and storage usage for your whole enterprise.

Asset reports:

Use Asset reports to view information about the instances, devices, databases, tablespaces, tables, containers, data files, fragments, control files, redo logs, and archive log directories, and general hardware inventory of RDBMS storage assets within your environment.

While running probes and scans, Data Manager for Databases itemizes your RDBMS storage assets and provides a hierarchical view of that information so you can drill down to view your those assets in greater detail. You can view information about RDBMS storage resources through the following views:

Asset Reporting View	Use these reports to...
Asset > By Computer	View the instances and their related objects based upon the computers on which they reside.
Asset > By DBMS Type	This node appears if you are using Data Manager for Databases to monitor more than one RDBMS. Expand this node to view asset information about instances organized according to their DBMS.
Asset > System-wide	Investigate all of your instances, databases, tablespaces, tables, container, data files, users and other assets for your *entire* system. **Note:** (UDB only) For the **All DBMSs** reports available under the **Asset > System-wide** reporting nodes, the Database-Tablespaces view shows information about UDB tablespaces only.

By Computer, By DBMS Type:

Learn about the reports available within the **By Computer** and **By DBMS Type** nodes within Asset reporting.

Asset Report Type	Use these reports to...	What's in these reports?
By Computer By Computer > *machine_name* > *instance/database_name* **By DBMS Type** By DBMS Type > *dbms_type* > *instance/database_name*	Drill down through your DBMS storage assets in a hierarchical manner according to the computers on which they reside or the DBMS to which they belong.	Depending on the DBMS of the Instance you want to view, you can drill down and expose its: • Tablespaces • Databases • Devices • Monitored tables • Containers • Data files • Fragments • Control files • Redo logs • Archive log directories

Asset Report Type	Use these reports to...	What's in these reports?
Instance Information **By Computer** By Computer > *machine_name* > *instance_name* **By DBMS Type** By DBMS Type > *dbms_type* > *instance_name*	View information about all of your monitored instances. Specifically, you can view information about the following information based upon DBMS type: **SQL Server** • Databases • Monitored Tables • Data Files **Oracle** • Tablespaces • Monitored Tables • Data Files • Control Files • Redo Logs • Archive Log Directories **Sybase** • Devices • Databases • Monitored Tables • Fragments **UDB** • Instances • Databases • Tablespaces • Tables • Containers • Users	**SQL Server** • RDBMS Type, Name, Version, Loginid • Number of Databases, Number of Datafiles, Number of Log Files • Total Size of all Storage, Capacity of Data Files, Free Space of Data Files, Capacity of Log Files, Free Space of Log Files • Probe Status, Last Probe Time **Oracle** *Database Information tab* • Name, Owner, Block Size, Log Mode, Status • Maximum Number of Instances, Number of Tablespaces, Maximum Number of Data Files, Number of Data Files, Maximum Number of Log Files, Number of Log Files • Total Size of all Storage, Capacity Of Data Files, Free Space of Data Files • Create Time, Discovered Time, Dropped Time *Instance Information tab* • RDBMS Type, Name, Version, Home Directory, Loginid, Port, Probe Status • Last Probe Time, Last Start Time, Discovered Time **Sybase** • RDBMS Type, Name, Version, Home Directory, Loginid, Port • Number of Databases, Number of Devices, Number of Fragments, Number of Log Fragments • Total Size of all Storage, Mirror Overhead, Capacity of Data Fragments, Free Space in Data Fragments, Capacity of Log Fragments, Free Space in Log Fragments • Probe Status, Last Probe Time **UDB** • RDBMS Type, Name, Version, Home Directory, Loginid, Port • Number of Tablespaces, Number of Containers, Number of Log Files • Total Size of all Storage, Container Capacity, DMS Container Free Space, Total Size of all Log Files • Probe Status, Last Probe Time

Asset Report Type	Use these reports to...	What's in these reports?
Oracle Detailed Instance Information **By Computer** By Computer > *machine_name* > *instance_name*	View information about the objects associated with a selected Oracle Instance.	• **Tablespaces**: When you expand the **Tablespaces** node, click the name of a tablespace to view information about that tablespace and its related data file. • **Monitored Tables**: When you expand the **Monitored Tables** node, click on the name of a table to view its information. Only the tables that belong to Table Groups included in scans will be available to view through this node. • **Data Files**: Click the **Data Files** node, click the name of a data file to view information about the data files associated with the Instance, including: Data File Path, Used Space, Percent Used, Free Space, Total Size, Coalesced Extents, Minimum Free Extent Size, Maximum Free Extent Size, and Create Time. • **Control Files**: Click this node to view information on the control files associated with the Instance, including the File Name (full path) and File Size. • **Redo Logs**: Click this node to view information on the redo logs associated with the Instance. – for Log Group: Sequence Number, Status, Archived, Total Size – for Log File: File Name (full path) • **Archive Log Directories**: Click this node to view information on the archive log directory associated with the Instance, including Directory Name, Size of Files, File Count, Status, Last Failure Time.
SQL Server Detailed Instance Information **By Computer** By Computer > *machine_name* > *server_name* **By DBMS Type** By DBMS Type > SQL/Server > *server_name*	View information about the objects associated with a selected SQL Server Instance.	• **Databases**: When you expand the **Databases** node, click the name of a Database to view information about that database, its related data file, and its related log file. • **Monitored Tables**: When you expand the **Monitored Tables** node, click on the name of a table to view its information including the computer, instances, and database on which is resides, as well as its creator, total size, empty used space, index size, number of indexes, row count, etc. Only the tables that belong to table groups included in scans will be available to view through this node. • **Data Files**: Click the **Data Files** node, click the name of a data file to view information about the data files associated with the Instance, including: Data File Path, Used Space, Percent Used, Free Space, Total Size, Coalesced Extents, Minimum Free Extent Size, Maximum Free Extent Size, and Create Time.

Asset Report Type	Use these reports to...	What's in these reports?
Sybase Detailed Instance Information **By Computer** By Computer > *machine_name* > *instance_name* **By DBMS Type** By DBMS Type > Sybase > *instance_name*	view information about the objects associated with a selected Sybase Instance.	• **Devices**: Click this node to view information about the Device associated with the Instance, including its name, computer on which it resides, path, free space, percent free, used space, total size, fragments, mirror path, etc. • **Databases**: When you expand the **Databases** node, click the name of a Database to view information about that database, its device information, and its related fragment information. • **Monitored Tables**: When you expand the **Monitored Tables** node, click on the name of a table to view its information including the computer, instances, and database on which is resides, as well as its creator, total size, empty used space, index size, number of indexes, row count, etc. Only the tables that belong to table groups included in scans will be available to view through this node. • **Fragments**: Click this node to view about the Fragments associated with the Instance including the computer, instances, and database on which they resides as well as their path, used space, percent used, free space, total size, data type, etc.
UDB Detailed Instance Information **By Computer** By Computer > *machine_name* > *instance_name* **By DBMS Type** >	view information about the objects associated with a selected UDB Instance.	• **Databases**: When you expand the **Databases** node, click the name of a databases to view information about that database. • **Monitored Tables**: When you expand the **Monitored Tables** node, click on the name of a table to view its information. Only the tables that belong to Table Groups included in scans will be available to view through this node. • **Containers**: Click the **Containers** node, click the name of a container to view information about the containers associated with the Instance, including: Database, Tablespace, Container Node, Container Name, Used Space, Percent Used, Free Space, Total Size, Container Type, Status, Discovered Time.

System-wide:

Use System-wide reports to view your RDBMS storage assets as a whole, across the storage environment

Asset Report Type	Use these reports to...	What's in these reports?
System-Wide **Asset** > **System-wide**	drill down through your DBMS storage assets in a hierarchical manner.	Depending on the DBMS of the Instance you want to view, you can drill down and display the following items. • Table spaces • Databases • Devices • Monitored tables • Containers • Data files • Fragments • Control files • Redo logs • Archive log directories
All DBMSs **Asset** > **System-wide** > **All DBMSs**	View information about all the DBMSs within your environment that is based on the objects that are related to those instances.	• **Instances**: Expand the **Instances** node and click the sorting method by which you want to view instances. • **Databases-Table spaces**: Expand **Databases-Table spaces** node and click the sorting method by which you want to view information about the databases and/or table spaces within your environment. • **Tables**: Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to table groups included in scans display within the reports that are viewed through this node. • **Files**: Expand the **Files** node and click the sorting method by which you want to view information about data files/containers. The viewable information includes File Path, Free Space, Percent Free, Used Space, Total Size, Create Time, and Discovered Time. • **Users**: Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored DBMS within your environment.

Asset Report Type	What's in these reports?
All DBMSs (continued)	• **Instances**: Expand the **Instances** node and click the sorting method by which you want to view instances. – By Instance: sort that is based on Instance name – By Version: sort that is based on the version of the DBMS under which an Instance is running – By Start Time: sort that is based on when an Instance was last started – By Probe Time: sort that is based on when an Instance was last probed – By RDBMS Type: sort that is based on the DBMS of an Instance • **Databases-Table spaces**: Expand the **Databases-Table spaces** node and click the sorting method by which you want to view information about the databases or table spaces within your environment. – By Database-Table space: sort that is based on DB/TB name – By Free space: sort that is based on the amount of free space on a DB/TB – By Total Size: sort that is based on the total size of a DB/TB – By Table Count: sort that is based on the number of tables on a DB/TB – By Scan Time: sort that is based on when a DB/TB was last scanned – By Discovered Time: sort that is based on when a DB/TB was discovered – Dropped Table spaces: sort that is based on dropped DBs or TBs. – By RDBMS Type: sort that is based on the DBMS of a DB/TB • **Tables**: Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans display within the reports that are viewed through this node. – By Table: sort that is based on table name – By Number of Indexes: sort that is based on the number of indexes that are associated with a table – By Total Size: sort that is based on the total size of a table – By Row Count: sort that is based on the number of rows in a table – By Chained Row Count: sort that is based on the number of chained rows in a table – By Discovered Time: sort that is based on when a table was discovered – Dropped Tables: sort that is based on dropped tables – By RDBMS Type: sort that is based on the DBMS of a table • **Files**: Expand the **Files** node and click the sorting method by which you want to view information about DBMS files. – By File: sort that is based on file name – By Free Space: sort that is based on file free space – By Total Size: sort that is based on the total size of a file – By Discovered Time: sort that is based on when a file was discovered – By RDBMS Type: sort that is based on the DBMS of a file • **Users**: Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored DBMS within your environment.

Asset Report Type	What's in these reports?
Oracle **Asset > System-wide > Oracle**	• **Instances**: Expand the **Instances** node and click the sorting method by which you want to view instances. – By Instance: sort that is based on Instance name – By Version: sort that is based on the version of the DBMS under which an Instance is running – By Start Time: sort that is based on when an Instance was last started – By Probe Time: sort that is based on when an Instance was last probed • **Table spaces**: Expand this node and click the sorting method by which you want to view information about table spaces within your environment. – By Table space: sort that is based on table space name – By Max Free Extent: sort that is based on the maximum free extent size – By Free space: sort that is based on the amount of free space on a table space – By Total Size: sort that is based on the total size of a table space – By Coalesced Extents: sort that is based on the number of coalesced extents – By Table Count: sort that is based on the number of tables – By Scan Time: sort that is based on when a table space was last scanned – By Discovered Time: sort that is based on when a table space was discovered – Dropped Table spaces: sort that is based on dropped table spaces • **Tables**: Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans display within the reports that are viewed through this node. – By Table: sort that is based on table name – By Number of Indexes: sort that is based on the number of indexes that are associated with a table – By Total Size: sort that is based on the total size of a table – By Row Count: sort that is based on the number of rows in a table – By Chained Row Count: sort that is based on the number of chained rows in a table – By Discovered Time: sort that is based on when a table was discovered – Dropped Tables: sort that is based on dropped tables • **Data Files**: Expand the **Data Files** node and click the sorting method by which you want to view information about DBMS files. – By Data File: sort that is based on file name – By Max Free Extent: sort data files that are sorted by their maximum free extent sizes – By Free Space: sort that is based on file free space – By Total Size: sort that is based on the total size of a file – With Auto-extend: sort data files that have auto-extend specified – Without Auto-extend: sort data files that do not have auto-extend specified – By Discovered Time: sort that is based on when a file was discovered • **Users**: Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored Oracle instances within your environment.

Asset Report Type	What's in these reports?
SQL Server **Asset** > **System-wide** > **SQL/Server**	• **Instances**: Expand the **Instances** node and click the sorting method by which you want to view instances. – By Instance: sort that is based on Instance name – By Version: sort that is based on the version of the DBMS under which an Instance is running – By Probe Time: sort that is based on when an Instance was last probed • **Databases**: Expand this node and click the sorting method by which you want to view information about databases within your environment. – By Database: sort that is based on database name – By Free space: sort that is based on the amount of free space on a database – By Total Size: sort that is based on the total size of a database – By Table Count: sort that is based on the number of tables in the database – By Scan Time: sort that is based on when a database was last scanned – By Discovered Time: sort that is based on when a database was discovered – Dropped Databases: sort that is based on dropped databases • **Tables**: Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans display within the reports that are viewed through this node. – By Table: sort that is based on table name – By Number of Indexes: sort that is based on the number of indexes that are associated with a table – By Total Size: sort that is based on the total size of a table – By Row Count: sort that is based on the number of rows in a table – By Discovered Time: sort that is based on when a table was discovered – Dropped Tables: sort that is based on dropped tables • **Data Files**: Expand the **Data Files** node and click the sorting method by which you want to view information about data files. – By Data File: sort that is based on file name – By Free Space: sort that is based on file free space – By Total Size: sort that is based on the total size of a file – With Auto-extend: sort data files that have auto-extend specified – Without Auto-extend: sort data files that do not have auto-extend specified – By Discovered Time: sort that is based on when a file was discovered • **Users**: Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored SQL Server instances. – By Users: sort that is based on user names – By Discovered Time: sort that is based when a user was discovered

Asset Report Type	What's in these reports?
Sybase **Asset > System-wide > Sybase**	• **Instances**: Expand the **Instances** node and click the sorting method by which you want to view instances. – By Instance: sort that is based on Instance name – By Version: sort that is based on the version of the DBMS under which an Instance is running – By Probe Time: sort that is based on when an Instance was last probed • **Devices**: Expand this node and click the sorting method by which you want to view information about devices within your environment. – By Device: sort that is based on device name – By Free space: sort that is based on the amount of free space on a device – By Total Size: sort that is based on the total size of a device – By Fragments: sort that is based on the number of fragments for a device – By Discovered Time: sort that is based on when a device was discovered – Dropped Devices: sort that is based on dropped devices • **Databases**: Expand this node and click the sorting method by which you want to view information about databases within your environment. – By Database: sort that is based on database name – By Free space: sort that is based on the amount of free space on a database – By Total Size: sort that is based on the total size of a database – By Table Count: sort that is based on the number of tables in the database – By Scan Time: sort that is based on when a database was last scanned – By Discovered Time: sort that is based on when a database was discovered – By Last Dump Time: sort that is based on the date/time when devices were last dumped – Dropped Databases: sort that is based on dropped databases • **Fragments**: Expand this node and click the sorting method by which you want to view information about fragments. – By Fragments: sort that is based on fragment name – By Free Space: sort that is based on the free space in fragments – By Total Size: sort that is based on the total size of fragments – By Discovered Time: sort that is based on when fragments were discovered • **Tables**: Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans display within the reports that are viewed through this node. – By Table: sort that is based on table name – By Number of Indexes: sort that is based on the number of indexes that are associated with a table – By Total Size: sort that is based on the total size of a table – By Row Count: sort that is based on the number of rows in a table – By Chained Row Count: sort that is based on the number of chained rows in a table – By Discovered Time: sort that is based on when a table was discovered – Dropped Tables: sort that is based on dropped tables • **Users**: Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored Sybase instances.

Asset Report Type	What's in these reports?
UDB **Asset** > **System-wide** > **UDB**	• **Instances**: Expand the **Instances** node and click the sorting method by which you want to view instances. – By Instance: sort that is based on Instance name – By Version: sort that is based on the version of the DBMS under which an Instance is running – By Probe Time: sort that is based on when an Instance was last probed • **Databases**: Expand this node and click the sorting method by which you want to view information about databases within your environment. – By Database: sort that is based on database name – By Free space: sort that is based on the amount of free space on a database – By Total Size: sort that is based on the total size of a database – By Table Count: sort that is based on the number of tables in the database – By Scan Time: sort that is based on when a database was last scanned – By Discovered Time: sort that is based on when a database was discovered – Dropped Databases: sort that is based on dropped databases • **Table spaces**: Expand this node and click the sorting method by which you want to view information about table spaces. – By Table space: sort that is based on table space name – By Free space: sort that is based on the free space in table spaces – By Total Size: sort that is based on the total size of table spaces – By Table Count: sort that is based on the number of tables that are associated with table spaces – By Scan Time: sort that is based on the date/time when table spaces were scanned – By Discovered Time: sort that is based on the date/time when table spaces were discovered – Dropped Table spaces: sort that is based on the names of dropped table spaces • **Tables**: Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans display within the reports that are viewed through this node. – By Table: sort that is based on table name – By Number of Indexes: sort that is based on the number of indexes that are associated with a table – By Total Size: sort that is based on the total size of a table – By Row Count: sort that is based on the number of rows in a table – By Chained Row Count: sort that is based on the number of chained rows in a table – By Discovered Time: sort that is based on the date/time when a table was discovered – Dropped Tables: sort that is based on the names of dropped tables

Asset Report Type	What's in these reports?
	• **Containers**: Expand the **Containers** node and click the sorting method by which you want to view containers. – By Container: sort that is based on container name – By Free space: sort that is based on the free space in containers – By Total Size: sort that is based on the total size of containers – With Auto–extend: sort that is based on whether containers with Auto-extend on – Without Auto–extend: sort that is based on whether containers with Auto-extend off – By Discovered Time: sort that is based on the date/time when container were discovered • **Users**: Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored UDB instances.

Capacity reports:

Use Capacity reports to view the storage capacity, used space, and free space for your RDBMS objects.

Capacity enable you to view storage capacity at the computer, instance, database, device, fragment, container, and data file level. For a specific Instance, computer, group of computers, database, group of databases, tablespace, group of tablespaces, datafiles, containers, fragments, or all the computers in your enterprise, you can view the following:

- How much storage you have for a DBMS and the objects it contains
- How much storage is being used for a DBMS and the objects it contains
- What percentage of DBMS storage is free for expansion

All DBMSs Capacity reports:

Use these reports to view storage capacity information about the different RDBMS instances that you are monitoring within your environment.

Capacity Reporting View	Use these reports to...
Capacity > All DBMSs > Total Instance Storage	If you have multiple RDBMSs that you monitor using Data Manager, use this node to view detailed storage information about the different RDBMS instances within your environment. Depending on the RDBMS of the instance you want to view, use this report to view information about its: • RDBMS instance • RDBMS type • total size • file capacity • file free space • DB-TS count • file count • log file count

Oracle Capacity reports:

Use these reports to view storage capacity information about the Oracle instances that you are monitoring within your environment.

Capacity Reporting View	Use these reports to...
Oracle Capacity > Oracle **By Clauses** for these reports: • By Instance: sort based on Instance name • By Database: sort a report by database name • By Tablespace: sort a report by tablespace name • By Tablespace Group: view the desired storage information organized by tablespace group • By Datafile: sort based on file name • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor Oracle instances. Expand this node to view the following information: • **Total Instance Storage.** Generate a report under this node to view detailed information about Oracle instances, including: – total size – file capacity – file free space – DB-TS count – file count – log file count • **Data File Used Space.** Generate a report under this node to view detailed information about the storage usage of datafiles associated with the tablespaces in an Instance, including: – tablespace – data file path – used space – percent used – free space – total size – free extents – coalesced extents – maximum free extent size • **Data File Free Space**. Generate a report under this node to view detailed information about the storage free space of datafiles associated with the tablespaces in an Instance, including: – tablespace – data free space – percent free – data used space – data size – free extents – coalesced extents – maximum free extent size

SQL Server Capacity reports:

Use these reports to view storage capacity information about the SQL Server instances that you are monitoring within your environment.

Capacity Reporting View	Use these reports to...
SQL Server Capacity > SQL/Server **By Clauses** for these reports: • By Instance: sort based on Instance name • By Database: sort a report by database name • By Database Group: view the desired storage information organized by database group • By Datafile: sort based on file name • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor SQL Server instances. Expand this node to view the following information: • **Total Instance Storage.** Generate a report under this node to view detailed information about SQL Server instances, including: – total size – data file capacity – data file free space – log file capacity – log file free space – database count – data file count – log file count • **Data File Used Space.** Generate a report under this node to view detailed information about the storage usage of datafiles associated with the databases in an Instance, including: – database – data used space – percent used – data free space – data size – log used space – log percent used – log free space – log size – number data files • **Data File Free Space**. Generate a report under this node to view detailed information about the storage free space of datafiles associated with the databases in an Instance, including: – database – data free space – percent free – data used space – data size – log free space – log percent used – log used space – log size – number data files

Sybase Capacity reports:

Use these reports to view storage capacity information about the Sybase instances that you are monitoring within your environment.

Capacity Reporting View	Use these reports to...
Sybase Capacity > Sybase **By Clauses** for these reports: • By Device: sort based on Device name • By Instance: sort based on Instance name • By Fragments: sort based upon fragments • By Database: sort a report by database name • By Database Group: view the desired storage information organized by database group • By Datafile: sort based on file name • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor Sybase instances. Expand this node to view the following information: • **Total Instance Storage.** Generate a report under this node to view detailed information about Sybase instances, including: – total device size – mirrored device size – device free space – data fragment capacity – data fragment free space – log fragment capacity – log fragment free space – device count – database count – fragment count – log fragment count • **Device Used Space.** Generate a report under this node to view detailed information about the storage usage of devices associated with instances, including: – Instance – device path – used space – percent used – free space – total size – number fragments – mirror path – number devices • **Device Free Space**. Generate a report under this node to view detailed information about the storage free space of devices associated instances, including: – Instance – device path – free space – percent free – used space – total size – number fragments – low, high – mirror path – number devices

Capacity Reporting View	Use these reports to...
Sybase (continued)	• **Fragment Used Space.** Generate a report under this node to view detailed information about the storage usage of fragments associated with databases, including: – Instance – fragment path – database – data used space – percent used – data free space – data size – data type – log used space – log percent used – log free space – log size – number fragments • **Fragment Free Space**. Generate a report under this node to view detailed information about the storage free space of fragments associated databases, including: – Instance – fragment path – database – data free space – percent free – data used space – data size – log free space – log used space – log percent used – log used space – log size – number fragments

UDB Capacity reports:

Use these reports to view storage capacity information about the UDB instances that you are monitoring within your environment.

Capacity Reporting View	Use these reports to...
UDB Capacity > UDB **By Clauses** for these reports: • By Container: sort based on Container name • By Tablespace: sort based on Tablespace name • By Database: sort a report by database name • By Database Group: view the desired storage information organized by database group • By Instance: sort based on Instance name • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor UDB instances. Expand this node to view the following information: • **Total Instance Storage.** Generate a report under this node to view detailed information about UDB instances, including: – total size – container capacity – container free space – log file capacity – tablespace count – container count – log file count – log mode • **Container Used Space.** Generate a report under this node to view detailed information about the storage usage of containers associated with a tablespace, including: – Instance – database – tablespace – container node – container name – used space – percent used – free space – total size – container type – status – discovered time • **Container Free Space**. Generate a report under this node to view detailed information about the storage free space of containers associated with a tablespace, including: – Instance – database – tablespace – container node – container name – free space – percent free – used space – total size – container type – status – discovered time

Usage reports:

Use Usage reports to monitor and report on the usage and growth of the storage consumption for your databases, tablespaces, tables, segments, and users.

Use these reports to do the following:

- View detailed information for tablespaces and databases including their total size, used space, empty used space, # of tables and indexes, etc.
- Identify the tables consuming the most space within your environment. You can view this information at the tablespace, database, and computer level, as well as at a network-wide level.
- (Oracle) Identify the segments with the most extents and most empty used space. You can view this information at the tablespace, database, and computer level, as well as at a network-wide level.
- Associate users with the amount of storage they are consuming on specific databases, groups of databases, tablespaces, groups of tablespaces, specific computers, groups of computers, and throughout the entire network.

All DBMS Usage reports:

Use these reports to view storage usage information about the different RDBMS instances that you are monitoring within your environment.

Usage Reporting View	**Use these reports to...**
Usage > All DBMSs	If you have multiple RDBMSs that you monitor using Data Manager, use this node to view detailed storage information about the different RDBMS instances within your environment.
Usage > All DBMSs > Databases-Tablespaces	view storage usage statistics about the databases and tablespaces with your monitored DBMSs. Some of the information you can view through these reports includes: • database name, tablespace name • computer • Instance • Instance type • total size • used space • empty used space • table count • index count • other count
Usage > All DBMSs > Tables	view storage usage statistics about the tables with your monitored DBMSs. Some of the information you can view through these reports includes the following: • Largest tables • Total size • Creator • Table name • Empty used space • Index size • Row count • Chained row count

Usage Reporting View	Use these reports to...
Usage > All DBMSs > Users	associate users with the amount of storage they are consuming on specific DBMS objects, specific computers, groups of computers, and throughout the entire network. • **Group by User Name**: select the users and then the resources on which you want to report storage consumption. For example, when setting up a report to view user storage consumption on an Oracle or UDB tablespace, you must do the following: – Select the users upon which you want to report – Select the tablespaces whose storage consumption by the selected users you want to view • **Group by Resource**: select the resources and then the users on which you want to report storage consumption. For example, when setting up a report to view storage consumption on Oracle or UDB tablespaces for users, you must do the following: – Select the tablespaces upon which you want to report – Select the users whose storage consumption on the selected tablespaces you want to view

Oracle Usage reports:

Use these reports to view storage usage information about the Oracle tablespaces, tables, segments, and users that you are monitoring within your environment.

Usage Reporting View	Use these reports to...
Oracle Usage > Oracle **By Clauses** for these reports: • By Database—sort a report by database name • By Tablespace—sort a report by tablespace name • By Tablespace Group—view the desired storage information organized by tablespace group • By Table—sort a report by table name • By Table Group—view the desired storage information organized by table group • By Computer—sort based on the computer where the object's Instance resides • By Computer Group—view the desired storage information organized by computer group • Network-wide—view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor Oracle instances. Expand this node to view the following information: • **Tablespaces.** View detailed information for tablespaces including the following: – Instance – Database – Total size – Used space – Empty used space – Freelist blocks – Table count – Index count • **Tables.** Identify the tables consuming the most space within your environment. You can view this information at the tablespace, database, and computer level, as well as at a network-wide level. Some of the information in these reports includes the following: – Table name – Total size – Empty used space – Index size – Freelist blocks – Number partitions – Number indexes – Number segments – Number extents – Percent free, used

Usage Reporting View	Use these reports to...
Oracle Usage → Oracle **By Clauses** for these reports: • By Database—sort a report by database name • By Tablespace—sort a report by tablespace name • By Tablespace Group—view the desired storage information organized by tablespace group • By Table—sort a report by table name • By Table Group—view the desired storage information organized by table group • By Computer—sort based on the computer where the object's Instance resides • By Computer Group—view the desired storage information organized by computer group • Network-wide—view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor Oracle instances. Expand this node to view the following information: • **Segments**. Identify the segments with the most extents and most empty used space. Some of the information in these reports includes the following: – Number extents – Segment name, creator – Tablespace – Parent name, creator – Total size – Empty used space – Number extents – Freelist blocks, groups • **Users**. Associate users with the amount of storage they are consuming on specific databases, tablespaces, groups of tablespaces, specific computers, groups of computers, throughout the entire network, and applied Quotas. Some of the information included in these reports includes the following: – Computer, instance, database, tablespace – Total size – Empty used space – Table count – Index count – Largest segment, 2nd largest segment – Segment sizes: <50Kb, 500Kb-100Mb, 1Mb-10Mb, 10Mb-100Mb, 100Mb-500MB, 500Mb-1Gb, >1Gb – Quota name – Percent of Quota – Quota threshold

SQL Server Usage reports:

Use these reports to view storage usage information about the SQL Server databases, tables, and users that you are monitoring within your environment.

Usage Reporting View	Use these reports to...
SQL Server Usage > SQL/Server **By Clauses** for these reports: • By Instance: sort based on Instance name • By Database: sort a report by database name • By Database Group: view the desired storage information organized by database group • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network • By User Quota: sort based upon the users to which Quotas are applied. A user who has not violated a quota will not appear in this report.	This node appears if you are using Data Manager for Databases to monitor storage usage on SQL Server instances. Expand this node to view the following information: • **Databases.** Generate a report under this node to view detailed information about the storage on SQL Server databases, including the following: – Total size – Used space – Empty used space – Table count – Index count • **Tables.** Identify the tables consuming the most space within your environment. You can view this information at the table, database, and computer level, as well as at a network-wide level. Some of the information in these reports includes the following: – Table name – Total size – Empty used space – Index count – Row count • **Users**. Associate users with the amount of storage they are consuming on specific databases, groups of databases, specific computers, groups of computers, throughout the entire network, and applied Quotas . Some of the information in these reports includes the following: – Computer, instance, database – Total size – Empty used space – Table count – Index count – Largest object, 2nd largest object – Object sizes: <50Kb, 500Kb-100Mb, 1Mb-10Mb, 10Mb-100Mb, 100Mb-500MB, 500Mb-1Gb, >1Gb – Quota name – Percent of Quota – Quota threshold

Sybase Usage reports:

Use these reports to view storage usage information about the Sybase databases, tables, and users that you are monitoring within your environment.

Usage Reporting View	Use these reports to...
Sybase Usage > Sybase • By Database: sort a report by database name • By Database Group: view the desired storage information organized by database group • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network • By User Quota: sort based upon the users to which Quotas are applied. A user who has not violated a quota will not appear in this report.	This node appears if you are using Data Manager for Databases to monitor storage usage on Sybase instances. Expand this node to view the following information: • **Databases.** Generate a report under this node to view detailed information about the storage on Sybase databases, including the following: – Database name, computer, instance – Total size – Used space – Empty used space – Table count – Index count • **Tables.** Identify the tables consuming the most space within your environment. You can view this information at the table, database, and computer level, as well as at a network-wide level. Some of the information in these reports includes the following: – Table name, database, computer, Instance, creator – Total size – Empty used space – Index size – Maximum rows per page – Expected row size – Fill factor – Page gap – Average row length – Row count • **Users**. Associate users with the amount of storage they are consuming on specific databases, groups of databases, specific computers, groups of computers, throughout the entire network, and applied Quotas. Some of the information included in these reports includes the following: – Computer, instance, database – Total size – Empty used space – Table, index count – Largest object, 2nd largest object – Object sizes: <50Kb, 500Kb-100Mb, 1Mb-10Mb, 10Mb-100Mb, 100Mb-500MB, 500Mb-1Gb, >1Gb – Quota name – Percent of Quota – Quota threshold

UDB reports:

Use these reports to view storage usage information about the UDB tablespaces, tables, and users that you are monitoring within your environment.

Usage Reporting View	Use these reports to...
UDB Usage > UDB • By Tablespace: sort a report by tablespace name • By Tablespace Group: view the desired storage information organized by tablespace group • By Database: sort a report by database name • By Instance: sort a report by Instance name • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor storage usage on UDB instances. Expand this node to view the following information: • **Tablespaces.** Generate a report under this node to view detailed information about the storage on UDB tablespaces, including the following: – Database name, computer, instance – Total size – Used space – Empty used space – Table count – Index count • **Tables.** Identify the tables consuming the most space within your environment. You can view this information at the table, database, and computer level, as well as at a network-wide level. Some of the information in these reports includes the following: – Table name, database, computer, Instance, creator – Total size – Empty used space – Index size – Maximum rows per page – Expected row size – Fill factor – Page gap – Average row length – Row count **Note:** When viewing reports under the **Monitored Tables** node, keep in mind that a table will appear in this report if: 1) it is included in the definition of a table group and 2) the table group in which it is included is defined as part of a scan. • **Users**. Associate users with the amount of storage they are consuming on specific UDB instances, tablespaces, groups of tablespaces, computers, groups of computers, and throughout the entire network. Some of the information included in these reports includes the following: – Computer, instance – Total size – Empty used space – Table, index count – Largest object, 2nd largest object – Object sizes: <50Kb, 500Kb-100Mb, 1Mb-10Mb, 10Mb-100Mb, 100Mb-500MB, 500Mb-1Gb, >1Gb – Quota name – Percent of Quota – Quota threshold

Usage Violation reports:

Use Usage Violation reports to view information about quota violations and help enforce corporate storage usage policies.

With quotas, you can control how much storage a user, or a group of users, can consume on the following:

- Database or group of databases
- Tablespace or group of tablespaces

- Instance
- Throughout the entire network

The following table contains information about Usage Violation reports:

Usage Violation Reporting Type	Use these reports to...	What's in these reports?
Quota Violations Usage Violations > Quota Violations	identify the quota violations within the environment	
All Quotas Usage Violations > Quota Violations > All Quotas	view all the quotas and quota violations in the environment • **By User**. View information about the users who have violated quotas. • **By User Quota**. View information about the quotas that have been violated.	• User name • RDBMS type • Instance • Database-tablespace • Space used • Threshold • % of threshold • Number of violations • Quota creator • Quota name • Quota type
Network Wide Quotas Usage Violations > Quota Violations > Network-wide Quotas	view quota violations at the network-wide level • **By User**. View information about the users who have violated quotas. • **By User Quota**. View information about the quotas that have been violated and which users have violated them.	• User name • Space used • Threshold • % of threshold • Number of violations • Quota creator • Quota name • Quota type
Instance Quotas Usage Violations > Quota Violations > Instance Quotas	view quota violations at the Instance level • **By Instance**. View Instance quota violations sorted by Instance • **By Computer**. View Instance quota violations sorted by computer • **By Computer Group**. View Instance quota violations sorted by computer group • **By User**. View Instance quota violations sorted by users • **By User Quota**. View Instance quota violations sorted by quota	• Computer • RDBMS type • Group creator (by group clauses only) • Group name (by group clauses only) • Instance • Database-tablespace • User name • Space used • Threshold • % of threshold • Number of violations • Quota creator • Quota name • Quota type

Usage Violation Reporting Type	Use these reports to...	What's in these reports?
Database-Tablespace Quotas Usage Violations > Quota Violations > Database-Tablespace Quotas	view quota violations at the tablespace level • **By Database-Tablespace**. View quota violations sorted by the database or tablespace name on which the violation occurred • **By Database-Tablespace Group**. View quota violations sorted by the database or tablespace group in which the violation occurred • **By Instance**. View database or tablespace quota violations sorted by Instance • **By Computer**. View database or tablespace quota violations sorted by computer • **By Computer Group**. View database or tablespace quota violations sorted by computer group • **By User**. View database or tablespace quota violations sorted by users • **By User Quota**. View database or tablespace quota violations sorted by quota	• Computer • Group creator (by group clauses only) • Group name (by group clauses only) • Instance • Database-tablespace • User name • Space used • Threshold • % of threshold • Number of violations • Quota creator • Quota name • Quota type

Data Manager for Chargeback reports

Use invoices generated by Data Manager for Chargeback to charge users or groups of users for storage usage within a storage environment.

See "Working with invoices" on page 620 for information about how to use Data Manager for Chargeback to generate invoices.

Disk Manager reports

Use Disk Manager reports to view the relationships among the components of a storage subsystem and the storage capacity. The capacity can be for a computer, file system, storage subsystem, volume, Array Sites, Ranks, and disk.

Reporting on storage subsystems:

Learn how to use Disk Manager to report on storage subsystems.

Through these reports, you can view the relationships among the components of a storage subsystem. Storage subsystems reports use data that is gathered by probes.

To ensure the accuracy and consistency of the data that is displayed in storage subsystem reports, run regularly scheduled probes against the hosts that use or import volumes, and the storage subsystems upon which you want to report. Run probes against storage subsystems after you run them against any host.

- A disk must be partitioned before it is considered a logical volume by Disk Manager. If a disk is not partitioned, it is not included in a file system or logical volumes report.
- Before you begin, you might want to refer to Generating and viewing reports. That section provides information on reports that applies all IBM Tivoli Storage

Productivity Center reports. It describes the types of reports: system, user-defined, and batch. It also describes how to save reports and how to customize reports, such as filtering information and displaying the information in chart form.

You can view storage subsystem reports in the following ways:

Disks Use these reports to view and chart the storage capacity of storage subsystem disks. Use By clauses to filter the storage entities that are displayed in a report. For example, if you want to view a report that displays specific disks, select **By Disk**; if you want to view the disks of specific storage subsystems, select **By Storage Subsystem**.

Volumes
Use these reports to view and chart the storage capacity of storage subsystem volumes. Use the **By Storage Subsystems** and **By Storage Subsystem Group** clauses to filter the storage entities that display in a report. Use the **By Volume Space** and **By Probe Time** clauses to sort the volumes that display in a report according to when a volume was probed or the capacity of a volume.

Storage Pools
Use these reports to view and chart the storage capacity of storage pools. Use By clauses to filter the storage entities that display in a report. For example, if you want to view a report that displays information about the storage pools for specific storage systems, select **By Storage Subsystem**.

Array Sites
Use these reports to view information about the array sites that are on storage subsystems. Use By clauses to filter the storage entities that display in a report. For example, if you want to view a report that displays about the array sites for specific groups of storage subsystems, select **By Storage Subsystem Group**.

Disk Groups
Use these reports to view information about the disk groups that are on storage subsystems. Use the By clauses to filter the storage entities that display in a report. For example, if you want to view a report that displays about the disk groups for specific storage subsystems, select **By Storage Subsystem**.

Ranks Use these reports to view information about the ranks that are on storage subsystems. Use the By clauses to filter the storage entities that display in a report. For example, if you want to view a report that displays about the ranks for specific groups of storage subsystems, select **By Storage Subsystem Group**.

Computer Views
Use these reports to view the relationships between the computer and the storage subsystems, volumes, and disks in your environment. The information is gathered from the viewpoint of the computer. For reporting purposes, *computer* is defined as a system that is running Tivoli Storage Productivity Center.

Restriction: These reports are available only if the following steps are done.

- A Storage Resource agent is installed on a system where a volume is surfaced
- A probe is run against the system where the volume is located

- A probe is run against the storage subsystem or IBM SAN Volume Controller cluster

By Computer

These reports view relationships between the computer and the storage subsystems, volumes, or disks in your environment. To generate a report by computer, select the **By Computer** node and choose an item from the **Relate Computers to:** field.

For example, to retrieve information on computers in your environment and their volumes, select the **By Computer** node and then choose **By volumes**. These settings generate a report that details all computers with agents on them, which use volumes from the monitored storage subsystems.

By Computer Group

These reports view relationships between the computer groups (both monitoring and reporting groups) and the storage subsystems, volumes, or disks in your environment. To generate a report by computer group, select the **By Computer group** node and choose an item from the **Relate Computers to:** field.

For example, to retrieve information on computers in a group that use storage from monitored storage subsystems, select the **By Computer Group** node and then choose **Storage Subsystems** from the **Relate Computers to:** field. These settings generate a report that details all computers in groups with agents on them, which use volumes from the monitored storage subsystems.

By File System/Logical Volume

These reports view the relationships between the file systems or logical volumes and the storage subsystems, volumes, and disks in your environment. To generate a report for the relationship you want to view, select the **By File System/Logical Volume** node and choose an item from the **Relate File System/Logical Volume to:** field.

For example, to retrieve information on file systems or logical volumes that use storage from monitored storage subsystems, select the **By File System/Logical Volume** node and choose **By Storage Subsystems**. These settings generate a report that details all file systems or logical volumes that use storage from monitored storage subsystems.

By File System Group

Use this report to view the relationships between the file system groups (both monitoring and reporting) and the storage subsystems, volumes, and disks in your environment. To generate a report for the relationship you want to view, select the **By File System Group** node and choose an item from the **Relate File System/Logical Volume to:** field.

For example, to retrieve information on file systems in a group that use storage from monitored storage subsystems, select the **By File System Group** node and choose **Storage Subsystems** from the **Relate File System/Logical Volume to:** field. These settings generate a report that details all file systems in a group that use storage from monitored storage subsystems.

Storage Subsystem Views

Use these reports to view the relationships among the storage subsystem components (storage subsystems, volumes, and disks) to the computers and file systems or logical volumes in your environment.

Note: These reports are available only if the following steps are performed.

- A Storage Resource agent is installed on a system where a volume is surfaced
- A probe is run against the system where the volume is located
- A probe is run against the storage subsystem or IBM SAN Volume Controller cluster
- For IBM XIV Storage System, there is no relationship between storage disks, pools, and volumes because of the unique storage disk distribution and architecture.

By Storage Subsystem

These reports view the relationships between the storage subsystems and the computers and file systems or logical volumes in your environment. To generate a report for the relationship you want to view, select the **By Storage Subsystem** node and choose an item from the **Relate Storage Subsystem to:** field.

For example, to view all storage subsystems that have storage that is being used by computers in the environment, select **By Storage Subsystem** and choose **Computers**. These settings generate a report that details all storage subsystems that have storage that is being used by computers in the environment.

By Storage Subsystem Group

These reports view the relationships between storage subsystem groups (both monitoring and reporting) and the computers and file systems or logical volumes in your environment. To generate a report for the relationship you want to view, select the **By Storage Subsystem Group** node and choose an item from the **Relate Storage Subsystem to:** field.

For example, to view all storage subsystem groups that have storage that is used by computers in the environment, select **By Storage Subsystem Group** and choose **Computers** from the **Relate Storage Subsystem to:** field. These settings generate a report that details all storage subsystem groups that have storage that is used by computers in the environment.

By Volume

Use this report to view the relationships between the volumes and the computers and file systems or logical volumes in your environment. To generate a report for the relationship you want to view, select the **By Volume** node and then choose an item from the **Relate Volumes to:** field.

For example, to get a listing of all volumes in the environment on which file systems or logical volumes are created, select **By Volume** and select **File Systems/Logical Volumes** from the **Relate Volumes to:** field. These settings generate a report that details all volumes in the environment on which file systems or logical volumes are created.

To get a listing of all volumes in the environment that are being used by computers in the environment, select **By Volume** and choose **Computers**. These settings generate a report that details all volumes in the environment that are being used by computers in the environment.

Disks Use this report to view the relationships between disks and the computers and file systems or logical volumes in your environment. To generate a report for the relationship you want to view, select the **Disks** node and then choose an item from the **Relate Disks to:** field.

For example, to get a listing of all storage subsystem disks in the environment on which file systems or logical volumes are created, select **Disks** and select **File Systems/Logical Volumes**. These settings generate a report that details all storage subsystem disks in the environment on which file systems or logical volumes are created.

Volume to HBA Assignment

Use these reports to view all the logical unit numbers (volumes) of each probed storage subsystem that is assigned to at least one host bus adapter (HBA) port worldwide name (WWN). These reports retrieve data only from systems that have HBAs.

The HBA is a Fibre Channel and volumes are physically attached to the storage subsystem.

By Storage Subsystem

This report shows all the volumes of each probed storage subsystem that is assigned to at least one HBA Port WWN. This report is sorted by storage subsystem name.

By Volume Space

This report views all the volumes of each probed storage subsystem that is assigned to at least one HBA Port WWN. This report is sorted by volume space.

Not on Monitored Computers

This report shows only volumes that are assigned to HBA Port WWNs on hosts with no Storage Resource agent installed on them.

Not Visible to Monitored Server

Use this report to find volumes that are assigned to hosts with Storage Resource agents, but not configured. This report can be used to generate a list of all volumes that must be configured before they are used.

Volume to Backend Volume Assignment By Volume

Use this report to view the relationships between the virtual disks (including mirror copies) and managed disks that are associated with a virtualizer storage subsystem, such as the SAN Volume Controller, and the back-end storage volumes that are associated with those managed disks.

Note: A virtual disk with one mirror copy is displayed as one row in the table. A virtual disk with multiple mirror copies is displayed in multiple rows in the table. Each row represents one of the copies.

Important: Some of the columns in the totals row for these reports can count storage multiple times if the report includes information.

- A storage virtualizer like a SAN Volume Controller
- The disk arrays that host the backend storage to a storage virtualizer like a SAN Volume Controller

In this situation, disks and logical unit numbers (LUNs) seem to belong to both a storage virtualizer like SAN Volume Controller and the disk arrays. LUNs are counted multiple times in some of the columns in the totals row. To avoid this situation, filter the report to display only the storage virtualizer or the disk arrays. If this situation is unavoidable, then some of the columns in the totals row can be adjusted as follows:

1. Use the **Physical Disk Space** column to determine the actual physical disk space on all arrays. This disk space is not counted multiple times.
2. Sum the **Disk Space** column for the disk arrays and subtract this result from the totals row to determine the total backend storage for a virtualizer.
3. To prevent counting the back-end volumes multiple times, subtract the **Backend Volume Space** column from the **Volume Space** column in the totals row to determine the amount of volume space.

In the topology of a subsystem, the **Backend Volume Space** column displays the total amount of back-end space available to the subsystem. This space is the total capacity of the managed disks on the SAN Volume Controller, not the internal disk storage space. In other reports, Backend Volume Space refers to the space that a storage subsystem makes available to a storage virtualizer. These reports include:

- Reports that are generated through Data Manager.
- Reports that are generated through Disk Manager.
- Rollup reports.
- **Data Manager** > **Reporting** > **Asset** > **By Storage Subsystem** reports.

Generating storage-subsystem reports:

Learn how to generate a storage subsystem report.

These reports are available only if the following steps have been performed:

1. An agent is installed on a machine where a volume is surfaced.
2. A probe has been run against the machine where the volume is located.
3. A probe has been run against the storage subsystem or SAN Volume Controller cluster.

To generate a storage-subsystems report, do the following:

1. In the navigation tree, expand **Disk Manager** > **Reporting** > **Storage Subsystems**.
2. Do one of the following:
   a. In the navigation tree, expand **Computer Views** and click one of the following:
      - By Computer
      - By Computer Group
      - By Filesystem/Logical Volume
      - By Filesystem Group
   b. In the navigation tree, expand **Storage Subsystems Views** and click one of the following:
      - By Storage Subsystem

- By Storage Subsystem Group
- By Volume
- Disks

c. In the navigation tree, expand **Volume to HBA Assignment** and click one of the following:
   - By Storage System
   - By Volume Space
   - Not On Monitored Computers
   - Not Visible to Monitored Computers

d. In the navigation tree, expand **Volume to Backend Volume Assignment** and click the following:
   - By Volume

3. The **Storage Subsystem Performance** pane for your selected storage subsystem opens the Select Resources pane. All valid available resources are listed in Included Columns. In Available Columns and Included Columns, use and to select the columns you want to include in your report.
4. Filters can be applied to the selected columns. In the **Report Filter Specification** area, the **Filter** button opens the **Edit Filter** window.
   a. Click **Add** to add a column to the filter list. The columns will be added in the order they appear on your report.
   b. Click **Delete** to delete all columns. Click the X for an individual column to delete it.
   c. In the **Records Must Meet** area, select **All Conditions** or **At Least One Condition**.
   d. Type the condition value in the **Value 1** text field.
      - If the text field is greyed-out, click **Edit**.
      - If you wish to add additional values for the same column using the same operator, click **Edit**.
   e. Select the Operator.
      - If you select the BETWEEN or NOT BETWEEN operators, two value text fields will open for a lower-limit value and an upper-limit value. Click **Edit** to enter the upper-limit value and lower-limit value. You must input values in both **Value 1** and **Value 2** fields, if shown.
      - For numeric columns, such as Free Space, a pull-down menu at the end of the each field assists you with numerical units.
   f. Once your selections have been made, click **OK**.
5. Click **Generate Report**. The report appears under a new tab window. If there is no data, as in if a probe has not been run, the window will say "No Data To Display."
6. Click **File** > **Save**. The report is saved under **IBM Tivoli Storage Productivity Center** > **My Reports**. Reports can be regenerated from that node.

To change the parameters for your report, or create a new report of the same type, do the following:

1. Click the **Selection** tab before closing the Storage Subsystem Performance window. Make the necessary changes.
2. Click **Generate Report**.
3. Click **File** > **Save** and enter the name of the new report.

4. Click **OK**.

Creating a storage subsystem performance report:

Learn how to create a storage-subsystem performance report.

Before you can view a performance report on a storage subsystem, that subsystem must have been discovered and had performance data successfully collected.

For a complete list of metrics available for performance reports, see "Performance metrics" on page 1701.

1. In the navigation tree pane, expand **Disk Manager** > **Reporting** > **Storage Subsystem Performance**.
2. Click the type of performance report that you want to view. You can choose from the following performance reports:

   By Storage Subsystem
   This report contains performance information for storage subsystems.

   By Controller
   This report contains performance information for controllers.

   By I/O Group
   This report contains performance information for I/O groups.

   By Node
   This report contains historical performance data for SAN Volume Controller (SVC) nodes.

   By Array
   This report contains performance information for arrays.

   By Managed Disk Group
   This report contains performance information for managed-disk groups.

   By Volume
   This report contains performance information for volumes.

   By Managed Disk
   This report contains performance information for managed disks.

   By Port
   This report contains performance information for ports.

   In the content pane, the Selection page opens.
3. Select from the following options to customize the report:

   Selection button
   Click this button to determine the resources upon which you want to report, such as the storage subsystems, the fibre-channel ports, and so forth. The Select Resources window displays. The type of resources displayed for selection depends on the type of report you selected. For example, if you selected a By Storage Subsystem report, all displayed resources will be storage subsystems. If you clear the box for a resource on the Select Resources window, a pushpin icon displays on the Selection button.

   Filter button
   Click this button to further defines the data that you want to include in the report. Filters enable you to apply general rules to the report. When you click this button, the Edit Filter window displays.

For example, to select only data associated with subsystems that have a name ending with "123–IBM", select the "Subsystem" column, the "LIKE" operator, and type "*123–IBM" (without the quotation marks) in the Value 1 field. To select only data which has a Read I/O Rate greater than 0, select the "Read I/O Rate (overall)" column, the ">" operator, and type "0" (without the quotation marks) in the Value 1 field.

If a filter is defined for this report, a pushpin icon displays on the **Filter** button.

Time range selectors

All performance reports display historical performance data. Use the time range selectors to specify the exact time range that you want the report to cover. Click one of the following radio buttons to select the type of historical report you want:

Display latest performance data

This report contains only the latest available performance data for the selected resources. Note that even the latest data available might still be somewhat old, depending on whether data was recently collected for the resources. Including the time column in the report will help you determine how old the latest available data actually is. Regardless of the age of the data, clicking this button ensures that only a single row (representing a single point in time) for each selected resource is included in the report. If no data has ever been collected for a particular resource, then that resource is not included in the report. If no data has ever been collected for any resource, then the resulting report will be empty.

Display historic performance data using absolute time

This report contains only data that falls within a specific time range. Select the time range using the **From** and **To** selectors. The specified from and to values are interpreted as inclusive. The resulting report contains all the performance data that is available within the specified time range, for each selected resource. This means that the data can consist of potentially many rows per resource, each with a unique timestamp. If no data has been collected for a particular resource during the specified time range, that resource is not included in the report. If no data has been collected for any resource during the specified time range, then the resulting report will be empty.

To determine which historical performance records fall within the specified time range, the **beginning timestamp** of the collection intervals is used for comparison. If an historical record contains performance data for an interval of time from 1:00 PM to 1:05 PM, that record would be included in the report for a time range of 1:00 PM to 8:00 PM. However, that same record would not be included if a time range of 1:02 PM to 8:02 PM was specified, because it is the beginning of the time period (for example, 1:00 PM in this example) that is used to determine inclusion in the specified time interval.

Display historic performance data using relative time

This report contains only data that falls within a specific time range extending backward from the current date and time. Specify the number of days in the desired time range, using the **days ago until now** field. The report will be exactly the same

as a report using absolute time, with the **To** selector set to the current date and time, and the **From** selector set to the current date and time minus the number of **days ago until now**.

Note that the same criteria applies as if determining which historical performance records fall within the specified time range. Only the beginning timestamp of the collection intervals is used for comparison. This means that data for the current hour or the current day is not included in the report when selecting hourly or daily summation data, because data for the current hour or current day is only saved in the database when the hour or day is complete. As an alternative, consider using the first radio button, **Display latest performance data**, to see a report showing the most recent daily summation data.

Summation Level

Use this selector to select which type of performance data records to include in the report. The Tivoli Storage Productivity Center Performance Manager collects and stores three different types of performance data:

By sample

Represents the most detailed data. There is usually one sample record per resource for every **n** minutes, where **n** is the performance monitor interval length that was specified when the monitor was defined. Each sample represents the average performance of the resource over those **n** minutes. A sample record with a timestamp of **t** will be saved at time **t+n**. In other words, the beginning of each **n** minute time interval is used for display and for time range matching.

Hourly

Represents an hourly average of the performance of the specified resources. There will be one hourly record written every hour, per resource. Each such record will be written at or shortly after each hour mark on the (device) clock. For example, a record with a timestamp of 12:00 PM and interval length of 1 hour will be written at 1:00 PM.

Daily Represents a daily average of the performance of the resources. There will be one daily record written every day, per resource. Each daily record will be written at or shortly after midnight, according to the device clock. For example, a record with a timestamp of April 21, 12:00 AM and interval length of 24 hours will be written on April 22, 12:00 AM.

Note that hourly and daily records are usually retained in the database for longer periods of time, compared to By Sample records. As a result, if a generated report indicates that sample data is no longer available for a particular time range in the past, it is possible that the sample data has already been purged, in which case you can try generating the report using hourly or daily summation levels.

Be aware that when selecting a time range using the **Display historic performance data using relative time** radio button, it is necessary to select the length of the range in days. This can lead to unexpected results when also selecting the daily summation level. Remember that a daily record is not written until the day is complete, which means that at any given time during the day, the current day's data does not yet

exist in the database. Therefore, selecting the daily summation level as well as a relative time range of n **days ago until now**, will usually result in n-1 records being returned per resource, assuming that at least n-1 days worth of data has been previously collected for those resources. Specifying a relative time range of 1 **days ago until now** will result in an empty report.

Note: The user interface displays the date and time of data in the Tivoli Storage Productivity Center server time zone, but the daily and hourly aggregations are calculated based on the device time zones. Hourly records are written at the end of each hour according to the devices' clocks and the daily records are written at midnight according the the devices' clocks, but both are displayed according to the Tivoli Storage Productivity Center server's clock. For example, if a subsystem clock is one hour ahead of the Tivoli Storage Productivity Center server clock, the hourly data for 5:00 PM according to the device's clock is displayed as 4:00 PM in the user interface and the daily data for 12:00 AM according to the device's clock is displayed as 11:00 PM of the previous day.

Available Columns, Included Columns

Use these items to specify which columns to include and which columns to exclude from the generated report.

Available Columns

Displays the columns that you can include in the generated report that are not already included. If a column appears in the **Included Columns** list box, it will not appear in the **Available Columns** list box.

Included Columns

Displays the columns that will appear in the generated report. By default, all columns for a report are listed in this list box and will be displayed in the generated report. You can highlight the columns and use the up and down arrows to rearrange the order of the columns in the report. The order of the items in this list box determines the order in which the columns will appear in the generated report. The columns with grayed-out names are the frozen columns, which cannot be selected, removed, or reordered for this type of report.

If a column appears in the **Available Columns** list box, it will not appear in the **Included Columns** list box.

To exclude a column from a report, click the name of the column in the **Included Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the left-arrow button (<<). The selected columns are removed from the **Included Columns** list box, and will appear in the **Available Columns** list box.

To include a column in a report, click the name of the column in the Available Columns list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the right-arrow button (>>). The selected columns will be removed from the **Available Columns** list box, and will appear in the **Included Columns** list box.

To rearrange the order of columns that will appear in the generated report, click the name of a column in the **Available Columns** list box.

Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the up-arrow button (^) or the down-arrow button (v) to move the selected columns up or down in the list, relative to the other columns.

4. Click **Generate Report**. A new tab opens and displays the report.
5. Optional: To view the report in chart format, click , specify charting options, and then click **OK**. Right-click on the chart and select **Customize this Chart** to further customize how the chart is displayed.
6. Optional: To save the report settings to generate later, click and enter a name for the report. The report is saved under theIBM Tivoli Storage Productivity Center node. See "Filtering the resources that appear on a report" on page 1326 for instructions on modifying saved report settings.

If the report is greater than 2500 rows, a dialog box will appear which will allow you to retrieve additional rows by clicking **Display More Rows**. An additional 2500 rows will be retrieved from the server and displayed on the panel for each click of the button. Do not try to display too much data or there is the possibility of running out of memory.

Generating constraint violation reports:

Learn how to generate a storage subsystem performance report.

Creating constraint violation reports:

Learn how to how to create constraint violation reports.

1. In the navigation tree pane, expand **Disk Manager** > **Reporting** > **Storage Subsystem Performance**.
2. Click **Constraint Violations**. The constraint violation report selection page appears in the content pane.
3. If you want to limit the report to a particular time period, select **Limit to this time range** and specify the time period.
4. Click **Selection** to choose which storage subsystems to include in the report.
5. Click **Filter** to modify the filter used in generating the report.
6. Customize the performance report. By default, all supported metrics are included. Click to move the selected metric up in the list and to move the selected metric down in the list.
   a. In the **Included Columns** list, select a metric and click to remove it from the report.
   b. To change the order in which the columns appear in the report, select a metric and click either or .
7. Click **Generate Report**. The **Constraint Violations** tab opens and displays the report.

The report might show information about DS6000, DS8000, TotalStorage Enterprise Storage Server, SAN Volume Controller, and other subsystems all on the same report. By default, the report displays one subsystem per row. Results are aggregated for the current day.

Viewing detailed constraint violation information:

Learn how to view detailed information about constraint violations.

Before you can view detailed information about constraint violations, you must generate a constraint violation report (see "Creating constraint violation reports" on page 1501).

1. Click on the **Constraint Violations** tab to view the report. The report might show information about DS6000, DS8000, TotalStorage Enterprise Storage Server,SAN Volume Controller, and other subsystems all on the same report. By default, the report displays one subsystem per row. Results are aggregated for the current day.
2. Click to view a detailed report of that subsystem.
3. Click , specify charting options, and then click **OK** to view the report in chart form. Right-click on the chart and select **Customize this Chart** to further customize how the chart is displayed.

Viewing affected volumes and affected hosts:

Learn how to view information on volumes and related hosts affected by constraint violations.

Before you can view affected volumes and hosts information, you must generate a constraint violation report (see "Creating constraint violation reports" on page 1501) and view detailed information on a subsystem (see "Viewing detailed constraint violation information").

1. Click for a constraint violation to display the Affected Volumes Selection page.
2. Under Volumes, select whether you want the report to show all volumes or only the most active volumes associated with the subsystem and component and type in the maximum number of volumes to display in the report.
3. Under Performance Data, select whether you do not want the report to display performance data for the volumes or you do want the report to display historic performance data for volumes.
4. Customize the performance report. By default, all supported metrics are included. Click to move the selected metric up in the list and to move the selected metric down in the list.
   a. In the **Included Columns** list, select a metric and click to remove it from the report.
   b. To change the order in which the columns appear in the report, select a metric and click either or .
5. Click **Generate Report**.

Fabric Manager reports

Fabric reports provide you with both overview and detailed information about your fabrics, switches, and ports. The reports are based on information collected by probes and performance monitors. You can view reports as tables of data or as

graphical charts. You can use pre-defined Fabric system reports or define your own switch performance reports for information by port or by constraint violations.

Before you can generate a report you need to gather information about your fabrics. You can do this through monitoring jobs, and discoveries.

Creating switch performance reports:

After performance monitors are in operation, you can generate switch performance reports. The procedure is similar for each report.

"Creating performance monitors" on page 300 describes how to create performance monitors for fabric switches.

For both reports, use the **Selection** tab to specify switches to report, report columns to include, and filters to apply. When you click **Generate Report**, the **Ports** page or **Constraint Violations** page is created and displays the report.

1. Expand **Fabric Manager → Reporting → Switch Performance**.
2. Click **By Switch**, **By Port** or **By Constraint Violations**. In the content pane, the Selection page opens.
3. Select from the following options to customize the report:

   Selection button
   Click this button to determine the resources upon which you want to report. The Select Resources window displays. The type of resources displayed for selection depends on the type of report you selected. If you clear the box for a resource on the Select Resources window, a pushpin icon displays on the Selection button.

   Filter button
   Click this button to further defines the data that you want to include in the report. Filters enable you to apply general rules to the report. When you click this button, the Edit Filter window displays.

   For example, to select only data associated with switches that have a name ending with "123–IBM", select the "Switch" column, the "LIKE" operator, and type "*123–IBM" (without the quotation marks) in the **Value 1** field.

   If a filter is defined for this report, a pushpin icon displays on the **Filter** button.

   Time range selectors
   All performance reports display historical performance data. Use the time range selectors to specify the exact time range that you want the report to cover. Click one of the following radio buttons to select the type of historical report you want:

   Display latest performance data
   This report contains only the latest available performance data for the selected resources. Note that even the latest data available might still be somewhat old, depending on whether data was recently collected for the resources. Including the time column in the report will help you determine how old the latest available data actually is. Regardless of the age of the data, clicking this button ensures that only a single row (representing a single point in time) for each selected resource is included in the report. If no data has ever been collected for

a particular resource, then that resource is not included in the report. If no data has ever been collected for any resource, then the resulting report will be empty.

Display historic performance data using absolute time
This report contains only data that falls within a specific time range. Select the time range using the **From** and **To** selectors. The specified from and to values are interpreted as inclusive. The resulting report contains all the performance data that is available within the specified time range, for each selected resource. This means that the data can consist of potentially many rows per resource, each with a unique timestamp. If no data has been collected for a particular resource during the specified time range, that resource is not included in the report. If no data has been collected for any resource during the specified time range, then the resulting report will be empty.

To determine which historical performance records fall within the specified time range, the **beginning timestamp** of the collection intervals is used for comparison. If an historical record contains performance data for an interval of time from 1:00 PM to 1:05 PM, that record would be included in the report for a time range of 1:00 PM to 8:00 PM. However, that same record would not be included if a time range of 1:02 PM to 8:02 PM was specified, because it is the beginning of the time period (for example, 1:00 PM in this example) that is used to determine inclusion in the specified time interval.

Display historic performance data using relative time
This report contains only data that falls within a specific time range extending backward from the current date and time. Specify the number of days in the desired time range, using the **days ago until now** field. The report will be exactly the same as a report using absolute time, with the **To** selector set to the current date and time, and the **From** selector set to the current date and time minus the number of **days ago until now**.

Note that the same criteria applies as if determining which historical performance records fall within the specified time range. Only the beginning timestamp of the collection intervals is used for comparison. This means that data for the current hour or the current day is not included in the report when selecting hourly or daily summation data, because data for the current hour or current day is only saved in the database when the hour or day is complete. As an alternative, consider using the first radio button, **Display latest performance data**, to see a report showing the most recent daily summation data.

Summation Level
Use this selector to select which type of performance data records to include in the report. The Tivoli Storage Productivity Center Performance Manager collects and stores three different types of performance data:

By sample
Represents the most detailed data. There is usually one sample record per resource for every **n** minutes, where **n** is the performance monitor interval length that was specified when the monitor was defined. Each sample represents the average

performance of the resource over those **n** minutes. A sample record with a timestamp of **t** will be saved at time **t+n**. In other words, the beginning of each **n** minute time interval is used for display and for time range matching.

Hourly
Represents an hourly average of the performance of the specified resources. There will be one hourly record written every hour, per resource. Each such record will be written at or shortly after each hour mark on the (device) clock. For example, a record with a timestamp of 12:00 PM and interval length of 1 hour will be written at 1:00 PM.

Daily Represents a daily average of the performance of the resources. There will be one daily record written every day, per resource. Each daily record will be written at or shortly after midnight, according to the device clock. For example, a record with a timestamp of April 21, 12:00 AM and interval length of 24 hours will be written on April 22, 12:00 AM.

Note that hourly and daily records are usually retained in the database for longer periods of time, compared to By Sample records. As a result, if a generated report indicates that sample data is no longer available for a particular time range in the past, it is possible that the sample data has already been purged, in which case you can try generating the report using hourly or daily summation levels.

Be aware that when selecting a time range using the **Display historic performance data using relative time** radio button, it is necessary to select the length of the range in days. This can lead to unexpected results when also selecting the daily summation level. Remember that a daily record is not written until the day is complete, which means that at any given time during the day, the current day's data does not yet exist in the database. Therefore, selecting the daily summation level as well as a relative time range of n **days ago until now**, will usually result in n-1 records being returned per resource, assuming that at least n-1 days worth of data has been previously collected for those resources. Specifying a relative time range of 1 **days ago until now** will result in an empty report.

Note: The user interface displays the date and time of data in the Tivoli Storage Productivity Center server time zone, but the daily and hourly aggregations are calculated based on the device time zones. Hourly records are written at the end of each hour according to the devices' clocks and the daily records are written at midnight according the the devices' clocks, but both are displayed according to the Tivoli Storage Productivity Center server's clock. For example, if a subsystem clock is one hour ahead of the Tivoli Storage Productivity Center server clock, the hourly data for 5:00 PM according to the device's clock is displayed as 4:00 PM in the user interface and the daily data for 12:00 AM according to the device's clock is displayed as 11:00 PM of the previous day.

Available Columns, Included Columns
Use these items to specify which columns to include and which columns to exclude from the generated report.

Available Columns
Displays the columns that you can include in the generated

report that are not already included. If a column appears in the **Included Columns** list box, it will not appear in the **Available Columns** list box.

Included Columns
Displays the columns that will appear in the generated report. By default, all columns for a report are listed in this list box and will be displayed in the generated report. You can highlight the columns and use the up and down arrows to rearrange the order of the columns in the report. The order of the items in this list box determines the order in which the columns will appear in the generated report. The columns with grayed-out names are the frozen columns, which cannot be selected, removed, or reordered for this type of report.

If a column appears in the **Available Columns** list box, it will not appear in the **Included Columns** list box.

To exclude a column from a report, click the name of the column in the **Included Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the left-arrow button (<<). The selected columns are removed from the **Included Columns** list box, and will appear in the **Available Columns** list box.

To include a column in a report, click the name of the column in the Available Columns list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the right-arrow button (>>). The selected columns will be removed from the **Available Columns** list box, and will appear in the **Included Columns** list box.

To rearrange the order of columns that will appear in the generated report, click the name of a column in the **Available Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the up-arrow button (^) or the down-arrow button (v) to move the selected columns up or down in the list, relative to the other columns.

4. Click **Generate** to run the report.
   The report is displayed in the **Ports** tab. You can display the information in chart form by clicking .
5. Click **File → Save as** to save the report in the **<user–id> Reports** node.

Switch Performance:

Use the Switch Performance report to generate and view reports that provide information about the free and occupied ports on the switches in your SAN. These reports provide information about the port connections and status, as well as the port and switch IDs, data and packet transfer rates, and other details. This report can be generated only for switches that have had performance monitors run on them.

The following list describes the columns that are on the Switch Performance report. All of the columns in the report are for display only.

Switch
The ID of the switch.

Time The date and time of the data collection.

Interval
The size of the sample interval, in seconds.

Port Send Packet Rate
The rate at which packets are sent through this port.

Port Receive Packet Rate
The rate at which packets are received through this port.

Total Port Packet Rate
The total rate at which packet operations are performed through this port.

Port Send Data Rate
The rate at which data is sent through this port.

Port Receive Data Rate
The rate at which data is received through this port.

Total Port Data Rate
The total rate at which data is processed through this port.

Port Peak Send Data Rate
The peak rate at which data is sent through this switch.

Port Peak Receive Data Rate
The peak rate at which data is received through this switch.

Port Send Packet Size
The size of the data block that is sent out through this switch.

Port Receive Packet Size
The size of the data block that is received in through this switch.

Overall Port Packet Size
The total size of the data block that is transferred by the switch.

Error Frame Rate
The rate at which error frames are generated by the switch.

Dumped Frame Rate
The rate at which frames are dumped by the switch.

Link Failure Rate
The rate at which failed links are generated by the switch.

Fabric Performance Report - By Port:

Click to display detailed information per computer in a **Network** tab. The **Network** tab displays a file summary by computer.

The following list describes the columns that are on the Fabric Performance Report - By Port report, and are for display only.

Port The ID of the port.

WWPN
The World Wide Port Name (WWPN) of the selected port.

Slot The number for the blade that the port is on.

Index The unique number for the port across the whole switch.

Time The date and time of the data collection.

Interval
The size of the sample interval, in seconds.

Port Send Packet Rate
The average number of packets sent per second, by a particular port over a particular time interval. A send operation is a read operation processed, or a write operation initiated by the particular port.

Port Receive Packet Rate
The average number of packets received per second, by a particular port over a particular time interval. A receive operation is a write operation processed, or a read operation initiated by the particular port.

Total Port Packet Rate
The average number of packets sent and received per second, by a particular port over a particular time interval.

Port Send Data Rate
The average number of megabytes sent per second, by a particular port over a particular time interval.

Port Receive Data Rate
The average number of megabytes received per second, by a particular port over a particular time interval.

Total Port Data Rate
The average number of megabytes transferred per second, by a particular port over a particular time interval.

Port Peak Send Data Rate
The peak number of megabytes sent per second, by a particular port over a particular time interval.

Port Peak Receive Data Rate
The peak number of megabytes received per second, by a particular port over a particular time interval.

Port Send Packet Size
The size of the data block that is sent out through this switch.

Port Receive Packet Size
The size of the data block that is received in through this switch.

Overall Port Packet Size
The total size of the data block that is transferred by the switch.

Error Frame Rate
The rate at which error frames are generated by the switch.

Dumped Frame Rate
The rate at which frames are dumped by the switch.

Link Failure Rate
The rate at which failed links are generated by the switch.

Loss of Sync Rate
The average number of times per second that synchronization was lost after the last reset of the device, for a particular component over a particular time interval.

Loss of Signal Rate
The average number of times per second that the signal was lost after the last reset of the device, for a particular component over a particular time interval.

CRC Error Rate
The average number of frames received per second in which the CRC in the frame did not match the CRC computed by the receiver, for a particular component over a particular time interval.

Short Frame Rate
The average number of frames received per second that were shorter than 28 octets (24 header + 4 CRC), not including any start of frame/end of frame bytes, for a particular component over a particular time interval.

Long Frame Rate
The average number of frames received per second that were longer than 2140 octets (24 header + 4 CRC + 2112 data), not including any start of frame/end of frame bytes, for a particular component over a particular time interval.

Encoding Disparity Error Rate
The average number of disparity errors received per second, for a particular component over a particular time interval.

Discarded Class3 Frame Rate
The average number of class-3 frames per second that were discarded by a particular component over a particular time interval.

F-BSY Frame Rate
The average number of F-BSY frames per second that were generated by a particular component over a particular time interval.

F-RJT Frame Rate
The average number of F-RJT frames per second that were generated by a particular component over a particular time interval.

Constraint Violations:

If you have defined switch performance alerts you can also generate the Constraint Violations report.

Click to display detailed information per computer in a **Network** tab. The **Network** tab displays a file summary by computer.

The following list describes the columns that are on the Top Switch Ports Packet Rate Performance report. The columns are for display only and are not editable.

Switch
The ID of the switch.

Total Port Data Rate Threshold
The total rate at which data is processed through this port.

Link Failure Rate Threshold
The rate at which failed links are generated by the switch.

Error Frame Rate Threshold
The rate at which error frames are generated by the switch.

Total Port Packet Rate Threshold
The total rate at which the packet operations are performed through this port.

System reports

Use the information in this section to learn how to generate and view system reports. These reports are available in the **IBM Tivoli Storage Productivity Center** > **Reporting** > **System Reports** and **Data Manager for Databases** > **My Reports** > **System Reports** nodes in the navigation tree. Data for system reports are gathered every time that monitoring jobs are run against your storage resources.

Generating and viewing system reports

Learn how to generate system reports, drill down through and filter their data, and generate charts.

To display a report, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** > **My Reports** > **System Reports** > **<manager>**.
2. Click the name of the report you want to view. The report will appear in the content pane. To regenerate a report, select the report name, right-click, and select **Regenerate Report**.

Many of the rows that are presented in the generated system reports enable you to:

- drill down into a row to get more detailed information
- view report information in a chart format

To drill down for more detailed report information:

- Click on the desired row. The drill down report is displayed and a new tab representing that report is added to the tabs on the report.

To view report information in a chart format:

- Click on the desired row. The chart is displayed and a new tab representing that chart is added to the tabs on the report.

  You can customize the charts that you generate from system reports by using the Chart Customization window. You can access this window by right-clicking on a system report chart and selecting **Customize this chart** from the pop-up menu. See “Customizing charts” on page 1314 for information on how to customize a system report chart.

To filter the data displayed on a report:

Click the **Filter** button, which displays an Edit Filter window. In this window, specify criteria for records (for example, Capacity in Data Reports), operators (for example, =) and values for each "column" (whatever you want to type). For example, in the Disk Space Summary report, you can specify only disks with a capacity greater than 70 GB.

If you change any element of a report definition (for example, the maximum number of rows in the report), you can save the new report definition (and regenerate it later) by clicking **File** → **Save** and then typing a name for the report. The report is saved with the name you specify under the ***user_id's Reports*** node.

System reports - Data

Use these predefined reports to view statistics about your storage resources.

Access Time Summary:

This report provides a summary of the number of files in your environment and when they were last accessed (for example, created or modified) during the last day, the last week, the last month, the last two months, the last three months, the last six months, the last nine months, the last year, and over a year.

The following list describes the columns that are on the Access Time Summary report. All of the fields on the report are for display only and are not editable.

Last Accessed <= 1 day
The number and total size of the files that were accessed within the last day.

Last Accessed 1 day — 1 week
The number and total size of the files that were accessed between 1 day and 1 week ago.

Last Accessed 1 week — 1 month
The number and total size of the files that were accessed between 1 week and 1 month ago.

Last Accessed 1 month — 2 months
The number and total size the files that were accessed between 1 month and 2 months ago.

Last Accessed 2 months — 3 months
The number and total size of the files that were accessed between 2 – 3 months ago.

Last Accessed 3 months — 6 months
The number and total size of the files that were accessed between 3 – 6 months ago.

Last Accessed 6 months — 9 months
The number and total size of the files that were accessed between 6 – 9 months ago.

Last Accessed 9 months — 1 year
The number and total size of the files that were accessed between 9 months and 1 year ago.

Overall File Count
Total number of files across a network.

Overall File Size
Total size of the space consumed by the files on a network.

Average Age
Average age of files on a network measured by days, hours, minutes, seconds.

- To view any of the following charts, click above the report and select the corresponding option from the pop-up menu:
  - pie chart of storage space distribution for a selected row
  - pie chart of file counts for a selected row
  - pie chart of storage space distribution for all the rows on the report
  - pie chart of file counts for all the rows on a report
  - history chart of file counts

- history chart of file count by percent
- history chart of file sizes
- history chart of files size by percent

- To view a pie chart of file access times, click to left of a report row.

- To view detailed file access times for the computers on the network, click to the left of a report row.

Disk Space Summary:

This report shows storage capacity for the whole network. You can drill down through this report to view storage capacity for each of the following levels: per disk, per computer, per cluster, per computer group, and per domain.

The following list describes the columns that are on the Disk Space Summary report. All of the fields on the report are for display only and are not editable.

Disk Space
Total storage capacity of the disks on the computers within a network.

Consumed File System Space
Amount of used storage space on the file systems within a network.

Available File System Space
Amount of unused storage space on the file systems within a network.

Raw Volume Space
Space on host-side logical volumes that is not occupied by file systems.

Unavailable Disk Space
Amount of storage subsystem volume or virtual storage volume space that is dedicated to redundancy. This value is dependent on the storage subsystem SMI-S provider returning a valid RAID value that is used to determine the overhead. If the value is zero for a storage subsystem, the overhead cannot be calculated. This amount does not include storage space information from storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.

Available Disk Space
Space that is assigned to a (monitored) host that is not part of any logical volume.

Unknown Storage Subsystem Volume Space
Volume space of unknown usage.

- To view any of the following charts click at the top of the report and select the corresponding option from the pop-up menu:
  - pie chart of network space distribution
  - pie chart of network row
  - history chart of the free space for a network
  - history chart of the percentage of free space for a network
  - history chart of the storage capacity on a network

- To view a pie chart of disk capacity, click to the left of a report row.

- To view detailed disk capacity for the computers on a network, click the to the left of a report row.

Access File Summary:

This report provides overview information for files in your storage environment.

The Access File Summary report provides overview information for files by directory, directory group, file system, file system group, cluster, computer, computer group, domain, and for the entire network. Through this report you can view historically the number of files for each resource in the report. The historical chart can be generated to show daily, weekly, or monthly history.

The following list describes the columns that are on the Access File Summary report. All of the fields on the report are for display only.

File Size
Total size of the storage space consumed by the files on a network.

File Count
Total number of files on a network.

Directory Count
Total number of directories on a network.

Average File Size
Average storage space consumed by each of the files on a network.

File System Space
Total storage capacity of the files on a network.

- To view any of the following charts click at the top of the report and select the corresponding option from the pop-up menu:
  - History chart of storage space usage for a selected row
  - History chart of file counts for a selected row
- To view a history chart of the storage space consumed by files on a network, click to the left of a report row.
- To view detailed file storage usage for the computers on the network, click to the left of a report row.

Storage Access Times:

This report shows when files were last accessed.

The following list describes the columns that are on the Storage Access Times report. All of the fields on the report are for display only.

Computer
Name of a computer against which the report was run.

Last Accessed <= 1 day
Number of files that were accessed within the last 24 hours and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes).

Last Accessed 1 day - 1 week
Number of files that were accessed between 1 day to 1 week previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes).

Last Accessed 1 week - 1 month
Number of files that were accessed between 1 week to 1 month previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes).

Last Accessed 1 month - 2 months
Number of files that were accessed between 1 month to 2 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes).

Last Accessed 2 months - 3 months
Number of files that were accessed between 2 months to 3 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes).

Last Accessed 3 months - 6 months
Number of files that were accessed between 3 months to 6 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes).

Last Accessed 6 months - 9 months
Number of files that were accessed between 6 months to 9 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes).

Last Accessed 9 months - 1 year
Number of files that were accessed between 9 months to 1 year previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes).

Last Accessed > 1 year
Number of files that were accessed over one year previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes).

Overall File Count
Total of all the counts.

Overall File Size
Total of all the sizes.

Average Age
Average time since the file was last accessed.

To view graphical information about specific computers' access times, click . A new tab displaying a chart for each selected computer is displayed.

Disk Defects:

This report shows any disk defects that are on the monitored computers.

The following list describes the columns that are on the Disk Defects report. All of the fields on the report are for display only and are not editable.

Device
Name of the computer containing a disk defect.

Path Full path to the location of a disk defect.

Primary Defects
Number of defects on the disk when it was new.

Manufacturer
Name of the manufacturer of the disk drive containing the disk defect.

Model Model number of the disk drive containing a disk defect.

In Storage Subsystem
Indicates Yes when a disk is an actual physical disk within the storage subsystem.

RPM Speed that the disk turns (revolutions per minute).

Read Cache
Hardware caching turned on/off for reads.

Write Cache
Hardware caching turned on/off for writes.

Failure Predicted
Disk predicts a failure.

Probe Last Run
Date and time when the last probe was run.

Discovered Time
Date and time when a disk defect was discovered.

Disk Space
Storage capacity of the disk drive containing a defect.

Available Disk Space
Amount of unused storage space on the disk drive containing a defect.

Grown Defects
Number of defects detected by the disk since new.

Recovered Errors
Number of errors (read, write, verify) from which the disk was able to recover.

Unrecovered Errors
Number of errors (read, write, verify) from which the disk was not able to recover.

Bytes Written
Number of bytes written to the disk.

Bytes Read
Number of bytes read from the disk.

Automatic Write Reallocation
Disk attempts to move the data sector on the disk in the event of a write error.

Desired Automatic Write Reallocation
Change automatic write allocation on the next probe.

Automatic Read Reallocation
Disk attempts to move the data sector on the disk in the event of a read error.

Desired Automatic Read Reallocation
Change automatic read allocation on the next probe.

Read Continuous (no error correction)
No error correction (setting is ON=no error correction).

Desired Read Continuous (no error correction)
Change automatic read continuous on the next probe.

Detectable
Indicates Yes when a storage volume is identified as being available to the storage subsystem following a successful probe.

Is Virtual
Indicates Yes when a disk is identified as being available to the server or the storage subsystem.

Mapped to Storage Subsystem
Indicates Yes when a computer disk or a backend storage volume can be mapped to a storage subsystem.

To view more information for a specific computer:

- Click . Detailed information regarding the computer will be available in three tabs:
  - **General tab.** This tab displays general information about the computer, such as computer information, discovered date and time, disk allocation information, and so on.
  - **Latest Probe tab.** This tab displays information about the latest probe, and provides information such as sector size, probe last run, primary and grown defects, and so on.
  - **Probe History tab.** This tab displays probe history for the selected machine, and provides information such as probe times, record type, bytes written, and so on.

Storage Modification Times:

This report shows information about files within the network that were modified with certain time frames.

The Storage Modification Times report provides information about files within the network that were modified:

- within the last 24 hours
- between 24 hours and one week previous
- between one week to one month previous
- between one month to two months previous
- between two months to three months previous
- between three months to six months previous
- between six months to nine months previous
- between nine months to one year previous
- more than one year previous

The following list describes the columns that are on the Storage Modification Times report. All of the fields on the report are for display only.

Last Modified <= 1 day
Number of files that were modified in the last 24 hours and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)

Last Modified 1 day - 1 week
Number of files that were modified between 1 day to 1 week previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)

Last Modified 1 week - 1 month
Number of files that were modified between 1 week to 1 month previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)

Last Modified 1 month - 2 months
Number of files that were modified between 1 month to 2 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)

Last Modified 2 months - 3 months
Number of files that were modified between 2 months to 3 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)

Last Modified 3 months - 6 months
Number of files that were modified between 3 months to 6 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)

Last Modified 6 months - 9 months
Number of files that were modified between 6 months to 9 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)

Last Modified 9 months - 1 year
Number of files that were modified between 9 months to 1 year previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)

Last Modified > 1 year
Number of files that were modified over a year previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)

Overall File Count
Sum of the files

Overall File Size
Sum of the size of the files

Average Age
Average age since files were modified

To view graphical information about storage modification times, click . A new tab displaying a chart for the network is displayed.

Most At Risk Files:

This report provides information on files that were modified the longest time ago and were not backed-up or archived since they were modified.

The following list describes the columns that are on the Most At Risk Files report. All of the fields on the report are for display only and are not editable.

Modification Time
Date and time when an at risk file was last modified

Computer
Name of the computer where an at risk file is stored

Filesystem
Filesystem where an at risk file is stored

Path Full path to the location of an at risk file

Filename
The name of the at risk file

Owner
ID of the user who owns an at risk file

OS User Group
OS User group of the user who owns an at risk file

Logical Size
Logical size of an at risk file

Physical Size
Physical size of an at risk file (measured in kilobytes, megabytes, or gigabytes)

Access Time
Date and time when an at risk file was last accessed

Create Time
Date and time when an at risk file was created

To view more information about a specific file:

- Click . Detailed information regarding the file is displayed in a **File Detail** tab. The **File Detail** tab displays more information for the specific file and provides information such as file ownership, modification and creation time, and any file attributes that were set.

Most Obsolete Files:

This report provides information about files that have not been accessed or modified for the longest period of time.

The following list describes the columns that are on the Most Obsolete Files report. All of the fields on the report are for display only and are not editable.

Access Time
Date and time when an obsolete file was last accessed

Computer
Name of the computer where an obsolete file is stored

Filesystem
Filesystem where an obsolete file is stored

Path Full path to the location of an obsolete file

Filename
File name of an obsolete file

Owner
ID of the user who owns an obsolete file

OS User Group
OS User Group to which the owner of the obsolete file belongs

Logical Size
Logical size of an obsolete file

Physical Size
Physical size of an obsolete file (measured in kilobytes, megabytes, or gigabytes)

Modification Time
Date and time when an obsolete file was last modified

Create Time
Date and time when an obsolete file was created

Viewing more information about a specific file

Click . Detailed information regarding the file is displayed in a **File Detail** tab. The **File Detail** tab displays more information for the specific file and provides information such as file ownership, modification and creation time, and any file attributes that are set.

Oldest Orphaned Files:

This report provides information about files that have the oldest creation date and no longer have owners who are registered as users on the computer or network.

The following list describes the columns that are on the Oldest Orphaned Files report. All of the fields on the report are for display only.

Access Time
Date and time when an orphaned file was last accessed

Computer
Name of the computer where an orphaned file is stored

Filesystem
Filesystem where an orphaned file is stored

Path Full path to the location of an orphaned file

Filename
File name of an orphaned file

Owner
Operating system internal ID of the user who owned the orphaned file. This is the internal ID that the operating system uses to identify the user, and *not* the user ID

OS User Group
OS User group to which the owner of an orphaned file belongs

Logical Size
Logical size of an orphaned file

Physical Size
Physical size of an orphaned file (measured in kilobytes, megabytes, or gigabytes)

Modification Time
Date and time when an orphaned file was last modified

Create Time
Data and time when an orphaned file was created

To view more information for a specific file:

- Click . Detailed information regarding the file is displayed in a **File Detail** tab. The **File Detail** tab displays information about a specific file such as file ownership, modification and creation time, and any file attributes that are set.

User Quota Violations:

This report shows which users have violated a Data Manager storage usage quota.

The following list describes the columns on the User Quota Violations report. All the fields on the report are for display only.

Quota Creator
Creator of a quota

Quota Name
Name of a quota

Computer
Computer that violated a quota

Filesystem
File system that violated a quota

User Name
ID of the user that violated a quota

Space Used
Total amount of space used

Threshold
Threshold defined by a quota

% of Threshold
Percentage of quota; always > 100

Number of Violations
Number of violations by a user

Quota Type
Type of quota: Network Wide, Filesystem, or Computer

Computer Storage Availability:

This report provides information about the availability of computers within your environment.

If the computer is not available (determined by a ping), the storage is likely not available for use. The following list describes the fields that are on the Storage Availability report. All of the fields on the report are for display only.

Computer
Name of the computer

% Available
Percentage of successful pings

Transmitted
Number of pings transmitted

Received
Number of responses from the ping

Damaged
Number of responses with errors

Avg Transit Time
Average time to send and receive a response from a ping

Max Transit Time
Maximum time it took to send and receive response from a ping

Min Transit Time
Minimum time it took to send and receive response from a ping

Report Start Time
Start time of information in this row

Report End Time
End time of information

To view the ping history for a specific computer, click . A new tab displaying ping history information for the selected computer is displayed.

Computer Disk Space:

This report provides storage capacity information about each computer within your environment.

The following list describes the columns that are on the Computer Disk Space report. All of the fields on the report are for display only.

Computer
Name of a computer against which the report was run

Disk Space
Total storage capacity for a computer

Available Disk Space
Amount of unused storage space on a computer (not in file systems seen by this operating system)

Owned Disk Space
Total amount of unique disk storage space for computers in a cluster that is owned by those computers. If Tivoli Storage Productivity Center determines that a disk is configured for a different host, that disk's capacity is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.

- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note:

- If the computer is an IBM SAN File System client, this value does not include the capacity of any volumes that are visible to it that are owned by the SAN File System.
- If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads from the host ID. If this does not agree with the local host ID, the disk is classified as belonging to a different (Solaris) host.

Owned Available Disk Space

Total amount of unique disk storage space for computers in a cluster that is not allocated to any logical volume within the computers and is owned by the computers. If Data Manager determines that a disk is configured for a different computer, that disk's unallocated space is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

If the computer is an IBM SAN File System client, then this value excludes the unallocated space on any storage subsystem volumes visible to it that are owned by the SAN File System.

If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product will examine the Veritas disk-label and read off the host-ID. If this does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

Non-Fibre Channel Attached Disk Space

Total amount of unique disk space on the computers in a cluster that is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Non-Fibre Channel Attached Available Disk Space

Total amount of disk storage space on the computers in a cluster that is not allocated to any logical volume within the computers and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Disk Space
Total amount of unique disk space on the computers in a cluster that is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Storage Resource agent.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Available Disk Space
Total amount of disk storage space on the computers in a cluster that is not allocated to any logical volume within the computers and is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Storage Resource agent.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

OS Type
Operating system running on a computer

Network Address
Network address of a computer

IP Address
IP address of a computer

Time Zone
Time zone in which a computer is running

CPU Architecture
Instruction set architecture on the machine where an agent is running

Available File System Space:

This report shows the total amount of unused storage across a network.

The following list describes the columns on the Total Free Space report. All the fields on the report are for display only.

Available File System Space
Total amount of available storage space available on a network.

Percent Available File System Space
Percentage of total space that is unused on a network.

Consumed File System Space
Amount of used storage space on a network.

File System Space
Total amount (capacity) of storage space on a network.

File Count
Total number of files on a network.

Directory Count
Total number of directories on a network.

Percent Free Inodes
Percentage of free Inodes on a network.

Used Inodes
Number of used Inodes on a network.

Free Inodes
Number of free Inodes on a network.

User Space Usage:

This report provides information about storage statistics related to users within your environment.

The following list describes the columns on the User Space Usage report. All the fields on the report are for display only.

User Name
ID of a user

File Size
Total amount of space used by a user

File Count
Number of files owned or created by a user

Directory Count
Number of directories owned or created by a user

Largest File
Largest file owned by a user

2nd Largest File
Second-largest file owned by a user

Files < 1KB
Number and total space usage of files under 1KB in size

Files 1KB - 10KB
Number and total space usage of files between 1KB and 10KB in size

Files 10KB - 100KB
Number and total space usage of files between 10KB and 100KB in size

Files 100KB - 1MB
Number and total space usage of files between 100KB and 1MB in size

Files 1MB - 10MB
Number and total space usage of files between 1MB and 10MB in size

Files 10MB - 100MB
Number and total space usage of files between 10MB and 100MB in size

Files 100MB - 500MB
Number and total space usage of files between 100MB and 500MB in size

Files > 500MB
Number and total space usage of files over 500MB in size

Wasted Space:

This report provides information about storage statistics on non-OS files not accessed in the first year and orphan files.

The following list describes the columns on the Wasted Space-File Summary by Computer report. All the fields on the report are for display only.

File Size
Total amount of space used by the obsolete and orphan files

File Count
Total number of obsolete and orphan files

Directory Count
Total number of orphan directories

Average File Size
Average size of obsolete and orphan files

File System Space
Amount of storage space available on a storage subsystem.

To view more information for a specific file, click . Detailed information per computer will be displayed in a **Network** tab. The **Network** tab displays file summary by computer.

System reports - Data Manager for Databases

Use these predefined reports to view statistics about your RDBMS storage resources. These reports are available in the **Data Manager for Databases** > **My Reports** > **System Reports** node of the navigation tree.

All DBMS - Database Storage By Computer:

This report shows storage information about the databases in your environment sorted by the computers on which they are stored.

The following list describes the columns that are displayed in the All DBMS Database Storage By Computer report.

Computer
Name of the computer where the databases are located.

Total Size
Amount of space that is consumed by the databases on the computers.

File Capacity
Storage capacity of the data files within the databases.

File Free Space
Amount of free space that is available on the data files within the databases.

DB-TS Count
Number tablespaces that are associated with the databases.

File Count
Number of data files that are associated with the tablespaces in the databases.

Log File Count
Number of log files that are associated with the databases.

All DBMS - User Database Space Usage:

This report shows information about the RDBMS storage space consumed by users across the network.

The following list describes the columns that appear on the All DBMS - User Database Space Usage report:

User Name
Name of a user upon whose storage usage information was collected.

Total Size
Amount of space consumed by the objects that are owned by the user.

Empty Used Space
Amount of empty used space that resides on the objects that are owned by the user.

Table Count
Number of tables that are owned by the user.

Index Count
Number of indexes that are owned by the user.

Other Count
Number of other objects owned by the user.

Largest Object
Size of the largest object owned by the user.

2nd Largest Object
Size of the second largest object owned by the user.

Segment Size < 50KB
Number and total space usage of objects owned by the user that are under 50KB in size.

Segment Size 50KB - 500KB
Number and total space usage of objects owned by the user that are between 50KB and 500KB in size.

Segment Size 500KB - 1Mb
Number and total space usage of objects owned by the user that are between 500KB and 1MB in size.

Segment Size 1Mb - 10Mb
Number and total space usage of objects owned by the user that are between 1MB and 10MB in size.

Segment Size 10Mb - 100Mb
Number and total space usage of objects owned by the user that are between 10MB and 100MB in size.

Segment Size 100Mb - 500Mb
Number and total space usage of objects owned by the user that are between 100MB and 500MB in size.

Segment Size 500Mb - 1GB
Number and total space usage of objects owned by the user that are between 500MB and 1GB in size.

Segment Size > 1Gb
Number and total space usage of objects owned by the user that are over 1GB in size.

All DBMS - User Database Quota Violations:

This report provides information on the users who violated any storage usage thresholds defined in quotas.

The following list describes the columns that are on the All DBMS - User Database Quota Violations report.

Quota Creator
ID of the user who created the quota

Quota Name
Name of the quota

Computer
Name of the computer where the quota violation occurred

RDBMS type
RDBMS to which an instance belongs, for example, Oracle, SQL Server, Sybase, or UDB

Instance
Name/SID of the Oracle Instance, SQL Server server, Sybase server, or UDB Instance where the quota violation occurred

Database-Tablespace
Name of the database where the tablespace resides and name of the tablespace whose quota was violated

User Name
ID of the user who violated the quota

Space Used
Amount of space consumed by the user

Threshold
Threshold of used space defined within the quota

% of Threshold
Percentage of space by which the user violated the quota

Number of Violations
Number of violations for the user

Quota Type
The type of quota: Database-Tablespace, Instance, or Network-wide

Oracle - Database Storage By Computer:

This report shows storage information about the databases in your environment sorted by the computers on which they are stored.

The following list describes the columns that are in the Oracle - Database Storage By Computer report.

Computer
Name of the computer where the databases are located.

Total Size
Amount of space consumed by the databases on the computers.

Data File Capacity
Storage capacity of the data files within the databases.

Data File Free space
Amount of free space available on the data files within the databases.

Tablespace Count
Number tablespaces associated with the databases.

Data File Count
Number of data files associated with the tablespaces in the databases.

Log File Count
Number of log files associated with the databases.

Oracle - Total Database Free Space:

This report provides information on the total free space for data files at a network-wide level, the percentage of free space, the total used space, number of free extents, and number of data files on a network. This report is available for Oracle and Microsoft SQL Server.

Free Space
Amount of free space on all data files in the network.

Percent Free
Percent of free space available on the data files in the network.

Used Space
Amount of used space on the data files in the network.

Total Size
Total size of data files in the network.

Free Extents
Number of free extents on the data files in the network.

Coalesced Extents
Number of coalesced extents on the databases in the network.

Minimum Free Extent Size
Minimum size of a free extent in a database.

Maximum Free Extent Size
Maximum size of a free extent in a database.

Number Data Files
Number of data files on the databases in the network.

Oracle - User Database Space Usage:

This report shows information about the RDBMS storage space consumed by users across the network.

The following list describes the columns that appear on the Oracle - User Database Space Usage report:

User Name
Name of a user upon whose storage usage information was collected.

Total Size
Amount of space consumed by the objects that are owned by the user.

Empty Used Space
Amount of empty used space that resides on the objects that are owned by the user.

Table Count
Number of tables that are owned by the user.

Index Count
Number of indexes that are owned by the user.

Other Count
Number of other objects owned by the user.

Largest Segment
Size of the largest segment owned by the user.

2nd Largest Segment
Size of the second largest segment owned by the user.

Segment Size < 50KB
Number and total space usage of objects owned by the user that are under 50KB in size.

Segment Size 50KB - 500KB
Number and total space usage of objects owned by the user that are between 50KB and 500KB in size.

Segment Size 500KB - 1Mb
Number and total space usage of objects owned by the user that are between 500KB and 1MB in size.

Segment Size 1Mb - 10Mb
Number and total space usage of objects owned by the user that are between 1MB and 10MB in size.

Segment Size 10Mb - 100Mb
Number and total space usage of objects owned by the user that are between 10MB and 100MB in size.

Segment Size 100Mb - 500Mb
Number and total space usage of objects owned by the user that are between 100MB and 500MB in size.

Segment Size 500Mb - 1GB
Number and total space usage of objects owned by the user that are between 500MB and 1GB in size.

Segment Size > 1Gb
Number and total space usage of objects owned by the user that are over 1GB in size.

Oracle - Segments With Wasted Space:

This report provides information on the segments containing allocated space that is currently empty/not being used. Use this information to gather information that can help you discover space that can be reclaimed and allocated to other objects.

Empty Used Space
The amount of empty used space within a segment (table, index, etc.)

Segment Creator
The owner of a segment

Segment Name
The name of a segment

Computer
The name of the computer on which the Instance of a segment Instance resides

Instance
The SID of an Oracle Instance

Database
The name of the database to which the segment belongs

Tablespace
The name of the tablespace to which the segment belongs

Partition
The partition on which the segment is stored

Segment Type
Type of the segment, including:
- TABLE
- TABLE PARTITION
- TABLE SUBPARTITION
- NESTED TABLE
- CLUSTER
- INDEX
- INDEX PARTITION
- INDEX SUBPARTITION
- LOBINDEX
- LOBSEGMENT
- LOB PARTITION
- LOB SUBPARTITION

Parent Type
Subset of the segment type

Parent Creator
The owner of a segment

Parent Name
The name of a segment

Total Size
The amount of space that is allocated to a segment

Number of Extents
The number of extents that is allocated to a segment

Freelist Blocks
The number of blocks that are on the freelist chain

Initial Extent
The size of the first extent that is allocated to a segment

Next Extent
The amount of space that Oracle will retrieve when allocating another extent

Maximum Extents
The maximum number of extents that Oracle can allocate to an object

Percent Increase
The size that Oracle allocates for the next extent

SqlServer - Database Storage By Computer:

This report shows storage information about the databases in your environment sorted by the computers on which they are stored.

The following list describes the columns that are displayed in the SqlServer Database Storage By Computer report.

Computer
Name of the computer on which SQL Server is running.

Total Size
Total amount of storage space allocated to SQL Server on a computer.

Data File Capacity
Amount of allocated storage space for data files that is unused on a computer where SQL Server is running.

Data File Free space
Amount of free space available on the data files associated with SQL Server on a computer.

Log File Capacity
Amount of allocated storage space for log files that is unused on a computer where SQL Server is running.

Log File Free space
Amount of free space available on the log files associated with SQL Server on a computer.

Database Count
Number of databases associated with SQL Server on a computer.

Data File Count
Number of data files associated with SQL Server on a computer.

Log File Count
Number of log files associated with SQL Server on a computer.

SqlServer - Total Database Free Space:

This report provides information on the total free space for data files at a network-wide level, the percentage of free space, the total used space, number of free extents, and number of data files on a network. This report is available for Oracle and Microsoft SQL Server.

Data Free Space
Amount of free space on all data files in the network.

Percent Free
Percent of free space available on the data files in the network.

Data Used Space
Amount of used space on the data files in the network.

Data Size
Total size of data files in the network.

Log Free Space
Amount of free space.

Log Percent Used
Percent of used space.

Log Used Space
Amount of space that is used by the log.

Log Size
Size of the log.

Number Data Files
Number of data files on the databases in the network.

SqlServer - User Database Space Usage:

This report shows information about the RDBMS storage space consumed by users across the network.

The following list describes the columns that appear on the SqlServer - User Database Space Usage report:

User Name
Name of a user upon whose storage usage information was collected.

Total Size
Amount of space consumed by the objects that are owned by the user.

Empty Used Space
Amount of empty used space that resides on the objects that are owned by the user.

Table Count
Number of tables that are owned by the user.

Index Count
Number of indexes that are owned by the user.

Other Count
Number of other objects owned by the user.

Largest Object
Size of the largest object owned by the user.

2nd Largest Object
Size of the second largest object owned by the user.

Object Size < 50KB
Number and total space usage of objects owned by the user that are under 50KB in size.

Object Size 50KB - 500KB
Number and total space usage of objects owned by the user that are between 50KB and 500KB in size.

Object Size 500KB - 1Mb
Number and total space usage of objects owned by the user that are between 500KB and 1MB in size.

Object Size 1Mb - 10Mb
Number and total space usage of objects owned by the user that are between 1MB and 10MB in size.

Object Size 10Mb - 100Mb
Number and total space usage of objects owned by the user that are between 10MB and 100MB in size.

Segment Size 100Mb - 500Mb
Number and total space usage of objects owned by the user that are between 100MB and 500MB in size.

Object Size 500Mb - 1GB
Number and total space usage of objects owned by the user that are between 500MB and 1GB in size.

Object Size > 1Gb
Number and total space usage of objects owned by the user that are over 1GB in size.

Sybase - Database Storage By Computer:

This report shows storage information about the databases in your environment sorted by the computers on which they are stored.

The following list describes the columns on the Sybase - Database Storage By Computer report.

Computer
Name of the computer on which Sybase is running.

Total Device Size
Total amount of storage space that is allocated to Sybase on a computer.

Mirrored Device Size
Total storage capacity of the mirrored device for an instance.

Device Free Space
Amount of free space available on the data files that are associated with Sybase on a computer.

Data Fragment Capacity
Amount of allocated storage space for log files that is unused on a computer where Sybase is running.

Data Fragment Free space
Amount of free space available on the data fragments that are associated with Sybase on a computer.

Log Fragment Capacity
Amount of allocated storage space for log fragments that is unused on a computer where Sybase is running.

Log Fragment Free space
Amount of free space available on the log files that are associated with Sybase on a computer.

Device Count
Number of devices associated with Sybase on a computer.

Database Count
Number of databases associated with Sybase on a computer.

Fragment Count
Number of fragments associated with Sybase on a computer.

Log File Count
Number of log files associated with Sybase on a computer.

Sybase - Total Device Free Space:

This report provides information on the total free space for Sybase devices at a network-wide level, the percentage of free space, the total used space, number of fragments, and number of devices.

Free Space
Amount of free space on all the devices in the network.

Percent Free
Percent of free space available on the devices in the network.

Used Space
Amount of used space on the devices in the network.

Total Size
Total size of devices in the network.

Number Fragments
Number of fragments that are associated with the devices in the network.

Number Devices
Number of devices in the network.

Sybase - Total Database Free Space:

This report provides information on the total free space for data files at a network-wide level, the percentage of free space, the total used space, number of free extents, and number of data files on a network. This report is available for Oracle and Microsoft SQL Server.

Data Free Space
Amount of free space on all data files in the network.

Percent Free
Percent of free space available on the data files in the network.

Data Used Space
Amount of used space on the data files in the network.

Data Size
Total size of data files in the network.

Log Free Space
Amount of free space available to a database log.

Log Percent Used
Percentage of allocated space that a database log is consuming.

Log Used Space
Amount of storage space that is consumed by a database log.

Log Size
Total amount of space that is allocated to a database log.

Number Fragments
Number of fragments that are associated with a database.

Sybase - User Database Space Usage:

This report shows information about the RDBMS storage space consumed by users across the network.

The following list describes the columns that appear on the Sybase- User Database Space Usage report:

User Name
Name of a user upon whose storage usage information was collected.

Total Size
Amount of space consumed by the objects that are owned by the user.

Empty Used Space
Amount of empty used space that resides on the objects that are owned by the user.

Table Count
Number of tables that are owned by the user.

Index Count
Number of indexes that are owned by the user.

Other Count
Number of other objects owned by the user.

Largest Object
Size of the largest object owned by the user.

2nd Largest Object
Size of the second largest object owned by the user.

Object Size < 50KB
Number and total space usage of objects owned by the user that are under 50KB in size.

Object Size 50KB - 500KB
Number and total space usage of objects owned by the user that are between 50KB and 500KB in size.

Object Size 500KB - 1Mb
Number and total space usage of objects owned by the user that are between 500KB and 1MB in size.

Object Size 1Mb - 10Mb
Number and total space usage of objects owned by the user that are between 1MB and 10MB in size.

Object Size 10Mb - 100Mb
Number and total space usage of objects owned by the user that are between 10MB and 100MB in size.

Object Size 100Mb - 500Mb
Number and total space usage of objects owned by the user that are between 100MB and 500MB in size.

Object Size 500Mb - 1GB
Number and total space usage of objects owned by the user that are between 500MB and 1GB in size.

Object Size > 1Gb
Number and total space usage of objects owned by the user that are over 1GB in size.

UDB - Instance Storage By Computer:

This report shows information about the UDB instances in your environment sorted by the computers on which they are stored.

The following list describes the columns that are displayed on the UDB - Instance Storage By Computer report.

Computer
Name of the computer on which a UDB Instance(s) resides

Total Size
Total amount of storage space allocated to UDB on a computer

Container Capacity
Amount of storage space allocated to the containers that reside on a computer

Container Free Space
Amount of free space available on the containers that reside on a computer

Log File Capacity
Amount of storage space allocated to the log files that reside on a computer

Tablespace Count
Number of tablespaces that reside on a computer

Container Count
Number of containers that reside on a computer

Log File Count
Number of log files that reside on a computer

UDB - Total DMS Container Free Space:

This report shows the total free space for the containers that are associated with DMS table spaces on UDB instances within your environment.

The following columns are displayed on the UDB - Total DMS Container Free Space report:

Free Space
Amount of free space available on the DMS containers within a network.

Percent Space
Percentage of free space available on the DMS containers within a network.

Used Space
Amount of storage space that is consumed on the DMS containers within a network.

Total Size
Total amount of space on the DMS containers within a network.

Number Containers
Number of DMS containers within a network.

UDB - User Database Space Usage:

This report shows information about the RDBMS storage space consumed by users across the network.

The following list describes the columns that appear on the UDB - User Database Space Usage report:

User Name
Name of a user upon whose storage usage information was collected.

Total Size
Amount of space consumed by the objects that are owned by the user.

Empty Used Space
Amount of empty used space that resides on the objects that are owned by the user.

Table Count
Number of tables that are owned by the user.

Index Count
Number of indexes that are owned by the user.

Other Count
Number of other objects owned by the user.

Largest Object
Size of the largest object owned by the user.

2nd Largest Object
Size of the second largest object owned by the user.

Object Size < 50KB
Number and total space usage of objects owned by the user that are under 50KB in size.

Object Size 50KB - 500KB
Number and total space usage of objects owned by the user that are between 50KB and 500KB in size.

Object Size 500KB - 1Mb
Number and total space usage of objects owned by the user that are between 500KB and 1MB in size.

Object Size 1Mb - 10Mb
Number and total space usage of objects owned by the user that are between 1MB and 10MB in size.

Object Size 10Mb - 100Mb
Number and total space usage of objects owned by the user that are between 10MB and 100MB in size.

Object Size 100Mb - 500Mb
Number and total space usage of objects owned by the user that are between 100MB and 500MB in size.

Object Size 500Mb - 1GB
Number and total space usage of objects owned by the user that are between 500MB and 1GB in size.

Object Size > 1Gb
Number and total space usage of objects owned by the user that are over 1GB in size.

System Reports - Disk Manager

Use these predefined reports to view performance information about storage subsystems.

The following system reports are available for Disk Manager .

Array Performance:

This report shows performance data for arrays. This report supports IBM System Storage DS8000, DS6000, and IBM TotalStorage Enterprise Storage Server systems only.

Column	Description
Subsystem	ID of the storage system
Array	The ID of the specified array.
Time	Date and time that the data was collected.
Interval	The actual duration of the sample interval (in seconds) for the performance data that was collected.
Disk Utilization Percentage	The approximate utilization percentage of a rank over a specified time interval (the average percent of time that the disks associated with the array were busy).
Sequential I/O Percentage	Percentage of all I/O operations that are performed for an array over a specified time interval that were sequential operations.
Backend Read I/O Rate	Average number of I/O operations per second for read operations.
Backend Write I/O Rate	Average number of I/O operations per second for write operations.
Total Backend I/O Rate	Average number of I/O operations per second for read and write operations.
Backend Read Data Rate	Average number of megabytes (2^20 bytes) that were transferred for read operations.
Backend Write Data Rate	Average number of megabytes (2^20 bytes) that were transferred for write operations.
Total Backend Data Rate	Average number of megabytes (2^20 bytes) that were transferred for read and write operations.
Backend Read Response Time	Average number of milliseconds that it took to respond to each read operation. For SAN Volume Controller models, this time is the external response time of the managed disks (MDisks).
Backend Write Response Time	Average number of milliseconds that it took to respond to each write operation. For SAN Volume Controller models, this time is the external response time of the managed disks (MDisks).
Overall Backend Response Time	Average number of milliseconds that it took to respond to each I/O operation (read and write). For SAN Volume Controller models, this time is the external response time of the managed disks (MDisks).
Write-cache Delay Percentage	Percentage of I/O operations that were delayed because of write-cache space constraints or other conditions. Write-cache Delay Percentage is the ratio of delayed I/O operations to total I/O operations.

Column	Description
Write-cache Delay I/O Rate	Average number of I/O operations per second that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval.

Controller Cache Performance:

This report shows performance data for storage system controller caches. This report supports IBM System Storage DS8000, DS6000, DS5000, DS4000, and IBM TotalStorage Enterprise Storage Server systems only.

Column	Description
Subsystem	ID of the storage system.
Controller	The ID of the controller.
Time	Date and time that the data was collected.
Interval	The actual duration of the sample interval (in seconds) for the performance data that was collected.
Read I/O Rate (normal)	Average number of I/O operations per second for non-sequential read operations for a component over a specified time interval.
Read I/O Rate (sequential)	Average number of I/O operations per second for sequential read operations for a component over a specified time interval.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.
Write I/O Rate (normal)	Average number of I/O operations per second for non-sequential write operations for a component over a specified time interval.
Write I/O Rate (sequential)	Average number of I/O operations per second for sequential write operations for a component over a specified time interval.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.
Total I/O Rate (normal)	Average number of I/O operations per second for non-sequential read and write operations for a component over a specified time interval.
Total I/O Rate (sequential)	Average number of I/O operations per second for sequential read and write operations for a component over a specified time interval.

Column	Description
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential read and write operations for a component over a specified time interval.
Cache Holding Time	Average cache holding time, in seconds, for I/O data in this subsystem controller (cluster). Shorter time periods indicate adverse performance.
Read Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and non-sequential read operations for a component over a specified time interval.
Write Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and non-sequential write operations for a component over a specified time interval.
Total Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and non-sequential read and write operations for a component over a specified time interval.
Read Cache Hits Percentage (normal)	Percentage of cache hits for non-sequential read operations for a component over a specified time interval.
Write Cache Hits Percentage (normal)	Percentage of cache hits for non-sequential write operations for a component over a specified time interval.
Total Cache Hits Percentage (normal)	Percentage of cache hits for non-sequential read and write operations for a component over a specified time interval.
Read Cache Hits Performance (sequential)	The sequential read cache allocated for the controller.
Write Cache Hits Performance (sequential)	The sequential write cache allocated for the controller.
Total Cache Hits Performance (sequential)	The total sequential cache percentage (read and write) allocated for the controller.
Record Mode Read I/O Rate	Average number of I/O operations per second for record mode read operations for a component over a specified time interval.
Record Mode Read Cache Hits Percentage	Percentage of cache hits for record mode read operations for a component over a specified time interval.
Disk to Cache I/O Rate	Average number of I/O operations (track transfers) per second for disk to cache transfers for a component over a specified time interval.
Cache to Disk I/O Rate	Average number of I/O operations (track transfers) per second for cache to disk transfers for a component over a specified time interval.
Write-cache Delay Percentage	Percentage of I/O operations that were delayed because of write-cache space constraints or other conditions. Write-cache Delay Percentage is the ratio of delayed I/O operations to total I/O operations.

Column	Description
Write-cache Delay I/O Rate	Average number of I/O operations per second that were delayed due to write-cache space constraints or other conditions for a component over a specified time interval.

Controller Performance:

This report shows performance data for storage system controllers. This report supports IBM System Storage DS8000, DS6000, DS5000, DS4000, and IBM TotalStorage Enterprise Storage Server systems only.

Column	Description
Subsystem	ID of the storage system.
Controller	The ID of the controller.
Time	Date and time that the data was collected.
Interval	The actual duration of the sample interval (in seconds) for the performance data that was collected.
Read I/O Rate (overall)	Average number of megabytes (2^20 bytes) per second that were transferred for read operations for a component over a specified time interval.
Write I/O Rate (overall)	Average number of megabytes (2^20 bytes) per second that were transferred for write operations for a component over a specified time interval.
Total I/O Rate (overall)	Average number of megabytes (2^20 bytes) per second that were transferred for read and write operations for a component over a specified time interval.
Read Data Rate	Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.
Write Data Rate	Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.
Total Data Rate	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Response Time	Average number of milliseconds that it took to service each read operation for a component over a specified time interval.
Write Response Time	Average number of milliseconds that it took to service each write operation for a component over a specified time interval.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.

I/O Group Performance:

This report shows performance data for I/O groups. This report supports IBM System Storage SAN Volume Controller systems only.

Column	Description
Subsystem	ID of the storage system
I/O Group	The ID of the specified input/output (I/O) group.
Time	Date and time that the data was collected.
Interval	The actual duration of the sample interval (in seconds) for the performance data that was collected.
CPU Utilization Percentage	Average utilization percentage of the CPUs.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read operations for a component over a specified time interval.
Write Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for write operations for a component over a specified time interval.
Total Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read and write operations for a component over a specified time interval.
Read Response Time	Average number of milliseconds that it took to service each read operation for a component over a specified time interval.
Write Response Time	Average number of milliseconds that it took to service each write operation for a component over a specified time interval.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.
Port to Host Send I/O Rate	Average number of exchanges (I/Os) per second sent to host computers by a component over a specified time interval.
Port to Host Receive I/O Rate	Average number of exchanges (I/Os) per second received from host computers by a component over a specified time interval.

Column	Description
Total Port to Host I/O Rate	Average number of exchanges (I/Os) per second transmitted between host computers and a component over a specified time interval.
Port to Disk Send I/O Rate	Average number of exchanges (I/Os) per second sent to storage subsystems by a component over a specified time interval.
Port to Disk Receive I/O Rate	Average number of exchanges (I/Os) per second received from storage subsystems by a component over a specified time interval.
Total Port to Disk I/O Rate	Average number of exchanges (I/Os) per second transmitted between storage subsystems and a component over a specified time interval.
Port to Host Send Data Rate	Average number of megabytes (2^20 bytes) per second sent to host computers by a component over a specified time interval.
Port to Host Receive Data Rate	Average number of megabytes (2^20 bytes) per second received from host computers by a component over a specified time interval.
Total Port to Host Data Rate	Average number of megabytes (2^20 bytes) per second transmitted between host computers and a component over a specified time interval.
Port to Disk Send Data Rate	Average number of megabytes (2^20 bytes) per second sent to storage subsystems by a component over a specified time interval.
Port to Disk Receive Data Rate	Average number of megabytes (2^20 bytes) per second received from storage subsystems by a component over a specified time interval.
Total Port to Disk Data Rate	Average number of megabytes (2^20 bytes) per second transmitted between storage subsystems and a component over a specified time interval.

Managed Disk Group Performance:

This report shows performance data for managed-disk (MDisk) groups. This report supports IBM System Storage SAN Volume Controller systems only.

Column	Description
Subsystem	ID of the storage system
Managed Disk Group	The ID of the specified managed disk.
Time	Date and time that the data was collected.
Interval	The actual duration of the sample interval (in seconds) for the performance data that was collected.
Read Transfer Size	Average number of KB per input/output (I/O) for read operations, for a particular component over an interval.

Column	Description
Write Transfer Size	Average number of KB per I/O for write operations, for a particular component over an interval.
Overall Transfer Size	Average number of KB per I/O for read and write operations, for a particular component over an interval.
Backend Read I/O Rate	Average number of I/O operations per second for read operations.
Backend Write I/O Rate	Average number of I/O operations per second for write operations.
Total Backend I/O Rate	Average number of I/O operations per second for read and write operations.
Backend Read Data Rate	Average number of Mbps that were transferred for read operations.
Backend Write Data Rate	Average number of Mbps that were transferred for write operations.
Total Backend Data Rate	Average number of Mbps that were transferred for read and write operations.

Node Cache Performance report:

This report lists the most recently gathered performance data for IBM System Storage SAN Volume Controller nodes, which is aggregated to the node level from the per-node VDisk sample data.

Column	Description
Subsystem	ID of the storage system
Node	The ID of the specified node.
Time	Date and time that the data was collected.
Interval	The actual duration of the sample interval (in seconds) for the performance data that was collected.
Read Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and nonsequential read operations, for a particular component over an interval.
Write Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and nonsequential write operations, for a particular component over an interval.
Total Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and nonsequential read and write operations, for a particular component over an interval.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.

Column	Description
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Disk to Cache Transfer Rate	Average number of I/O operations (track transfers) per second for disk to cache transfers, for a particular component over an interval.
Cache to Disk Transfer Rate	Average number of I/O operations (track transfers) per second for cache to disk transfers, for a particular component over an interval.
Write-cache Delay Percentage	Percentage of I/O operations that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval. (The ratio of delayed operations to total I/Os.)
Write-cache Overflow Percentage	Percentage of write operations that were delayed because of lack of write-cache space, for a particular component over an interval.
Write-cache Flush-through Percentage	Percentage of write operations that were processed in flush-through write mode, for a particular component over an interval.
Write Cache Write-through Percentage	Percentage of write operations that were processed in write-through write mode, for a particular component over an interval.
Write-cache Delay I/O Rate	Average number of I/O operations per second that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval.
Write-cache Overflow I/O Rate	Average number of tracks per second that were delayed because of lack of write-cache space, for a particular component over an interval.
Write Cache Flush-through I/O Rate	Average number of tracks per second that were processed in flush-through write mode, for a particular component over an interval.
Write Cache Write-through I/O Rate	Average number of tracks per second that were processed in write-through write mode, for a particular component over an interval.
Readahead Percentage of Cache Hits	Percentage of all read cache hits, which occurred on prestaged data.
Dirty Write Percentage of Cache Hits	Percentage of all write cache hits, which occurred on already dirty data in the cache.

Port Performance:

This report shows performance data for ports. This report supports all storage systems.

Column	Description
Subsystem	ID of the storage system
Port	The ID of the specified port.
WWPN	The World Wide Port Name (WWPN) of the selected port.
Time	Date and time that the data was collected.
Interval	The actual duration of the sample interval (in seconds) for the performance data that was collected.
Port Send Bandwidth Percentage	The approximate bandwidth utilization percentage for send operations by a port based on its current negotiated speed.
Port Receive Bandwidth Percentage	The approximate bandwidth utilization percentage for receive operations by this port, based on its current negotiated speed.
Overall Port Bandwidth Percentage	The approximate bandwidth utilization percentage for send and receive operations by this port.
Port Send Utilization Percentage	The average amount of time that the port was busy sending data over a time interval.
Port Receive Utilization Percentage	The average amount of time that the port was busy receiving data over a time interval.
Overall Port Utilization Percentage	The average amount of time that the port was busy sending or receiving data over a time interval.
Port Send I/O Rate	The average number of input/output (I/O) operations per second for send operations for a port over a time interval.
Port Receive I/O Rate	Average number of I/O operations per second of receive operations for a port over a time interval.
Total Port I/O Rate	Average number of I/O operations per second for send and receive operations for a port over a time interval.
Port Send Data Rate	The average number of megabytes (2^20 bytes) per second that were transferred for send (read) operations for a port over a time interval.
Port Receive Data Rate	The average number of megabytes (2^20 bytes) per second that were transferred for receive (write) operations for a port over a time interval.
Total Port Data Rate	The average number of megabytes (2^20 bytes) per second that were transferred for send and receive operations for a port over a time interval.
Port to Host Send I/O Rate	The average number of exchanges (I/Os) per second sent to host computers by a component over a time interval.
Port to Host Receive I/O Rate	The average number of exchanges (I/Os) per second received from host computers by a component over a time interval.

Column	Description
Total Port to Host I/O Rate	The average number of exchanges (I/Os) per second transmitted between host computers and a component over a time interval.
Port to Disk Send I/O Rate	The average number of exchanges (I/Os) per second sent to storage subsystems by a component over a time interval.
Port to Disk Receive I/O Rate	The average number of exchanges (I/Os) per second received from storage subsystems by a component over a time interval.
Total Port to Disk I/O Rate	The average number of exchanges (I/Os) per second transmitted between storage subsystems and a component over a time interval.
Port to Host Send Data Rate	The average number of megabytes (2^20 bytes) per second sent to host computers by a component over a time interval.
Port to Host Receive Data Rate	The average number of megabytes (2^20 bytes) per second received from host computers by a component over a time interval.
Total Port to Host Data Rate	The average number of megabytes (2^20 bytes) per second transmitted between host computers and a component over a time interval.
Port to Disk Send Data Rate	Average number of megabytes (2^20 bytes) per second sent to storage subsystems by a component over a time interval.
Port to Disk Receive Data Rate	Average number of megabytes (2^20 bytes) per second received from storage subsystems by a component over a time interval.
Total Port Data Rate	Average number of megabytes (2^20 bytes) per second transmitted between storage subsystems and a component over a time interval.
Port Send Response Time	The average number of milliseconds that it took to service each send (read) operation for a port over a time interval.
Port Receive Response Time	The average number of milliseconds that it took to service each receive (write) operation for a port over a time interval.
Total Port Response Time	The average number of milliseconds that it took to service each operation (send and receive) for a port over a time interval.
Port Send Transfer Size	The average number of KB sent per I/O by a port over a time interval.
Port Receive Transfer Size	The average number of KB received per I/O by a port over a time interval.
Total Port Transfer Size	The average number of KB transferred per I/O by a port over a time interval.

Subsystem Performance:

This report shows performance data for storage systems. This report supports all storage systems.

Column	Description
Subsystem	ID of the storage system.
Time	Date and time that the data was collected.
Interval	The actual duration of the sample interval (in seconds) for the performance data that was collected.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read operations for a component over a specified time interval.
Write Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for write operations for a component over a specified time interval.
Total Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read and write operations for a component over a specified time interval.
Read Response Time	Average number of milliseconds that it took to service each read operation for a component over a specified time interval.
Write Response Time	Average number of milliseconds that it took to service each write operation for a component over a specified time interval.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.

Top Active Volumes Cache Hit Performance:

This report summarizes the primary metrics for active volumes, prioritized by cache hit ratio. Only volumes with an input/output (I/O) rate that is greater than 0 will be included in this report. If all volumes are inactive, this report will be empty. This report supports all storage systems.

Column	Description
Subsystem	ID of the storage system.
Volume	The ID of the volume.
Time	Date and time that the data was collected.
Interval	The actual duration of the sample interval (in seconds) for the performance data that was collected.
Read Cache Hits Percentage (overall)	Overall read cache percentage for both sequential and nonsequential read operations.
Write Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and nonsequential write operations.
Total Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and nonsequential read and write operations.
Readahead Percentage of Cache Hits	Percentage of all read cache hits which occurred on prestaged data.
Dirty Write Percentage of Cache Hits	Percentage of all write cache hits which occurred on already dirty data in the cache.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.
Total I/O Rate (overall	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read operations for a component over a specified time interval.
Write Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for write operations for a component over a specified time interval.
Total Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read and write operations for a component over a specified time interval.
Read Response Time	Average number of milliseconds that it took to service each read operation for a component over a specified time interval.
Write Response Time	Average number of milliseconds that it took to service each write operation for a component over a specified time interval.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.

Top Volumes Data Rate Performance:

This report summarizes the primary performance metrics for volumes, prioritized by data rate. This report supports all storage systems.

Column	Description
Subsystem	ID of the storage system
Volume	The ID of the specified volume.
Time	Date and time that the data was collected.
Interval	The actual duration of the sample interval (in seconds) for the performance data that was collected.
Read Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read operations for a component over a specified time interval.
Write Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for write operations for a component over a specified time interval.
Total Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read and write operations for a component over a specified time interval.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Response Time	Average number of milliseconds that it took to service each read operation for a component over a specified time interval.
Write Response Time	Average number of milliseconds that it took to service each write operation for a component over a specified time interval.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.

Top Volumes Disk Performance:

This report summarizes the primary performance metrics for volumes, prioritized by disk rate. This report supports all storage systems.

Column	Description
Subsystem	ID of the storage system.

Column	Description
Volume	The ID of the volume.
Time	Date and time that the data was collected.
Interval	The actual duration of the sample interval (in seconds) for the performance data that was collected.
Disk to Cache Transfer Rate	Average number of I/O operations (track transfers) per second for disk to cache transfers, for a particular component over an interval.
Cache to Disk Transfer Rate	Average number of I/O operations (track transfers) per second for cache to disk transfers, for a particular component over an interval.
Write-cache Delay Percentage	Percentage of I/O operations that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval. (The ratio of delayed operations to total I/Os.)
Write-cache Overflow Percentage	Percentage of write operations that were delayed because of lack of write-cache space, for a particular component over an interval.
Write-cache Flush-through Percentage	Percentage of write operations that were processed in flush-through write mode, for a particular component over an interval.
Write Cache Write-through Percentage	Percentage of write operations that were processed in write-through write mode, for a particular component over an interval.
Write-cache Delay I/O Rate	Average number of I/O operations per second that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval.
Write-cache Overflow I/O Rate	Average number of tracks per second that were delayed because of lack of write-cache space, for a particular component over an interval.
Write Cache Flush-through I/O Rate	Average number of tracks per second that were processed in flush-through write mode, for a particular component over an interval.
Write Cache Write-through I/O Rate	Average number of tracks per second that were processed in write-through write mode, for a particular component over an interval.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.

Column	Description
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read operations for a component over a specified time interval.
Write Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for write operations for a component over a specified time interval.
Total Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read and write operations for a component over a specified time interval.
Read Response Time	Average number of milliseconds that it took to service each read operation for a component over a specified time interval.
Write Response Time	Average number of milliseconds that it took to service each write operation for a component over a specified time interval.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.

Top Volumes I/O Rate Performance:

This report summarizes the primary performance metrics for volumes, prioritized by I/O rate. This report supports all storage systems.

Column	Description
Subsystem	ID of the storage system.
Volume	The ID of the specified volume.
Time	Date and time that the data was collected.
Interval	The actual duration of the sample interval (in seconds) for the performance data that was collected.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.

Column	Description
Read Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read operations for a component over a specified time interval.
Write Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for write operations for a component over a specified time interval.
Total Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read and write operations for a component over a specified time interval.
Read Response Time	Average number of milliseconds that it took to service each read operation for a component over a specified time interval.
Write Response Time	Average number of milliseconds that it took to service each write operation for a component over a specified time interval.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.

Top Volumes Response Performance:

This report summarizes the primary performance metrics for volumes, prioritized by response rate. This report supports all storage systems.

Column	Description
Subsystem	ID of the storage system.
Volume	The ID of the volume.
Time	Date and time that the data was collected.
Interval	The actual duration of the sample interval (in seconds) for the performance data that was collected.
Read Response Time	Average number of milliseconds that it took to service each read operation for a component over a specified time interval.
Write Response Time	Average number of milliseconds that it took to service each write operation for a component over a specified time interval.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.
Peak Read Response Time	The peak (worst) response time for disk reads for a component over a time interval.
Peak Write Response Time	The peak (worst) response time for disk writes for a component over a time interval.

Column	Description
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read operations for a component over a specified time interval.
Write Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for write operations for a component over a specified time interval.
Total Data Rate	Average number of megabytes (2^20 bytes) per second that were transferred for read and write operations for a component over a specified time interval.

System reports - Fabric Manager

Use these predefined reports to view performance information about fabrics.

Port Connections:

Use the Port Connections report to view information about the ports on the switches in your storage area network (SAN). This report includes information about the port connections and status, as well as the port and switch IDs.

Fabric ID
The ID of the fabric that the switch is in.

Switch
The label value of the switch.

Switch Port
The ID of the port (*domainID,port_number*).

Slot The number for the blade that the port is on.

Index The unique number for the port across the whole switch.

Switch Port State
The state of the switch port: Normal, Warning, or Critical.

Connection State
The state of the connection between the two ports: Normal, Warning, or Critical.

Connected Port
The port that the switch port is connected to. This value is the WWN of the device connected to that port.

Connected Port State
The state of the connected port: Normal, Warning, or Critical.

Device
The device that the connected port is on.

Device State
The state of the device that the connected port is on: Normal, Warning, or Critical.

SAN Assets (ALL):

Use the SAN Assets (All) report to generate and view reports that provide information about all the assets discovered by IBM Tivoli Storage Productivity Center.

The following list describes the columns that are on the SAN Assets (ALL) report. All of the columns in the report are for display only.

Fabric ID
The ID of the fabric that the asset is in.

Type The type of discovered asset.

Label The label of the asset.

Identifier
The unique identifier of the storage asset. For example, a switch WWN.

Status The state of the device, either Normal, Warning, or Critical.

SAN Assets (Connected Devices):

Use the SAN Assets (Connected Devices) report to generate and view reports that provide information about all the connected devices discovered by IBM Tivoli Storage Productivity Center.

The following table lists the columns that are on the SAN Assets (Connected Devices) report. All of the columns in the report are for display only.

Fabric ID
The ID of the fabric that the asset is in.

Type The type of discovered asset.

Label The label of the asset.

Identifier
The unique identifier of the asset. For example, the device worldwide name.

Status The status of the asset: Normal, Warning, or Critical.

SAN Assets (Switches):

Use the SAN Assets (Switches) report to generate and view reports that provide information about switches in your SAN. These reports provide information about the switch ID, type, manufacturer, WWN, location, and other details.

The following list describes the columns that are on the SAN Assets (Switches) report. All of the columns in the report are for display only.

Fabric ID
The ID of the fabric that the switch is in.

Type The type is always "Switch".

Switch Type
The type of discovered switch. For example, a physical switch.

Label The label of the switch.

Status The status of the switch. The possible states are **Normal** and **Missing**.

IP Address
The IP address of the switch.

Vendor ID
The manufacturer of the switch.

Model The model name and number of the switch.

Version
The version information of the switch.

Serial Number
The serial number of the switch.

Object ID
The object ID of the switch.

Management ID
The management ID of the switch.

Management Address
The URL of the switch element manager.

Domain
The ID of the domain that the switch is in.

Contact
The contact information for the switch.

WWN The world-wide name (WWN) of the switch.

Location
The physical location of the switch.

Physical Switch WWN
The WWN for the parent or physical switch of a virtual switch. If this is not a virtual switch then this column is empty.

Switch Performance:

Use the Switch Performance report to generate and view reports that provide information about the free and occupied ports on the switches in your SAN. These reports provide information about the port connections and status, as well as the port and switch IDs, data and packet transfer rates, and other details. This report can be generated only for switches that have had performance monitors run on them.

The following list describes the columns that are on the Switch Performance report. All of the columns in the report are for display only.

Switch
The ID of the switch.

Time The date and time of the data collection.

Interval
The size of the sample interval, in seconds.

Port Send Packet Rate
The rate at which packets are sent through this port.

Port Receive Packet Rate
The rate at which packets are received through this port.

Total Port Packet Rate
The total rate at which packet operations are performed through this port.

Port Send Data Rate
The rate at which data is sent through this port.

Port Receive Data Rate
The rate at which data is received through this port.

Total Port Data Rate
The total rate at which data is processed through this port.

Port Peak Send Data Rate
The peak rate at which data is sent through this switch.

Port Peak Receive Data Rate
The peak rate at which data is received through this switch.

Port Send Packet Size
The size of the data block that is sent out through this switch.

Port Receive Packet Size
The size of the data block that is received in through this switch.

Overall Port Packet Size
The total size of the data block that is transferred by the switch.

Error Frame Rate
The rate at which error frames are generated by the switch.

Dumped Frame Rate
The rate at which frames are dumped by the switch.

Link Failure Rate
The rate at which failed links are generated by the switch.

Switch Port Errors:

Use the Switch Port Errors report to generate and view reports that provide information about the errors that are being generated by ports on the switches in your SAN. These reports provide information about the error frames, dumped frames, link failures, port and switch IDs, and other details. This report can be generated only for switches that have had performance monitors run on them.

To view more information for a specific file, click . Detailed information per computer is displayed in a **Network** tab. The **Network** tab displays file summary by computer.

The following list describes the fields on the Switch Port Errors report. All the fields on the report are for display only.

Switch
The ID of the switch.

Port The ID of the port.

WWPN
The World Wide Port Name (WWPN) of the selected port.

Slot The number for the blade that the port is on.

Index The unique number for the port across the whole switch.

Time The date and time of the data collection.

Interval
The size of the sample interval, in seconds.

Port Send Packet Size
The size of the data block that is sent out through this switch.

Port Receive Packet Size
The size of the data block that is received in through this switch.

Overall Port Packet Size
The total size of the data block that is transferred by the switch.

Error Frame Rate
The rate at which error frames are generated by the switch.

Dumped Frame Rate
The rate at which frames are dumped by the switch.

Link Failure Rate
The rate at which failed links are generated by the switch.

Loss of Sync Rate
The average number of times per second that synchronization was lost after the last reset of the device, for a particular component over a particular time interval.

Loss of Signal Rate
The average number of times per second that the signal was lost after the last reset of the device, for a particular component over a particular time interval.

CRC Error Rate
The average number of frames received per second in which the CRC in the frame did not match the CRC computed by the receiver, for a particular component over a particular time interval.

Short Frame Rate
The average number of frames received per second that were shorter than 28 octets (24 header + 4 CRC), not including any start of frame/end of frame bytes, for a particular component over a particular time interval.

Long Frame Rate
The average number of frames received per second that were longer than 2140 octets (24 header + 4 CRC + 2112 data), not including any start of frame/end of frame bytes, for a particular component over a particular time interval.

Encoding Disparity Error Rate
The average number of disparity errors received per second, for a particular component over a particular time interval.

The following list describes columns that are not part of the Switch Port Errors system report. However, you can include them, and save the resulting report in the *<user-id> Reports* node.

Port Send Packet Rate
The rate at which packets are sent through this port.

Port Receive Packet Rate
The rate at which packets are received through this port.

Total Port Packet Rate
The total rate at which packet operations are performed through this port.

Port Send Data Rate
The rate at which data is sent through this port.

Port Receive Data Rate
The rate at which data is received through this port.

Total Port Data Rate
The total rate at which data is processed through this port.

Port Peak Send Data Rate
The peak rate at which data is sent through this switch.

Port Peak Receive Data Rate
The peak rate at which data is received through this switch.

Top Switch Ports Data Rate Performance:

Use the Top Switch Port Data Rate Performance report to generate and view reports that provide information about the switches in your SAN that have the top data rates. This report can be generated only for switches that have had performance monitors run on them.

This report only displays rows for the 25 highest data rates from the switch ports monitored. It displays those rows in descending order by Total Port Data Rate. You can adjust the number of rows shown on the **Selection** tab.

To view more information for a specific file, click . Detailed information per computer is displayed in a **Network** tab. The **Network** tab displays the file summary by computer.

The following list describes the fields that are on the Top Switch Ports Data Rate Performance report. All of the fields on the report are for display only.

Switch
The ID of the switch.

Port The ID of the port.

Time The date and time of the data collection.

Interval
The size of the sample interval, in seconds.

Total Port Data Rate
The total rate at which data is processed through this port.

Total Port Packet Rate
The total rate at which the packet operations are performed through this port.

The following list describes columns that are not part of the Top Switch Ports Data Rate Performance report. However, you can include them, and save the resulting report in the ***<user> Reports*** node.

Port Send Packet Rate
The rate at which packets are sent through this port.

Port Receive Packet Rate
The rate at which packets are received through this port.

Port Send Data Rate
The rate at which data is sent through this port.

Port Receive Data Rate
The rate at which data is received through this port.

Port Peak Send Data Rate
The peak rate at which data is sent through this switch.

Port Peak Receive Data Rate
The peak rate at which data is received through this switch.

Port Send Packet Size
The size of the data block that is sent out through this switch.

Port Receive Packet Size
The size of the data block that is received in through this switch.

Overall Port Packet Size
The total size of the data block that is transferred by the switch

Error Frame Rate
The rate at which error frames are generated by the switch.

Dumped Frame Rate
The rate at which frames are dumped by the switch.

Link Failure Rate
The rate at which failed links are generated by the switch.

WWPN
The World Wide Name identifier for the FC port.

Slot The number for the blade that the port is on.

Index The unique number for the port across the whole switch.

Top Switch Ports Packet Rate Performance:

Use the Top Switch Ports Packet Rate Performance report to generate and view reports that provide information about the switches in your SAN that have the top packet rates. This report can only be generated for switches that have had performance monitors run on them.

This report only displays rows for the 25 highest packet rates from the switch ports monitored. It displays those rows in descending order by Total Port Packet Rate. You can adjust the number of rows shown on the Selection tab.

To view more information for a specific file:

- Click . Detailed information per computer will be displayed in a **Network** tab. The **Network** tab displays file summary by computer.

This section describes the fields on the Top Switch Ports Packet Rate Performance report. All the fields on the report are for display only.

Switch
The ID of the switch.

Port The ID of the port.

Time The date and time of the data collection.

Interval
The size of the sample interval, in seconds.

Total Port Packet Rate
The total rate at which the packet operations are performed through this port.

Total Port Data Rate
The total rate at which data is processed through this port.

The following list describes the columns that are not part of the Top Switch Ports Packet Rate Performance report. However, you can include them, and save the resulting report in the ***<user> Reports*** node.

Port Send Packet Rate
The rate at which packets are sent through this port.

Port Receive Packet Rate
The rate at which packets are received through this port.

Port Send Data Rate
The rate at which data is sent through this port.

Port Receive Data Rate
The rate at which data is received through this port.

Port Peak Send Data Rate
The peak rate at which data is sent through this switch.

Port Peak Receive Data Rate
The peak rate at which data is received through this switch.

Port Send Packet Size
The size of the data block that is sent out through this switch.

Port Receive Packet Size
The size of the data block that is received in through this switch.

Overall Port Packet Size
The total size of the data block that is transferred by the switch.

Error Frame Rate
The rate at which error frames are generated by the switch.

Dumped Frame Rate
The rate at which frames are dumped by the switch.

Link Failure Rate
The rate at which failed links are generated by the switch.

WWPN
The World Wide Name identifier for the FC port.

Slot The number for the blade that the port is on.

Index The unique number for the port across the whole switch.

Batch reports

Use the information in this section to learn how to generate and view batch reports. These reports are available in the IBM Tivoli Storage Productivity Center > **Reporting > My Reports** > **Batch Reports** node in the navigation tree.

A batch report represents any Tivoli Storage Productivity Center report that you want to run on a regular basis and save its data to a file. Schedule batch reports to run immediately, run once at a specified date and time, or run repeatedly at a specified date, time, cycle, and day. You can also select the time zone to use when running the report.

You can view the information in a generated batch report file directly or import it into an application of your choice. Batch reports provide a convenient and powerful way for you to save report definitions and schedule when to run those reports.

Creating batch reports

Create a batch report.

To create a batch report follow this procedure:

1. Expand **IBM Tivoli Storage Productivity Center → Reporting → My Reports**.
2. Right-click the **Batch Reports** node and select **Create Batch Report**.

   Note: If the batch report you want to create is similar to an existing one, click the name of the existing report, edit it and save it using a different name.
3. Click the **Report** tab and select the type of report you want to create and move it from **Available** to **Current Selections** by clicking . See "Selecting batch report options" on page 1563 for details of the tab contents.
4. Click the **Selection** tab and refine the criteria for selecting and displaying report data. Some options might not be available for all reports. Use the **Selection** page to:
   - Select a profile for controlling which data is collected. This option is not available for all reports. See "Use profiles" on page 315 for information about profiles.
   - Click **Selection** to specify the resources (for example, computers) that will be displayed in the report. This option is not available for all reports. For details about selecting resources, see "Selecting resources to display on a report" on page 1324.
   - Click **Filters** to apply filters the data that you want to display (for example, show only the machines that contain over 80% used space). For details about applying filters, see "Filtering the resources that appear on a report" on page 1326.
   - Select the columns you want to include in the report and move them from **Available Columns** to **Included Columns** by clicking .

   Note:
   - The values you set in **Available Columns** and **Included Columns** are not used for the following types of batch reports: History CSV File, PDF Chart, and HTML Chart.
   - The value you set in the **Summation Level** field is not used for the following types of batch reports: History CSV File, PDF Chart, and HTML Chart.

5. Click the **Options** tab of the Batch Report job window, specify the following:
   - The machine on which the report file will be saved
   - The report format (for example, HTML)
   - Whether to run a script when the report process is complete
   - The format for the name of the batch report

   See "Selecting batch report options" for details of the tab contents.
6. In the **When to Run** tab of the Batch Report job window, specify when and how often to run the batch report. See "Scheduling a batch report" on page 1567 for details of the tab contents.
7. In the **Alert** tab of the Batch Report job window, define the conditions that will trigger an alert and the method of notification (for example, e-mail). See "Defining alerts for batch reports" on page 1568 for details of the tab contents.
8. Click **File** → **Save** and type a name for the batch report. The batch report will be saved using your user ID as a prefix. For example, if you are logged onto Tivoli Storage Productivity Center as "Smith," the name of the batch report will be Smith.*report_name*.

Selecting batch report options:

Define the agent computer that you want to run the batch report on and specify the format of the file to be saved.

Use the **Options** page to define where the batch report is run and specify information about the output file.

1. Use the following fields to specify the location where a batch report is saved.

   TPC Server
   Save the batch report on the computer where the Tivoli Storage Productivity Center server is located. Use the **Report Location** field to specify the directory where the batch report file is saved.

   Agent Save the batch report on the computer where an agent is located. Use the **Report Location** field to specify the directory where the batch report file is saved.

   You can save batch reports to computers on which Storage Resource agents that are v4.2 or higher are deployed.

   The **Select Agent** field becomes active if you select to save a batch report to an agent.

   Select Agent
   Select the name of the computer on which an agent is deployed. This menu only displays computers on which Storage Resource agents that are v4.2 or higher are deployed:

   To save a batch report to a computer on which v4.1 of a Storage Resource agent is deployed, you must upgrade that agent to v4.2 or higher.

   Note:
   - The type of report that is selected for the batch report determines the agents that are available in this list. Agents that do not support the selected report are not displayed. For example, agents for versions of the product before v4.1 are not available for Data Source reports.

- If the agent you select from this menu is removed from the computer before the batch report is run, an error occurs and the batch report is not generated. If this error occurs, edit the definition of the batch report to specify an existing agent and then run the report again.
- There are differences in the batch reports depending on whether a Storage Resource agent is selected. If you are using a Storage Resource agent to collect data, these additional columns are available in the batch reports using the Selection tab:
  - Is Encryptable
  - Is Encrypted

Report Location
Specify the directory where you want to save the batch report. This field is only available for batch reports that are sent to the following locations:

- TPC Server
- Storage Resource agents that are v4.2 or higher.

For Storage Resource agent batch reports, when you select a Storage Resource agent from the **Select Agent** drop-down list, you see the location for the batch reports.

Windows
SRA_location\agent\log*computername*\reports

UNIX or Linux
SRA_location/agent/log/*computername*/reports

Note:

You can modify the location of the batch reports for Storage Resource agents and the TPC server only.

You also have the option of saving the report on the Tivoli Storage Productivity Center server. The default path for the Tivoli Storage Productivity Center server is:

Windows
C:\Program Files\IBM\TPC\data\log\reports

UNIX or Linux
/opt/IBM/TPC/data/log/reports

2. In the **Report Type Specification** section, select the type of batch report you want to generate. You are limited to one type of output per batch report job and can select from the following output formats:

CSV File
Create a CSV (comma delimited) file of the batch report. Select **Include Headers** to include column headers in the batch report output file. Select **Include Totals** to include the Totals column in the report output file.

Formatted File
Create a formatted text file of the batch report.

HTML File
Create an HTML file of the batch report.

History CSV File
For reports with historical trending, you can export the historical data

to a CSV file. The CSV file has the following format: <*object name, date, amount, date, amount*> and so on, where <*date, amount*> corresponds to each data point on the history chart.

PDF Chart

Create a chart for the batch report in PDF format. After you select this option, you can determine the type of chart you want to create. The list of available chart types vary by the report you selected on the **Report** tab:

- Chart space usage
- History chart: File Count
- History chart: Free Space
- History chart: Free Space %
- History chart: Used Space
- History chart: Used Space %
- History chart: Capacity

Note:

- You must use Adobe Acrobat Reader to view the resulting PDF. You can download Adobe Acrobat Reader from: http://get.adobe.com/reader/.
- This option is available for reports that support charting only.
- For agents that run on Linux and UNIX computers, you must perform additional configuration steps to enable this option. For information about how to configure the agent system, go to the Tivoli Storage Productivity Center Information Center. Search for the topic "Configuring IBM Tivoli Storage Productivity Center for batch reports on UNIX or Linux."
- The batch chart uses any customization settings that are defined for the corresponding "online" chart. "Online" charts are those charts generated using the Reporting facility.

  For example, if you choose to generate a **TPC-wide Storage Space > File System Space > By Filesystem > PDF Chart > Chart space usage** batch report chart, any customization settings defined in the Reporting facility for **Data Manager > Reporting > TPC-wide Storage Space > File System Space > By Filesystem > Chart space usage for selected** is applied to the batch report chart.

  Any batch chart with an associated online report that provides a choice between "selected" and "all" for its generated charts use the customization settings defined for the "selected" chart.
- If you select **PDF Chart > Pie Chart** from the Chart Customization window and one of the elements represented in the pie chart is much smaller than the other elements, the pie chart displays a graphical error. For example, if the pie chart for a batch report shows disk space, and the disk space used is 4 KB and the free space is 10 GB, a graphical error is displayed.

HTML Chart

Select this option if you want to create a chart for the batch report in HTML format. When HTML Chart is selected, the drop-down list to the right becomes active. Use this drop-down list to select the type of chart you want to create. The list of available chart types vary by the report you selected on the **Report** tab:

- Chart space usage
- History chart: File Count
- History chart: Free Space
- History chart: Free Space %
- History chart: Used Space
- History chart: Used Space %
- History chart: Capacity

Note:

- If you generate charts for reports and these charts contain long path names or a large number of segments, the text on the charts might overlap and make them difficult to read.
- For agents that run on Linux and UNIX computers, you must perform additional configuration steps to enable this option. For information about how to configure the agent system, go to the Tivoli Storage Productivity Center Information Center. Search for the topic "Configuring IBM Tivoli Storage Productivity Center for batch reports on UNIX or Linux."
- The batch chart you generate uses any customization settings that are defined for the corresponding "online" chart. "Online" charts are those generated using the Reporting facility.

  For example, if you choose to generate a **TPC-wide Storage Space > File System Space > By Filesystem > PDF Chart > Chart space usage** batch report chart, any customization settings defined in the Reporting facility for **Data Manager > Reporting > TPC-wide Storage Space > File System Space > By Filesystem > Chart space usage for selected** is applied to the batch report chart.

  Any batch chart with an associated online report that provides a choice between "selected" and "all" for its generated charts use the customization settings defined for the "selected" chart.

Customize this chart

Click this button to use the Chart Customization window to customize the chart that is generated for the batch report. You will be able to select a chart type and chart orientation. If there is a customized chart for the corresponding online report, the customized chart is used for this batch report. However, you can override that customized chart by using this function.

The button is active only when **PDF Chart** or **HTML Chart** is selected.

3. In the **Classic Column Names Specification** section, check **Use Classic Column Names** to use the original names of report columns in the output for batch reports. If you do not select this option, the new names of columns are used. For your reference, the online help displays the original names of those columns within parentheses ( ) next to their new names.

   Note: The titles of some reports have been updated. These new titles are always used for batch reports
4. In the **Script** section, check **Run Script** to run a script when the batch report is completed.

   Select this option and enter the name of the script you want to run against the results of the batch report. For example, if you create a batch report for the

most obsolete files on monitored computers, you can indicate a script that runs a program to archive all the files that appear in that report.

5. In the **Output File Specification** section, define the output file naming format for the batch report file.

   Enter Mask that will be used to create the output file name
   Displays the parameters that determine the name of the output file. For example, if you use the default parameters, {Report creator} {Report name} {Report run number}, Tivoli Storage Productivity Center retrieves the values for these parameters and uses them to name the output file. Click any of the following parameters to include them in the mask for the name of an output file:

   - Report creator (default)
   - Report name (default)
   - Report run number (default)
   - Report title
   - Report run date
   - Report run time
   - Unique number

   Note: If you include {Report name} in the mask, Tivoli Storage Productivity Center uses the name of the batch report definition in the name of the output file. If you included NLS characters in the name of the batch report definition, errors might occur if you selected a Storage Resource agent as a destination for the batch report. You can use NLS characters in the name of a batch report definition in the following situations:

   - TPC server is selected as the destination for a batch report.
   - A Storage Resource agent is selected as the destination for a batch report and the {Report name} parameter is not included in the mask for the output file name. If {Report name} is not used in the mask, it is recommended that other parameters are included in the mask to ensure that the output file name is unique. For example, `{Report Creator} {Report Run Number} {Report Title} {Unique number}`.

Scheduling a batch report:

Determine when and how often to run a batch report.

Use the **When to Run** page to specify when and how often to run the batch reporting job.

You have the following options for determining when to generate a report:

- Generate the report immediately
- Generate the report once at a specified time and date
- Generate the report repeatedly according to a schedule you specify

Select the option you want and then edit the time and date information accordingly. You can also use the fields in the How to handle time zones section to indicate the time zone that IBM Tivoli Storage Productivity Center should use when scheduling and running the batch report. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run the batch report at:

- 9:00 AM local time of the agent running the batch report

- 9:00 AM of the time zone where the server resides
- 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the batch report will run at 7:00 AM on an agent located in the PST time zone.

Local time in each timezone. Select this option to use the time zone of the computer where the agent that collects data for the job is deployed. For probes of storage systems and fabrics, the time zone of the Tivoli Storage Productivity Center server is used.

Same Global time across all timezones. Use the options in this section to apply the same global time across all the time zones where the probe is being run.

- **Use the timezone that the server runs in.** Select this option to use the time zone of the location where the Tivoli Storage Productivity Center server resides.
- **Use this timezone.** Select this option to indicate a specific time zone for when to run the job. Select a timezone from the list box.

Defining alerts for batch reports:

Define the alert that will be triggered if the batch report fails to run.

Use the **Alert** tab to define an alert that will be triggered if the report does not run on schedule.

To define alert values for a batch report, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** > **Reporting -> My Reports** and right-click on **Batch Reports**.
2. Select **Create Batch Reports**. The Create Batch Reports window opens.
3. For information on creating the batch report, see "Batch reports" on page 1562.
4. Click the **Alerts** tab.
5. Define the conditions that will trigger the alert. The default setting is **Report Failed**. No other alerting conditions are currently available.
6. Select alert mechanisms in addition to writing an error to the error log:
   - If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** > **Configuration** > **Alert Disposition**.
   - If you want a script to be run after an event has been triggered, you must create the script before you create the alert.
7. Save the file:
   a. Click **File** > **Save**.
   b. The **Save As** window opens. Enter a name for the batch report and click **OK**.

Editing alert values for a batch report

1. Expand **IBM Tivoli Storage Productivity Center** > **My Reports** > **Batch Reports**. The list of existing batch reports is displayed.
2. Click the batch report. The **Edit Batch Report** window opens.
3. In the **Alerts** tab, make any necessary changes.
4. Click **File** > **Save**.

Related topics:

"Batch reports" on page 1562

Creating batch reports for storage subsystem or switch performance:

This topic describes how to create a batch report for storage subsystem performance or switch performance.

To create a batch report for storage subsystem performance or switch performance, follow these steps:

1. Expand **IBM Tivoli Storage Productivity Center > Reporting > My Reports**.
2. Right-click the **Batch Reports** node and select **Create Batch Report**.
3. Click the **Report** tab and select the type of report you want to create and move it from **Available** to **Current Selections** by clicking the right-arrow button (>>). Select an item from **Storage Subsystem Performance** or **Switch Performance**.
4. Click the **Selection** tab. The Report Filter Specifications window displays.

   Selection button

   This determines the resources upon which you want to report, such as the storage subsystems, the fibre-channel ports, and so forth. The Select Resources window displays. The type of resources displayed for selection depends on the type of report you previously selected on the **Report** tab. For example, if you selected a By Storage Subsystem report, all displayed resources will be storage subsystems. If you clear the box for a resource on the Select Resources window, a warning icon displays on the Selection button.

   Filter button

   This further defines the data that you want to include in the report. Filters enable you to apply general rules to the report. When you click this button, the Edit Filter window displays.

   For example, to select only data associated with subsystems that have a name ending with "123–IBM", select the "Subsystem" column, the "LIKE" operator, and type "*123–IBM" (without the quotation marks) in the Value 1 field. To select only data which has a Read I/O Rate greater than 0, select the "Read I/O Rate (overall)" column, the ">" operator, and type "0" (without the quotation marks) in the Value 1 field.

   If a filter is defined for this report, a warning icon displays on the **Filter** button.

   Time range selectors

   All performance reports display historical performance data. Use the time range selectors to specify the exact time range that you want the report to cover. Click one of the following radio buttons to select the type of historical report you want:

   Display latest performance data

   This report contains only the latest available performance data for the selected resources. Note that even the latest data available might still be somewhat old, depending on whether data was recently collected for the resources. Including the time column in the report will help you determine how old the latest available data actually is. Regardless of the age of the data, clicking this button ensures that only a single row (representing a single point in time) for each selected resource is included in the report. If no data has ever been collected for

a particular resource, then that resource is not included in the report. If no data has ever been collected for any resource, then the resulting report will be empty.

Display historic performance data using absolute time
This report contains only data that falls within a specific time range. Select the time range using the **From** and **To** selectors. The specified from and to values are interpreted as inclusive. The resulting report contains all the performance data that is available within the specified time range, for each selected resource. This means that the data can consist of potentially many rows per resource, each with a unique timestamp. If no data has been collected for a particular resource during the specified time range, that resource is not included in the report. If no data has been collected for any resource during the specified time range, then the resulting report will be empty.

To determine which historical performance records fall within the specified time range, the **beginning timestamp** of the collection intervals is used for comparison. If an historical record contains performance data for an interval of time from 1:00 PM to 1:05 PM, that record would be included in the report for a time range of 1:00 PM to 8:00 PM. However, that same record would not be included if a time range of 1:02 PM to 8:02 PM was specified, because it is the beginning of the time period (for example, 1:00 PM in this example) that is used to determine inclusion in the specified time interval.

Specifying an absolute time range is not suitable for batch reports that are run on a repeating schedule, because the specified time range is fixed and will not automatically be updated based on the most current date and time. The result will be that the same report is generated every time that the batch report is run. Instead, use a relative time range for this type of report.

Display historic performance data using relative time
This report contains only data that falls within a specific time range extending backward from the current date and time. Specify the number of days in the desired time range, using the **days ago until now** field. The report will be exactly the same as a report using absolute time, with the **To** selector set to the current date and time, and the **From** selector set to the current date and time minus the number of **days ago until now**. However, when regenerating this report periodically, as the current date and time advances, the data included in the report will advance as well. As a result, this is the suitable option to select when defining batch reports that are run on a repeated schedule.

Note that the same criteria applies as if determining which historical performance records fall within the specified time range. Only the beginning timestamp of the collection intervals is used for comparison. This means that data for the current hour or the current day is not included in the report when selecting hourly or daily summation data, because data for the current hour or current day is only saved in the database when the hour or day is complete. As an alternative, consider using

the first radio button, **Display latest performance data**, to see a report showing the most recent daily summation data.

Summation Level

Use this selector to select which type of performance data records to include in the report. The Tivoli Storage Productivity Center Performance Manager collects and stores three different types of performance data:

By sample

Represents the most detailed data. There is usually one sample record per resource for every **n** minutes, where **n** is the performance monitor interval length that was specified when the monitor was defined. Each sample represents the average performance of the resource over those **n** minutes. A sample record with a timestamp of **t** will be saved at time **t+n**. In other words, the beginning of each **n** minute time interval is used for display and for time range matching.

Hourly

Represents an hourly average of the performance of the specified resources. There will be one hourly record written every hour, per resource. Each such record will be written at or shortly after each hour mark on the (device) clock. For example, a record with a timestamp of 12:00 PM and interval length of 1 hour will be written at 1:00 PM.

Daily Represents a daily average of the performance of the resources. There will be one daily record written every day, per resource. Each daily record will be written at or shortly after midnight, according to the device clock. For example, a record with a timestamp of April 21, 12:00 AM and interval length of 24 hours will be written on April 22, 12:00 AM.

Note that hourly and daily records are usually retained in the database for longer periods of time, compared to By Sample records. As a result, if a generated report indicates that sample data is no longer available for a particular time range in the past, it is possible that the sample data has already been purged, in which case you can try generating the report using hourly or daily summation levels.

Be aware that when selecting a time range using the **Display historic performance data using relative time** radio button, it is necessary to select the length of the range in days. This can lead to unexpected results when also selecting the daily summation level. Remember that a daily record is not written until the day is complete, which means that at any given time during the day, the current day's data does not yet exist in the database. Therefore, selecting the daily summation level as well as a relative time range of n **days ago until now**, will usually result in n-1 records being returned per resource, assuming that at least n-1 days worth of data has been previously collected for those resources. Specifying a relative time range of 1 **days ago until now** will result in an empty report.

Available Columns, Included Columns

Use these items to specify which columns to include and which columns to exclude from the generated report.

Available Columns

Displays the columns that you can include in the generated

report that are not already included. If a column appears in the **Included Columns** list box, it will not appear in the **Available Columns** list box.

Included Columns
Displays the columns that will appear in the generated report. By default, all columns for a report are listed in this list box and will be displayed in the generated report. You can highlight the columns and use the up and down arrows to rearrange the order of the columns in the report. The order of the items in this list box determines the order in which the columns will appear in the generated report. The columns with grayed-out names are the frozen columns, which cannot be selected, removed, or reordered for this type of report.

If a column appears in the **Available Columns** list box, it will not appear in the **Included Columns** list box.

To exclude a column from a report, click the name of the column in the **Included Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the left-arrow button (<<). The selected columns are removed from the **Included Columns** list box, and will appear in the **Available Columns** list box.

To include a column in a report, click the name of the column in the Available Columns list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the right-arrow button (>>). The selected columns will be removed from the **Available Columns** list box, and will appear in the **Included Columns** list box.

To rearrange the order of columns that will appear in the generated report, click the name of a column in the **Available Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the up-arrow button (^) or the down-arrow button (v) to move the selected columns up or down in the list, relative to the other columns.

5. Click the **Options** tab of the Batch Report job window, specify the following:
   - The machine on which the report file will be saved
   - The report format (for example, HTML)
   - Whether to run a script when the report process is complete
   - The format for the name of the batch report
6. In the **When to Run** tab of the Batch Report job window, specify when and how often to run the batch report.
7. In the **Alert** tab of the Batch Report job window, define the conditions that will trigger an alert and the method of notification.
8. Click **File > Save** and type a name for the batch report. The batch report will be saved using your user ID as a prefix. For example, if you are logged onto Tivoli Storage Productivity Center as "Smith," the name of the batch report will be Smith.*report_name.*

Editing batch reports

Edit existing batch report definitions.

To edit a batch-report definition, follow this procedure:

1. Expand **IBM Tivoli Storage Productivity Center → Reporting → My Reports → Batch Reports**.
2. Click the report name, make your changes and then click **File → Save**.

Viewing job information for a batch report

View job information for batch reports on the Job Management panel. Use this panel to check the status of a batch report schedule and view the runs, jobs, and log files associated with it.

To display a list of the runs and jobs for a batch report, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center → Reporting → My Reports → Batch Reports**.
2. Click the name of a batch report. The content pane displays information about the batch report schedule.
3. Click **View Job History**. The Job Management panel is displayed and the batch report schedule is highlighted in the **Schedules** section.
4. View the runs and jobs for the highlighted batch report schedule in the **Jobs for Selected Schedule** section. See "Managing schedules, jobs, and runs" on page 605 for more information about how to use the Job Management panel.

Displaying log files for a batch report

Display the log files batch reports.

To display the log files for a batch report schedule, complete these steps:

1. Expand **IBM Tivoli Storage Productivity Center → Reporting → My Reports → Batch Reports**.
2. Click the name of a batch report. The content pane displays information about the batch report schedule.
3. Click **View Job History**. The Job Management panel is displayed and the batch report schedule is highlighted in the **Schedules** section.
4. In the **Jobs for Selected Schedule**, expand the run whose log files you want to view. Click **+** or **-** to expand or collapse the job listing under a run.
5. Select the job whose log file you want to view and click **View Log File(s)**. See "Managing schedules, jobs, and runs" on page 605 for more information about how to use the Job Management panel.

Displaying batch reports from the saved directory

Access the report file that is generated from a batch report schedule. This file contains the data for the report that you generated as a batch report.

After you run a batch report, batch report files are saved in the directory that you specified on the **Options** page of the Create Batch Report panel. Batch reports for Storage Resource agents are automatically saved to the following default locations:

- Windows: C:\Program Files\IBM\TPC\agent\log\computername\reports
- UNIX/Linux: /opt/IBM/TPC/agent/log/computername/reports

Viewing CSV batch reports with spreadsheet applications

View CSV-formatted batch reports in third-party spreadsheet applications.

If you are using a spreadsheet program (such as Excel) to view a CSV-type batch report file, that program might display data in the wrong columns either because

the data was exported from Data Manager, the data contains commas, or the file was opened or formatted in the spreadsheet program by double-clicking it.

If data is displaying in the wrong columns, open the CSV file using the spreadsheet program's open command.

Exporting and importing settings that you define

This topic describes how to export batch report definitions, constraints, groups, and profiles.

You can export settings that you define from one instance of IBM Tivoli Storage Productivity Center and import the settings into another instance. To export and import your definitions, complete the following steps:

1. On the system running the server component, change to the following directory:

   On Linux and UNIX operating systems

   ```
   /opt/IBM/TPC/data/server/tools/
   ```

   On Windows operating systems

   ```
   C:\Program files\IBM\TPC\data\server\tools\
   ```

2. To export the definitions from a command prompt, type the following command and press **Enter**:

   ```
   export -e <file>
   ```

   where <file> is the fully qualified name of the file where the definition information is written.

3. On the system to which you want to migrate the definitions, change to the **tools** directory.
4. To import the definitions, type the following command and press Enter:

   ```
   export -i <file>
   ```

   where <file> is the fully qualified name of the file that you created when exporting.

User ID (saved) reports

Use the information in this section to learn how to generate, save, and view user ID reports. You can create a user ID reports based on a system report or any of the online reports available in Data Manager, Data Manager for Databases, Disk Manager, or Fabric Manager. Once saved, user ID reports are available in the **IBM Tivoli Storage Productivity Center** > **Reporting >My Reports** > **user ID name** node in the navigation tree.

Creating a user-defined report

Create a custom report based on a system report or any of the online reports available in Data Manager, Data Manager for Databases, Disk Manager, or Fabric Manager.

To create a custom report, follow this procedure:

1. Specify new report criteria by doing one of the following:
   - Access a system report and change its format according to your organization's needs.
   - Create a new report using the reporting component in any of the IBM Tivoli Storage Productivity Center managers. For example, **Data Manager → Reporting → Usage → Access Load → Access Time → By Filesystem**.

2. Click **File → Save** and a type a name for the report. The report is saved in the **Tivoli Storage Productivity Center > Reporting > My Reports > *user name's*** node, where ***user name's*** is your login name. For example, if you are logged in as "Smith", your custom reports appear under the **Smith's Reports** node.

Displaying user-defined reports

Display a custom report saved under your user ID.

To display a list of custom reports, expand the **IBM Tivoli Storage Productivity Center > Reporting > My Reports > <*your user ID*>** node. Click the name of the report you want to view.

Editing user-defined reports

Change the settings for an existing user ID report and save those changes.

To edit a custom report definition, follow this procedure:

1. Expand **My Reports → *user name's* Reports**.
2. Click the report name, make your changes and then click **File → Save**.

Rollup reports

Use rollup reports to view the storage information that is collected by the master IBM Tivoli Storage Productivity Center server from its associated subordinate servers. Rollup reports offer a network-wide perspective of storage usage in an environment where multiple Tivoli Storage Productivity Center servers are deployed and collecting storage metrics.

Before you can generate rollup reports, you must perform the following tasks:

- Associate subordinate servers with the master server on the **Administrative Services** > **Data Sources** > **Tivoli Storage Productivity Center Servers window**.
- Run discoveries, probes, and scans from the subordinate servers to gather information about their locally monitored storage entities.
- Run discoveries, probes, and scans from the master server to gather information about its locally monitored storage entities. Information about the storage entities that are monitored by a master server are automatically shown in rollup reports.
- Run TPC Server probes from the master against the subordinate servers to collect their storage information for the rollup reports. Use the Tivoli Storage Productivity Center **Monitoring** > **TPC Server Probes** node to define and schedule probe jobs.

  Note:

  – If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

  – After upgrading to Tivoli Storage Productivity Center 5.1, you must run probes against the monitored storage assets for the master server to have information about those assets displayed in **Tivoli Storage Productivity Center** > **Reporting** > **Rollup Reports**.

Tivoli Storage Productivity Center includes different rollup reporting categories that provide you with the flexibility to view data about your storage resources according to the needs of your environment. The following rollup reporting categories are available:

Generating rollup reports

Learn how to generate and view rollup reports.

1. Expand the navigation tree to display the type of report you want to generate. For example, if you want to generate a database capacity report, expand **IBM Tivoli Storage Productivity Center** > **Reporting** > **Rollup Reports** > **Database Capacity**.
2. Highlight a By Clause to determine how you want to view a report. For example, to view rollup capacity information according to computer group, click **Tivoli Storage Productivity Center** > **Reporting** > **Rollup Reports** > **All DBMSs** > **By Computer Group**. The **Selection** page displays.
3. Use the **Selection** page to determine what columns appear within a report, what storage assets to report upon (Capacity and Database Capacity reports only), and how to filter that values that appear report columns.
4. (Capacity and Database Capacity reports only) Click **Selection...** to select the objects that you want to report upon from the Select Resources window. See "Selecting resources to display on a report" on page 1324 for information.
5. Click **Filter...** to further filter the objects that appear in a report. Filters enable you to apply general rules to the report based on the rows in that report. See "Filtering the resources that appear on a report" on page 1326 for more information on filtering.
6. Use the **Available Columns** and **Included Columns** list boxes to determine what columns are displayed in a generated report.
7. Click **Generate Report**. A new tab will be added to the tab dialog representing the report that you generated.
8. Click the new tab to view the report. Reports are tabular in format and composed of rows and columns. You can scroll the report to view the entire report. Use the **View** menu to hide/show the navigation tree to increase the viewable area of the report or drag the divider bar in the middle of the screen back and forth to reallocate the amount of space that is given each pane.

Rollup reports - fragmentation:

Fragmentation occurs when related storage entities are being monitored by different subordinate IBM Tivoli Storage Productivity Center servers.

The following example shows when fragmentation can occur using subordinate servers and rollup reports:

- **Example:** Host A is assigned a volume from storage subsystem 1. Host A is being monitored only by one Tivoli Storage Productivity Center subordinate server (Server 1) and storage subsystem 1 is being monitored by another subordinate server (Server 2). Server 1 and Server 2 are subordinate servers to the master server (Server 3).

  Reason for fragmentation: In this example, the host and subsystem are fragmented because they are not being monitored by the same Tivoli Storage Productivity Center server. This will cause the **Correlated Volume Space** value in the rollup reports for storage subsystems and volumes to be incorrect.

To prevent fragmentation when using master and subordinate server relationships to generate rollup reports, we recommend the following:

- Do not monitor a storage subsystem and the hosts to which its volumes are assigned with different subordinate servers that are associated with the same master server.

- Do not monitor a fabric and its related switches with different subordinate servers that are associated with the same master server.
- If an agent machine is using storage from a subsystem, then both the agent and subsystem should be monitored by the same subordinate server.
- Configure all the subordinate servers to discover all the storage subsystems within your environment. Note that if a subordinate server is monitoring a subsystem with an agent other than the agent that is using storage on that subsystem, the value for the **Correlated Volume Space** column in rollup reports will be incorrect. For example:
  - Host A has storage from Subsystem 1
  - Subordinate server B is Monitoring Host A and Subsystem 1
  - Subordinate server C is Monitoring Subsystem 1 and probed the subsystem more recently than server B.
  - The master server is monitoring subordinate servers B and server C.

  The master server will report the storage subsystem data from server C. Because the master server does not know about Host A, the value for the **Correlated Volume Space** column will be incorrect.

Rollup Reports - Asset

Use these reports to view detailed statistics about agents (IBM Tivoli Storage Productivity Center), computers, storage subsystems, disk and volume groups, disks, file systems, logical volumes, volumes, and fabrics that are monitored by subordinate servers.

The information that is shown in these reports is gathered by the master server when it runs Tivoli Storage Productivity Center server probes against associated subordinate servers. If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

Rollup Reports - Asset - Agents:

Use this report to view information about Fabric agents and Storage Resource agents that are associated with subordinate servers in your environment.

You can generate and sort the data within this report using the following By clauses:

- By Agent: Shows agents sorted according to the name of the computer on which they are installed.
- By OS Type: Shows agents sorted according to the operating system under which they run. Click the magnifying glass icon next to an operating system type to view more detailed information about each of the agents running under that operating system.
- By TPC Server: Shows agents sorted according to the subordinate server that manages them. Click the magnifying glass icon next to subordinate server to view more detailed information about each of the agents associated with that server.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Agent

Agent Name
Name of the computer on which an agent is installed.

Data/SR Agent Version
Tivoli Storage Productivity Center version number of a Storage Resource agent (for example, 4.2.0).

Fabric Agent Version
Tivoli Storage Productivity Center version number of a Fabric agent (for example, 4.1.1).

OS Type
Operating system on the computer where an agent is installed.

Data/SR Agent Status
Status of a Storage Resource agent. Possible statuses include Up and Down. A down agent represents an agent that is still running but with which the server is not communicating.

Fabric Agent Status
Status of a Fabric agent. Possible statuses include Up and Down. A down agent represents an agent that is still running but with which the server is not communicating.

Data/SR Agent Port
Port number on which a Storage Resource agent listens for requests. Port 9549 is registered for use by IBM.

Fabric Agent Port
Port number on which a Fabric agent listens for requests. Port 9550 is registered for use by IBM.

Host Address
Network name or IP address of an agent as seen by subordinate server.

Time Zone
Time zone where the agent is installed (for example, CDT).

CPU Architecture
Instruction set architecture on the computer where an agent is running.

By OS Type

OS Type
Operating system on the computer where an agent is installed.

Number of Data Agents/SR Agents
Total number of Storage Resource agents that are running on the operating system.

Number of Data Agents/SR Agents (Status: Up)
Number of Storage Resource agents that are running on the operating system that have a status of Up.

Number of Data Agents/SR Agents (Status: Down)
Number of Storage Resource agents that are running on the operating system that have a status of Down. A down agent represents an agent that is still running but with which the server is not communicating.

Number of Data Agents/SR Agents (Status: Defected)
Number of Storage Resource agents that are running on the operating system that have been reassigned to another server.

Number of Data Agents/SR Agents (Status: Upgrading)
Number of Storage Resource agents that are running on the operating system and are currently in the process of upgrading to the same version as the subordinate server.

Number of Data Agents/SR Agents (Status: Unreachable)
Number of Storage Resource agents that are running on the operating system that are unreachable. This status indicates that an agent that is unreachable and not responding to the server. This status might occur when the agent is not running but did not inform the server that it was shutting down, or when a network problem has occurred.

Number of Data Agents/SR Agents (Status: Upgrade Required)
Number of Storage Resource agents that are running on the operating system that require an upgrade to the same version as the subordinate server.

Number of Data Agents/SR Agents (Status: Deleted)
Number of Storage Resource agents that are running on the operating system that have a status of Deleted.

Number of Fabric Agents
Total number of Fabric agents that are running on the operating system.

Number of Fabric Agents (Status: Up)
Number of Fabric agents that are running on the operating system that have a status of Up.

Number of Fabric Agents (Status: Down)
Number of Fabric agents that are running on the operating system that have a status of Down. A down agent represents an agent that is still running but with which the server is not communicating.

By TPC Server

TPC Server
Name of the TPC server on which an agent is associated.

Number of Data/SR Agents
Total number of Storage Resource agents that are associated with the TPC server.

Number of Data Agents/SR Agents (Status: Up)
Number of Storage Resource agents that are associated with the TPC server that have a status of Up.

Number of Data Agents/SR Agents (Status: Down)
Number of Storage Resource agents that are associated with the TPC server that have a status of Down. A down agent represents an agent that is still running but with which the server is not communicating.

Number of Data Agents/SR Agents (Status: Defected)
Number of Storage Resource agents that are associated with the TPC server that have been reassigned to another server.

Number of Data Agents/SR Agents (Status: Upgrading)
Number of Storage Resource agents that are associated with the TPC server and are currently in the process of upgrading to the same version as the subordinate server.

Number of Data Agents/SR Agents (Status: Unreachable)
Number of Storage Resource agents that are associated with the TPC

server that are unreachable. This status indicates that an agent that is unreachable and not responding to the server. This status might occur when the agent is not running but did not inform the server that it was shutting down, or when a network problem has occurred.

Number of Data Agents/SR Agents (Status: Upgrade Required)
Number of Storage Resource agents that are associated with the TPC server that require an upgrade to the same version as the subordinate server.

Number of Data Agents/SR Agents (Status: Deleted)
Number of Storage Resource agents that are associated with the TPC server that have a status of Deleted.

Number of Fabric Agents
Total number of Fabric agents associated with the TPC server.

Number of Fabric Agents (Status: Up)
Number of Fabric agents associated with the TPC server that have a status of Up.

Number of Fabric Agents (Status: Down)
Number of Fabric agents associated with the TPC server that have a status of Down.

Rollup Reports - Asset - Computers:

Use this report to view information about the hosts and computers (including NetApp) that are monitored by subordinate servers in your environment. This includes VMWare-related entities such as hypervisors and virtual machines.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

You can generate and sort the data within this report using the following by clauses:

- By Computer: Shows computers monitored by subordinate servers. Sorted by computer name.
- By TPC Server: Shows subordinate servers that are monitoring computers. Sorted by subordinate server name. Click the magnifying glass icon next to subordinate server to view more detailed information about each of the computers associated with that server.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Computer

Note: Two entries for the same host might appear in this report and the values for that host are counted twice in the **TOTAL** column. This scenario occurs under the following circumstances:

- You install two Storage Resource agents on the same host (using the long and short names of the host) that point to different Tivoli Storage Productivity Center servers. For example, you install an agent from serverA using the short name of a host (for example, hostC). You install another agent from serverB using the long name of the same host (for example, hostC.storage.chicago.com). You then set up the servers in a master and subordinate relationship. For

example, serverA is the master and serverB is the subordinate. In this example, there are two entries in the report for the host: hostC and hostC.storage.chicago.com. The values for hostC are counted twice in the **TOTAL** column.

- A Storage Resource agent on a host (hostC) points to two different Tivoli Storage Productivity Center servers (serverA and serverB). The host has a dynamic IP address. You set up serverA as the master and serverB as the subordinate. Server serverA probes hostC, but then the IP address of hostC changes before serverB probes the host. In this example, the report might display two entries for hostC: one entry with the IP address that was probed by serverA; the other entry with the changed IP address that was probed by serverB. The values for hostC are counted twice in the **TOTAL** column.
- A Storage Resource agent on a host (hostC) points to two different Tivoli Storage Productivity Center servers (serverA and serverB). The host has a name that resolves to hostC.storage.chicago.com on serverA and hostC.srm.storage.chicago.com on serverB. You then set up the servers in a master and subordinate relationship. For example, serverA is the master and serverB is the subordinate. In this example, there are two entries in the report for the host: hostC.storage.chicago.com and hostC.srm.storage.chicago.com. The values for hostC are counted twice in the **TOTAL** column.

Computer
Name of a computer monitored by a subordinate server.

OS Type
Operating system running on a computer.

OS Version
Version of the operating system running on a computer.

RAM Amount of RAM installed on a computer.

Disk Space (Capacity)
Total storage capacity for a computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Note: If the computer is an IBM Tivoli Storage SAN File Systems client, then this value will also include the capacity of any storage subsystem volumes visible to this computers that are owned by the SAN File System.

Number of Disks
Total number of disks for a computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value in this column:

- Computer disks on a computer that has not been probed by a Storage Resource agent.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Processor Count
Number of CPUs running on a computer.

IP Address
IP address of a computer.

Serial Number
Serial number of a computer.

Model Model number or name of a computer.

Manufacturer
Manufacturer of a computer (for example, IBM).

Time Zone
Time zone where the computer is located.

Available File System Space (Filesystem Free Space)
Total amount unique file system storage space that is not used or available to the computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Network Address
Network name of a computer.

Processor Type
Type of processor running on a computer (for example, GenuineIntel: x86 Family 15).

Processor Speed
Speed at which a computer's processor is running (in MHz).

CPU Arch
Architecture of the processor (for example, IA32).

Swap Space
Amount of swap space (virtual memory) defined on a computer.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring computers in your environment.

# of Computers
Number of computers monitored by the subordinate server.

Disk Space (Capacity)
Total storage capacity of the computers monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Note: If a computer is an IBM Tivoli Storage SAN File Systems client, then this value will also include the capacity of any storage subsystem volumes visible to this computers that are owned by the SAN File System.

Available File System Space (Filesystem Free Space)
Total amount unique file system storage space that is not used or available to the computers that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Number of Disks
Number of disks associated with the computers monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value in this column:

- Computer disks on a computer that has not been probed by the Storage Resource agent.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Rollup Reports - Asset - Storage Subsystems:

Use this report to view information about the storage subsystems that are monitored by subordinate servers in your environment.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

You can generate and sort the data within this report using the following By clauses:

- By Storage Subsystem: Sorted according to storage systems name. The names of storage systems monitored by subordinate servers appear in the first column and remain displayed when scrolling through the other columns on the report.
- By Disk Space: Sorted according to the storage capacity of storage systems. The names of storage systems monitored by subordinate servers and their disk capacity appear in the first two columns and remain displayed when scrolling through the other columns on the report.
- By TPC Server: Shows storage systems aggregated at subordinate server level. The names of subordinate servers that are monitoring storage systems appear in the first column and remain displayed when scrolling through the other columns on the report. Click the magnifying glass icon next to subordinate server to view more detailed information about each of the storage systems monitored by that server.

Note:

- Storage systems that have been discovered but not probed are not displayed in these reports.
- This report includes information about storage system volumes. For IBM System Storage SAN Volume Controller systems, IBM Tivoli Storage Productivity Center uses the term "Volume" to represent "virtual disks" for those systems.

By Storage Subsystem, and By Disk Space

Name Name or alias of a storage system that is monitored by a subordinate server or the master server.

Disk Space (Capacity)
Total raw (unformatted) disk capacity of a storage system. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Capacity of spare disks identified on IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 storage systems.
- Capacity of storage system disks that become missing after a storage system probe.

Note: For Hitachi Data Systems storage systems, this value is the capacity of all the volumes in the storage system.

Available Disk Space (Unallocated Space)
Total unformatted disk free space of the storage system. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- RAID overhead on storage systems that are identified by Tivoli Storage Productivity Center.
- External back-end storage for TagmaStore systems.
- Storage space that is consumed by volumes that reside on external disks when Hitachi Data Systems (HDS) storage systems are configured as storage virtualizers.

Note: For Hitachi Data Systems storage systems, internal disk information is not available. This field is set to 0.

Number of Disks
Total number of disks in the storage system. Tivoli Storage Productivity Center does not include disks that become missing after a successful storage system probe in its calculation of the value for this column. For Hitachi Data Systems storage systems, internal disk information is not available. This field is to 1.

Number of Volumes (LUNs)
Number of storage volumes and virtual storage volumes discovered by a probe. Tivoli Storage Productivity Center does not include storage volumes or virtual storage volumes that become missing after a successful probe in its calculation of the value for this column.

Volume Space (LUN Capacity)
Total amount of unique storage system volume space or virtual storage volume space discovered on monitored storage system arrays. Tivoli Storage Productivity Center does not include storage volumes or virtual storage volumes that become missing after a successful probe in its calculation of the value for this column.

Note: For space-efficient volumes, this value is the amount of storage space that is requested for these volumes, not the actual allocated amount. This value can result in discrepancies in the overall storage space reported for a storage system that uses space-efficient volumes. This situation also

applies to other space calculations, such as the calculations for the storage system's Consumable (Available) Volume Space and FlashCopy Target Volume Space.

Consumable Volume Space (Usable LUN Capacity)
Total amount of unique storage system volume space or virtual storage volume space discovered on monitored storage system arrays. Tivoli Storage Productivity Center does not include the following storage entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe
- Storage system volumes or virtual storage volumes that are used as flash copy targets.
- Storage system volumes or virtual storage volumes that are identified as a Business Continuance Volume extended (BCVx) (EMC only).
- Storage system volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database (EMC only).

FlashCopy Target Volume Space (Capacity)
Total amount of unique storage system volume space and virtual storage volume space discovered on monitored storage system arrays and has been identified as flash copy target space. Tivoli Storage Productivity Center does not include the following storage entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes or virtual storage volumes that are not used as flash copy targets.

Note: TotalStorage Enterprise Storage Server systems must have at least the following microcode levels: ess800mincodelevel = 2.4.3.56, essf20mincodelevel = 2.3.3.89. Tivoli Storage Productivity Center does not report FlashCopy information if the TotalStorage Enterprise Storage Server systems do not meet this requirement.

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Total amount of storage system volume space or virtual storage volume space on systems where the systems disks are identified by Tivoli Storage Productivity Center as residing within the monitored storage systems in a group. This includes storage system volumes or virtual storage volumes identified as back-end storage volumes on monitored storage system virtualizers.Tivoli Storage Productivity Center does not include the following entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that are assigned to a monitored storage system but are not identified by Tivoli Storage Productivity Center as a disk residing within the storage system. Reasons why these disks are not identified as residing within the storage system include:
  - Tivoli Storage Productivity Center does not support the identification of disks for some storage system types.
  - The connection of the disk to the storage system is lost and is no longer detectable during a probe.
- Storage system volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).

- Storage system volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.

Volume Space Not Correlated (LUN Capacity Not Visible to Data Manager)
Total amount of unique storage system volume space and virtual storage volume space discovered on monitored storage system arrays in a group that can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that reside within a monitored storage system.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers.

Unavailable Disk Space (Overhead)
Amount of storage system volume or virtual storage volume space that is dedicated to redundancy. This value is dependent on the storage system SMI-S provider returning a valid RAID value that can be used to determine the overhead. If the value is zero for a storage system, the overhead cannot be calculated.Tivoli Storage Productivity Center does not include the storage system volumes or virtual storage volumes that become missing after a probe in its calculation of the value for this column.

Formatted Space (Capacity)
Total amount of formatted storage space associated with the storage system that is used or can be used for system storage volumes or virtual storage volumes. This value is dependent on the storage system SMI-S provider returning a formatted value.

For some storage systems, this value is the total managed space which includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the total size of the managed disk groups created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

Formatted Space with No Volumes (Formatted Space with No Volumes)
Total amount of formatted storage space associated with the storage system that can be used for system storage volumes or virtual storage volumes. This value is dependent on the storage system SMI-S provider returning a formatted value.

For some systems, this value represents the remaining managed space that includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the space within the managed disk groups that does not include the size of the virtual disks created and the corresponding virtual disk extent overhead.

Overall Unavailable Disk Space (Total Overhead)
(TotalStorage Enterprise Storage Server, DS6000, DS8000,and SAN Volume Controller storage systems only.) Total amount of unformatted space within disk groups or array sites.

Unformatted Disk Space (Undefined Disk Group or Array Site Capacity)
This value does not apply to all storage systems. For TotalStorage Enterprise Storage Server, DS6000, DS8000 storage systems, this value is the total raw disk space of any unformatted disk groups or array sites. For IBM XIV Storage System, this value is the total available physical (hard) space that remains for pool creation. For other storage systems, this value is the total disk space that is not within a storage pool or**N/A** if the value is not applicable for the storage system.

Cache Amount of RAM or other cache for a storage system.

Note: This column is blank for SAN Volume Controllers.

Manufacturer
Manufacturer of a storage system.

Model Model number/name of a storage system.

Serial Number
Serial number of a storage system.

Firmware Revision
Microcode level of a storage system.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage system volume space within the storage system that is mapped or assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes that are not mapped or are not assigned to host systems.

Unassigned Volume Space (Subsystem Unassigned LUN Capacity)
Total storage system volume space within a storage system that is not mapped or not assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes that are mapped or assigned to host systems.

Unassigned and Available Formatted Space (Subsystem Total Formatted Free Space) Total storage system space that is not mapped or not assigned to host systems and the total amount of formatted storage space associated with the storage system that can be used for system storage volumes.

This total amount of formatted storage space is dependent on the storage system SMI-S provider returning a formatted value. For some systems, this value represents the remaining managed space that might include space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a probe.
- Storage system volumes that are mapped or assigned to host systems.

zOS Volume Space (LUN Capacity)
Total storage system space that is reserved for use on mainframe storage such as used by zOS. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on open system server or non-mainframe storage.

Open System Volume Space (LUN Capacity)
Total storage system space that is reserved for use on open system storage such as used by operating systems like Linux, AIX, HP/UX, Oracle Solaris, and Windows. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on mainframe storage.

Last Probed Timestamp
Date and time when a storage system was last probed by the subordinate or master server.

Physical Disk Space (Capacity)
Total amount of physical disk space discovered on the monitored storage system arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

Backend Volume Space (LUN Capacity)
Total amount of storage system volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage system virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are back-end disks on a storage system virtualizer that is not monitored.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring storage systems in your environment.

Number of Subsystems (Total)
Number of storage systems monitored by the subordinate server.

Disk Space (Total Disk Capacity)
Total raw (unformatted) disk capacity of the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Capacity of spare disks identified on TotalStorage Enterprise Storage Server, DS8000, and DS6000 storage systems.
- Capacity of storage system disks that become missing after a storage system probe.

Note: For Hitachi Data Systems storage systems, this value is the capacity of all the Volumes in the storage system.

Available Disk Space (Total Disk Unallocated Space)
Total unformatted disk free space of the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- RAID overhead on storage systems that are identified by Tivoli Storage Productivity Center.
- External back-end storage for TagmaStore systems.
- Storage space that is consumed by volumes that reside on external disks when Hitachi Data Systems (HDS) storage systems are configured as storage virtualizers.

Note: For Hitachi Data Systems storage systems, internal disk information is not available. This field is set to 0.

Number of Disks (Total)
Total number of disks in the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include disks that become missing after a successful storage system probe in its calculation of the value for this column. For Hitachi Data Systems storage systems, internal disk information is not available. This field is set to 1.

Number of Volumes (Total Number of LUNs)
Total number of storage volumes and virtual storage volumes discovered by the subordinate server. Tivoli Storage Productivity Center does not include storage volumes or virtual storage volumes that become missing after a successful probe in its calculation of the value for this column.

Volume Space (Total LUN Capacity)
Total amount of unique storage system volume space or virtual storage volume space discovered on storage system arrays that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include storage volumes or virtual storage volumes that become missing after a successful probe in its calculation of the value for this column.

Note: For space-efficient volumes, this value is the amount of storage space that is requested for these volumes, not the actual allocated amount. This value can result in discrepancies in the overall storage space reported for a storage system that uses space-efficient volumes. This situation also applies to other space calculations, such as the calculations for the storage system's Consumable (Available) Volume Space and FlashCopy Target Volume Space.

Consumable Volume Space (Total Usable LUN Capacity)
Total amount of unique storage system volume space or virtual storage volume space discovered on storage system arrays that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following storage entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe
- Storage system volumes or virtual storage volumes that are used as flash copy targets.
- Storage system volumes or virtual storage volumes that are identified as a Business Continuance Volume extended (BCVx) (EMC only).
- Storage system volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database (EMC only).

FlashCopy Target Volume Space (Total FlashCopy Target Capacity)
Total amount of unique storage system volume space and virtual storage volume space discovered on storage system arrays that are monitored by the subordinate server and have been identified as flash copy target space. Tivoli Storage Productivity Center does not include the following storage entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes or virtual storage volumes that are not used as flash copy targets.

Note: TotalStorage Enterprise Storage Server systems must have at least the following microcode levels: ess800mincodelevel = 2.4.3.56, essf20mincodelevel = 2.3.3.89. Tivoli Storage Productivity Center does not report FlashCopy information if the TotalStorage Enterprise Storage Server systems do not meet this requirement.

Unavailable Disk Space (Overhead)
Amount of storage system volume or virtual storage volume space that is dedicated to redundancy and is monitored by the subordinate server. This value is dependent on the storage system SMI-S provider returning a valid RAID value that can be used to determine the overhead. If the value is zero for a storage system, the overhead cannot be calculated.Tivoli Storage Productivity Center does not include the storage system volumes or virtual storage volumes that become missing after a probe in its calculation of the value for this column.

Correlated Volume Space (Total LUN Capacity Visible to Data Manager)
Total amount of storage system volume space or virtual storage volume space on systems where the systems disks are identified by Tivoli Storage Productivity Center as residing within the storage systems that are monitored by the subordinate server. This includes storage system volumes or virtual storage volumes identified as back-end storage volumes on monitored storage system virtualizers.

The value in this column reflects the systems that are monitored by the subordinate server.Tivoli Storage Productivity Center does not include the following entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that are assigned to a monitored storage system but are not identified by Tivoli Storage Productivity Center as a disk residing within the storage system. Reasons why these disks are not identified as residing within the storage system include:
  - Tivoli Storage Productivity Center does not support the identification of disks for some storage system types.

– The connection of the disk to the storage system is lost and is no longer detectable during a probe.

- Storage system volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage system volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.

Volume Space Not Correlated (Total LUN Capacity Not Visible to Data Manager)

Total amount of unique storage system volume space and virtual storage volume space discovered on monitored storage system arrays that can be assigned or are assigned to servers within the network. The value in this column reflects the storage systems that are monitored by the subordinate server.Tivoli Storage Productivity Center does not include the following entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that reside within a monitored storage system.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Centeridentifies as back-end storage volumes on monitored storage system virtualizers.

Overall Unavailable Disk Space (Total Formatted Capacity)

Total amount of formatted storage space that is used or can be used for system storage volumes or virtual storage volumes. The value in this column reflects the storage systems that are monitored by the subordinate server. This value is dependent on the storage system SMI-S provider returning a formatted value.

For some storage systems, this value is the total managed space which includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Note: This value does not apply to SAN Volume Controller.

Formatted Space (Total Formatted Space with No Volumes)

Total amount of formatted storage space that can be used for system storage volumes or virtual storage volumes. The value in this column reflects the storage systems that are monitored by the subordinate server. This value is dependent on the storage system SMI-S provider returning a formatted value.

For some systems, this value represents the remaining managed space that includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the total size of the managed disk groups created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

Formatted Space with No Volumes (Total Subsystem Total Formatted Free Space) Total storage system space that is not mapped or not assigned to host systems and the total amount of formatted storage space associated with

the storage system that can be used for system storage volumes. The value in this column reflects the storage systems that are monitored by the subordinate server.

This total amount of formatted storage space is dependent on the storage system SMI-S provider returning a formatted value. For some systems, this value represents the remaining managed space that might include space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the space within the managed disk groups that does not include the size of the virtual disks created and the corresponding virtual disk extent overhead.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a probe.
- Storage system volumes that are mapped or assigned to host systems.

Capacity Overhead (Total)
(TotalStorage Enterprise Storage Server and SAN Volume Controller only.) Sum of all the space in a disk group minus the formatted capacity of the storage systems monitored by a subordinate server.

Unformatted Disk Space (Total Undefined Disk Group or Array Site Capacity)
This value does not apply to all storage systems. For TotalStorage Enterprise Storage Server, DS6000, DS8000 storage systems, this value is the total raw disk space of any unformatted disk groups or array sites that are monitored by the subordinate server. For the XIV system, this value is the total available physical (hard) space that remains for pool creation. For other storage systems, this value is the total disk space that is not within a storage pool or **N/A** if the value is not applicable for the storage system.

Cache (Total)
Amount of RAM or other cache for the storage systems that are monitored by the subordinate server.

Assigned Volume Space
Total storage system volume space within the storage system that is mapped or assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes that are not mapped or are not assigned to host systems.

Unassigned Volume Space (Total Subsystem Assigned LUN Capacity)
Total storage system volume space within storage systems that is mapped or assigned to host systems. The value in this column reflects the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes that are not mapped or are not assigned to host systems.

Unassigned and Available Formatted Disk Space (Total Subsystem Unassigned LUN Capacity)

Total storage system volume space within storage systems that is not mapped or not assigned to host systems. The value in this column reflects the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes that are mapped or assigned to host systems.

zOS Volume Space (Total zOS LUN Capacity)

Total storage system space that is reserved for use on mainframe storage such as used by zOS. The value in this column reflects the storage systems that are monitored by the subordinate server.Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on open system server or non-mainframe storage.

Open System Volume Space (Total Open System LUN Capacity)

Total storage system space that is reserved for use on open system storage such as used by operating systems like Linux, AIX, HP/UX, Oracle Solaris, and Windows. The value in this column reflects the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on mainframe storage.

Physical Disk Space (Capacity)

Total amount of physical disk space discovered on the storage system arrays that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

Backend Volume Space (LUN Capacity)

Total amount of storage system volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage system virtualizer (for example, SAN Volume Controller). The value in this column reflects the volumes that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are back-end disks on a storage system virtualizer that is not monitored.

Rollup Reports - Asset - Disk/Volume Groups:

Use this report to view information about the disk and volume groups that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report by using the following By-clauses.

- By Computer: Shows disk and volume groups that are aggregated at the computer level and sorted by computer name.
- By Disk/Volume Group Space: Shows disk and volume groups that are aggregated at the computer level and sorted by storage capacity.
- By Available Disk/Volume Group Space: Shows disk and volume groups that are aggregated at the computer level and sorted by storage free space.
- By TPC Server: Shows disk and volume groups that are aggregated at the subordinate server level and sorted by subordinate server name. Click the magnifying glass icon next to the name of a subordinate server to view information about the computers it monitors.

Note: The first column in each report remains displayed when you scroll through the other columns on the report.

By Computer

Note: The names of some columns are changed. For your reference, the original names of those columns display within parentheses ( ) next to their new names.

Computer Name
Name of a computer with disk and volume groups that is monitored by a subordinate server.

Number of Disk/Volume Groups (Volume Groups)
Number of volume groups on a computer.

Disk/Volume Group Space (Total Capacity)
Total storage capacity of the disks and volume groups on a computer.

Available Disk/Volume Group Space (Total Free Space)
Amount of storage space available on the disk and volume groups that are monitored by a subordinate server.

OS Type
Operating system that runs on a computer.

Number of Disks
Number of disks on a computer.

Last Probe Time
Date and time when a computer was last probed.

By Disk/Volume Group Space (Free Space)

Computer Name
Name of a computer with disk and volume groups that are monitored by a subordinate server.

Disk/Volume Group Space (Total Free Space)
Amount of storage space available on the disks and volume groups that are associated with a computer.

Number of Disk/Volume Groups (Number of Volume Groups)
Number of volume groups on a computer.

Available Disk/Volume Group Space (Total Capacity)
Total storage capacity of the disks and volume groups on a computer.

OS Type
Operating system that runs on a computer.

Number of Disks
Number of disks on a computer.

Last Probe Time
Date and time when a computer was last probed.

By Available Disk/Volume Group Space (Capacity)

Computer Name
Name of a computer with disk and volume groups that is monitored by a subordinate server.

Available Disk/Volume Group Space (Total Capacity)
Total storage capacity of the disks and volume groups on a computer.

Number of Disk/Volume Groups (Volume Groups)
Number of volume groups on a computer.

Disk/Volume Group Space (Total Free Space)
Amount of storage space available on the disks and volume groups that are associated with a computer.

OS Type
Operating system that runs on a computer.

Number of Disks
Number of disks on a computer.

Last Probe Time
Date and time when a computer was last probed.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring computers with disk and volume groups.

Number of Disk/Volume Groups
Number of disk and volume groups that are monitored by a subordinate server.

Disk/Volume Group Space (Total Free Space)
Amount of storage space available on the disks and volume groups that are monitored by a subordinate server.

Available Disk/Volume Group Space (Total Capacity)
Total storage capacity of the disks and volume groups that are monitored by a subordinate server.

Rollup Reports - Asset - Disks:

Use this report to view information about the disks that are monitored by subordinate servers in your environment.

You can generate and sort the data within Disk reports using the following by clauses:

- By Available Disk Space (Computers): Shows disks aggregated at the computer level and sorted by unallocated space. The names of computers and their system type appear in the first two columns and remain displayed when scrolling through the other columns on the report.
- By Available Disk Space (Subsystems): Shows disks aggregated at the storage subsystem level and sorted by unallocated space. The names of storage subsystems and their system type appear in the first two columns and remain displayed when scrolling through the other columns on the report.
- By TPC Server (computers): Shows computer disks aggregated at the subordinate server level. The names of subordinate servers appear in the first column and remain displayed when scrolling through the other columns on the report. Click the magnifying glass icon next to a subordinate server name to view more information about the computers monitored by that subordinate server.
- By TPC Server (subsystems): Shows storage subsystem disks aggregated at the subordinate server level. The names of subordinate servers appear in the first column and remain displayed when scrolling through the other columns on the report. Click the magnifying glass icon next to a subordinate server name to view more information about the storage subsystems monitored by that subordinate server.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

By Available Disk Space (Computers)

TOTAL
Shows the system-wide totals for each of the columns in the report.

System Name
Name of a computer with disks that are monitored by a subordinate server.

OS Type
Operating system running on a computer.

Available Disk Space (Total Unallocated Space)
Amount of storage space on disks that is not allocated to any logical volume. Note: the value that appears in the TOTAL row for this column includes the unallocated capacity of all disks within the computers monitored by subordinate servers.

Disk Space (Total Capacity)
Amount of storage space available on the disks associated with a computer.

Number of Disks
Number of disks on a computer.

By Available Disk Space (Subsystems)

System Name
Name of a storage subsystem with disks that are monitored by a subordinate server.

Subsystem Type
Operating system running on a storage subsystem.

Available Disk Space (Total Unallocated Space)
Amount of space on the disk that is not allocated to any volumes for disks within storage subsystems). Note: the value that appears in the TOTAL row for this column includes the unallocated capacity of all disks within the storage subsystem, including the capacity of spare disks.

Disk Space (Total Capacity)
Amount of storage space available on the disks associated with a storage subsystem.

Note: For space-efficient volumes, this value is the amount of storage space that is requested for these volumes, not the actual allocated amount. This value can result in discrepancies in the overall storage space reported for a storage system that uses space-efficient volumes. This situation also applies to other space calculations, such as the calculations for the storage system's Consumable (Available) Volume Space and FlashCopy Target Volume Space.

Number of Disks
Total number of disks in a storage subsystem.

Physical Disk Space
Total amount of physical disk space discovered on the monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, IBM SAN Volume Controller) that are attached from a storage subsystem.

By TPC Server (Computers)

TOTAL
Shows the system-wide totals for each of the columns in the report.

TPC Server
Name of a subordinate server that is monitoring disks on computers.

Available Disk Space (Total Unallocated Space)
Amount of storage space on computer disks that is not allocated to any logical volume and are monitored by a subordinate server.

Disk Space (Total Capacity)
Amount of storage space available on all the computer disks that are monitored by a subordinate server.

Note: For space-efficient volumes, this value is the amount of storage space that is requested for these volumes, not the actual allocated amount. This value can result in discrepancies in the overall storage space reported for a storage system that uses space-efficient volumes. This situation also applies to other space calculations, such as the calculations for the storage system's Consumable (Available) Volume Space and FlashCopy Target Volume Space.

Number of Disks
Number of computer disks monitored by a subordinate server.

By TPC Server (Subsystems)

TPC Server
Name of a subordinate server that is monitoring disks on storage subsystems.

Available Disk Space (Total Unallocated Space)
Amount of storage space on storage subsystem disks that is not allocated to any volume and are monitored by a subordinate server.

Disk Space (Total Capacity)
Amount of storage space available on all the storage subsystem disks that are monitored by a subordinate server.

Number of Disks
Number of storage subsystem disks monitored by a subordinate server.

Physical Disk Space
Total amount of physical disk space discovered on the monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, IBM SAN Volume Controller) that are attached from a storage subsystem.

Rollup Reports - Asset - File Systems or Logical Volumes:

Use this report to view information about the file systems or logical volumes that are monitored by subordinate servers in your environment.

You can generate and sort the data within File Systems or Logical Volumes reports by using the following by-clauses.

- By Available File System or Logical Volume Space: Shows file systems that are aggregated at the computer level and sorted by free space and computer name.
- By TPC Server: Shows file systems that are aggregated at the subordinate server level and sorted by subordinate server name. Click the magnifying glass icon next to a subordinate server name to view more information about the file systems or logical volumes that are monitored by that subordinate server.

Note: The first column in each report remains displayed when you scroll through the other columns on the report.

Note: The names of some columns are changed. For your reference, the original names of those columns are within parentheses ( ) next to their new names.

By Available File System or Logical Volume Space

Computer Name
Name of a computer with file systems or logical volumes that are monitored by a subordinate server.

Available File System/Logical Volume Space (Free Space)
Amount of storage space available in the file systems or logical volumes on a computer.

Consumed File System/Logical Volume Space (Used Space)
Amount of storage space that is used by the file systems or logical volumes on a computer.

File System/Logical Volume Space (Capacity)
Total storage capacity of the file systems or logical volumes on a computer.

OS Type
Operating system on a computer.

Number of File Systems/Logical Volumes (Number of FSs/LVs)
Number of file systems or logical volumes on a computer.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring the file systems and logical volumes in your environment.

Available File System/Logical Volume Space (Total Free Space)
Amount of storage space available in all the file systems or logical volumes on the computers that are monitored by the subordinate server.

Consumed File System/Logical Volume Space (Total Used Space)
Amount of storage space that is used by all the file systems or logical volumes on the computers that are monitored by the subordinate server.

File System/Logical Volume Space (Total Capacity)
Total storage capacity of all the file systems or logical volumes on the computers that are monitored by the subordinate server.

Number of File Systems/Logical Volumes (Number of FSs/LVs)
Number of file systems or logical volumes on the computers that are monitored by the subordinate server.

Rollup Reports - Asset - Storage Subsystem Volumes:

Use this report to view information about the Storage Subsystem Volumes that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Storage Subsystem: Shows Storage Subsystem Volume capacity and count aggregated at the storage subsystem level.
- By Storage Subsystem Volume Space: Shows Storage Subsystem Volume capacity and count aggregated at the subsystem level. Sorted according to the storage capacity of Storage Subsystem Volumes.
- By TPC Server: Shows Storage Subsystem Volume capacity and count aggregated at the subordinate server level. Sorted by subordinate server name. Click the magnifying glass icon next to a subordinate server to view information about the storage subsystems it monitors.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

By Storage Subsystem

Storage Subsystem Name
Name of a storage subsystem containing Storage Subsystem Volumes that is monitored by a subordinate server.

Storage Subsystem Type
Type of storage subsystem, such as the IBM SAN Volume Controller.

Volume Space (LUN Capacity)
Total storage capacity of the Storage Subsystem Volumes in a storage subsystem.

Consumable Volume Space (Usable LUN Capacity)
Storage capacity of Storage Subsystem Volumes in the created disk array (not including BCVs, VCM database, and Replication targets).

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Sum of the capacity of all monitored Storage Subsystem Volumes visible to a monitored host computer and the capacity of all monitored Storage Subsystem Volumes that are back-end disks on a monitored Storage Subsystem Virtualizer.

Formatted Space (Formatted Capacity)
(IBM TotalStorage Enterprise Storage Server, IBM FAStT, HP StorageWorks, and Hitachi Data Systems only) Total amount of formatted storage space in the storage pools associated with a storage subsystem. For example, a RAID volume space has more capacity than a stripe-mirror volume space.

For SAN Volume Controller, this value represents the total size of the managed disk groups created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage capacity of all Storage Subsystem Volumes that the storage subsystem knows are assigned to a host computer.

zOS Volume Space (zOS LUN Capacity)
Total storage capacity for all the Storage Subsystem Volumes allocated to zOS. Data for this column is calculated for TotalStorage Enterprise Storage Server only. The value for this column is set to N/A for other storage subsystems.

Number of Volumes (Number of LUNs)
Number of Storage Subsystem Volumes associated with a storage subsystem.

Backend Volume Space
Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage subsystem virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are back-end disks on a storage subsystem virtualizer that is not monitored.

By Storage Subsystem Volume Space

Volume Space (LUN Capacity)
Total storage capacity of the Storage Subsystem Volumes in a storage subsystem.

Consumable Volume Space (Usable LUN Capacity)
Storage capacity of Storage Subsystem Volumes in the created disk array (not including BCVs, VCM database, and Replication targets).

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Sum of the capacity of all monitored Storage Subsystem Volumes visible to a monitored host computer and the capacity of all monitored Storage Subsystem Volumes that are back-end disks on a monitored Storage Subsystem Virtualizer.

Formatted Space (Formatted Capacity)
(IBM TotalStorage Enterprise Storage Server, IBM FAStT, HP StorageWorks, and Hitachi Data Systems only) Total amount of formatted storage space in the storage pools associated with a storage subsystem. For example, a RAID volume space has more capacity than a stripe-mirror volume space.

For SAN Volume Controller, this value represents the total size of the managed disk groups created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage capacity of all Storage Subsystem Volumes that the storage subsystem knows are assigned to a host computer.

zOS Volume Space (zOS LUN Capacity)
Total storage capacity for all the storage subsystem volumes allocated to zOS. Data for this column is calculated for TotalStorage Enterprise Storage Server and SAN Volume Controller and SAN Volume Controller only. The value for this column is set to N/A for other storage subsystems.

Number of Volumes (Number of LUNs)
Number of storage subsystem volumes associated with a storage subsystem.

Storage Subsystem Name
Name of a storage subsystem containing storage subsystem volumes that is monitored by a subordinate server.

Storage Subsystem Type
Type of storage subsystem, such as the IBM SAN Volume Controller.

Backend Volume Space
Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the backend disk space on a monitored storage subsystem virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are backend disks on a storage subsystem virtualizer that is not monitored.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring storage subsystems.

Storage Subsystem Count
Number of storage subsystems that are monitored by a subordinate server.

Volume Space (LUN Capacity)
Total storage capacity of the Storage Subsystem Volumes that are monitored by a subordinate server.

Consumable Volume Space (Usable LUN Capacity)
Storage capacity of Storage Subsystem Volumes in the created disk array (not including BCVs, VCM database, and Replication targets) that are monitored by a subordinate server.

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Sum of the capacity of all monitored Storage Subsystem Volumes visible to a monitored host computer and the capacity of all monitored Storage Subsystem Volumes that are back-end disks on a monitored Storage Subsystem Virtualizer associated with a subordinate server.

Formatted Space (Formatted Capacity)
(IBM TotalStorage Enterprise Storage Server, IBM FAStT, HP StorageWorks, and Hitachi Data Systems only) Total amount of formatted storage space in the storage pools associated with the storage subsystems that are monitored by a subordinate server. For example, a RAID volume space has more capacity than a stripe-mirror volume space.

For SAN Volume Controller, this value represents the total size of the managed disk groups created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage capacity of all Storage Subsystem Volumes that the monitored storage subsystems know are assigned to a host computer.

zOS Volume Space (zOS LUN Capacity)
Total storage capacity for all the Storage Subsystem Volumes allocated to zOS that are monitored by a subordinate server. Data for this column is calculated for TotalStorage Enterprise Storage Server and SAN Volume Controller only. The value for this column is set to N/A for other storage subsystems.

Number of Volumes (Number of LUNs)
Number of Storage Subsystem Volumes associated with the storage subsystems monitored by a subordinate server.

Backend Volume Space
Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the backend disk space on a monitored storage subsystem virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are backend disks on a storage subsystem virtualizer that is not monitored.

Rollup Reports - Asset - Fabrics:

Use this report to view information about the fabrics that are monitored by subordinate servers in your environment.

You can view the following reports about fabrics:

- SAN Assets (ALL): Use this report to view information about all the SAN assets for each subordinate server.
- SAN Assets (Connected Dev): Use this report to view information about all the connected devices for each subordinate server.
- SAN Assets (Switches): Use this report to view information about switches in your SAN for each subordinate server.

SAN Assets (ALL)

TPC Server
Name of a subordinate server that is monitoring fabrics.

Fabric ID
The ID of the fabric that the asset is in.

Type The type of asset discovered. For example: switch, host, host bus adapter, storage subsystem.

Label The label of the asset. For example: the host name, switch name, storage subsystem name, node name of the port.

Identifier
The unique identifier of the switch. For example, the switch WWN.

Status The status of the asset: Normal, Warning, or Critical.

SAN Assets (Connected Dev)

TPC Server
Name of a subordinate server that is monitoring a connected device.

Fabric ID
The ID of the fabric that the asset is in.

Type The type of asset discovered. For example: switch, host, host bus adapter, storage subsystem.

Label The label of the asset. For example: the host name, switch name, storage subsystem name, node name of the port.

Identifier
The unique identifier of the switch. For example, the switch WWN.

Status The status of the asset: Normal, Warning, or Critical.

SAN Assets (Switches)

TPC Server
Name of a subordinate server that is monitoring a switch.

Fabric ID
The ID of the fabric that a switch is in.

Type The type of switch discovered.

Switch Type
The switch type, such as physical.

Label The label of a switch.

Status The status of a switch: Normal, Warning, or Critical.

IP Address
The IP address of a switch.

Vendor ID
The manufacturer of the switch.

Model The model name and number of a switch.

Version
The version information of the switch.

Serial Number
The serial number of the switch.

Object ID
The object ID of the switch.

Management ID
The management ID of the switch.

Management Address
The URL of the switch element manager.

Domain
The name of the domain that the switch is in.

Contact
The contact information for the switch.

WWN The world-wide name (WWN) of the switch.

Location
The physical location of the switch.

Physical Switch WWN
The WWN for the parent or physical switch of a virtual switch. If this is not a virtual switch then this column is empty.

Rollup Reports - Asset - Hypervisors:

Use this report to view information about the hypervisors that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clause:

- By Hypervisor: Sorted according to hypervisor name. The names of hypervisors monitored by subordinate servers appear in the first column and remain displayed when scrolling through the other columns on the report.

By Hypervisors

Hypervisor Name
Name of a hypervisor monitored by a subordinate server.

OS Version
Version of the operating system running on a computer.

RAM The amount of RAM installed on a computer.

Disk Space
Total raw (unformatted) disk capacity of a storage subsystem. Tivoli

Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Capacity of spare disks identified on IBM TotalStorage Enterprise Storage Server, IBM DS8000, and IBM DS6000 storage subsystems.
- Capacity of storage subsystem disks that become missing after a storage subsystem probe.

Note: For Hitachi Data Systems storage subsystems, this is the capacity of all the Volumes in the storage subsystem.

Number of Disks
Total number of disks for a computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value in this column:

- Computer disks on a computer that has not been probed by a Storage Resource agent.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Processor Count
Number of CPUs running on a computer.

IP Address
IP address of a computer.

Serial Number
Serial number of a computer.

Model Model number or name of a computer.

Manufacturer
Manufacturer of a computer (for example, IBM).

Time Zone
Time zone where the computer is located.

Available File System Space
Total amount unique file system storage space that is not used or available to the computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

OS Type
Operating system running on a computer.

Network
Network name of a computer.

Processor Type
Type of processor running on a computer (for example, GenuineIntel: x86 Family 15).

Processor Speed
Speed at which a computer's processor is running (in MHz).

CPU Architecture
Architecture of the processor (for example, IA32).

Swap Space
Amount of swap space (virtual memory) defined on a computer.

Rollup Reports - Database Asset

Use these reports to view detailed statistics that have been collected by subordinate servers about the RDBMSs in your environment, including Oracle, SQL Server, Sybase, and UDB/DB2.

The information shown in these reports is gathered by the master server when it runs Tivoli Storage Productivity Center server probes against associated subordinate servers. If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

Rollup Reports - Database Asset - All DBMSs:

Use this report to view information about the RDBMSs (including Oracle, Sybase, SQL Server, and UDB/DB2) that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Instance: Shows databases for each computer sorted by computer name.
- By Version: Shows databases for each computer sorted by RDBMS version.
- By DBMS Type: Shows databases for each computer sorted by RDBMS type.
- By TPC Server: Shows databases aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

By Instance

Computer Name
Name of the computer on which an Instance resides.

Instance
Name of an instance.

RDBMS Type
RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

RDBMS Version
Version of the RDBMS under which an instance is running.

Port Port on which an instance listens for requests.

Home Directory
Path of the home directory for an instance.

Logonid
ID used by Tivoli Storage Productivity Center to log in to an instance.

Probe Status
Status of the last probe run against an instance.

Last Start Time
Date and time when an instance was last started.

Last Probe Time
Date and time when an instance was last probed.

Tablespace count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

By Version

RDBMS Version
Version of the RDBMS under which an instance is running.

Computer Name
Name of the computer on which an instance resides.

Instance
Name of an instance.

RDBMS Type
RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

Port Port on which an instance listens for requests.

Home Directory
Path of the home directory for an instance.

Logonid
ID used by Tivoli Storage Productivity Center to log in to an instance.

Probe Status
Status of the last probe run against an instance.

Last Start Time
Date and time when an instance was last started.

Last Probe Time
Date and time when an instance was last probed.

Tablespace Count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

By DBMS Type

RDBMS Type
RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

Computer Name
Name of the computer on which an instance resides.

Instance
Name of an instance.

RDBMS Version
Version of the RDBMS under which an instance is running.

Port Port on which an instance listens for requests.

Home Directory
Path of the home directory for an instance.

Logonid
ID used by Tivoli Storage Productivity Center to log in to an instance.

Probe Status
Status of the last probe run against an instance.

Last Start Time
Date and time when an instance was last started.

Last Probe Time
Date and time when an instance was last probed.

Tablespace count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring DBMSs in your environment.

Total Number of DBMSs
Number of DBMSs monitored by a subordinate server.

Number of UDB Instances
Number of UDB instances monitored by a subordinate server.

Number of Oracle Instances
Number of Oracle instances monitored by a subordinate server.

Number of SQL/Server Instances
Number of SQL/Server instances monitored by a subordinate server.

Number of Sybase Instances
Number of Sybase instances monitored by a subordinate server.

Tablespace Count
Number of tablespaces associated with the instances monitored by a subordinate server.

File Count
Number of data files associated with the instances monitored by a subordinate server.

Log File Count
Number of log files associated with the instances monitored by a subordinate server.

File Free Space
Amount of storage free space available on the data files associated with the instances monitored by a subordinate server.

Total Size
Amount of storage space allocated to the instances monitored by a subordinate server.

Total File Capacity (File Capacity)
Amount of storage space allocated to the data files associated with the instances monitored by a subordinate server.

Rollup Reports - Database Asset - Oracle, SQL/Server, Sybase, UDB:

Use this report to view information about Oracle, Sybase, SQL Server, and UDB/DB2 instances that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Instance: Shows databases for each computer sorted by computer name.
- By Version: Shows databases for each computer sorted by the version of the RDBMS under which a database is running.
- By TPC Server: Shows databases aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Instance

Computer Name
Name of the computer on which an Instance resides.

Instance
Name of an instance.

RDBMS Type
RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

RDBMS Version
Version of the RDBMS under which an instance is running.

Port Port on which an instance listens for requests.

Home Directory
Path of the home directory for an instance.

Logonid
ID used by Tivoli Storage Productivity Center to log in to an instance.

Probe Status
Status of the last probe run against an instance.

Last Start Time
Date and time when an instance was last started.

Last Probe Time
Date and time when an instance was last probed.

Tablespace count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

By Version

RDBMS Version
Version of the RDBMS under which an instance is running.

Computer Name
Name of the computer on which an instance resides.

Instance
Name of an instance.

RDBMS Type
RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

Port Port on which an instance listens for requests.

Home Directory
Path of the home directory for an instance.

Logonid
ID used by Tivoli Storage Productivity Center to log in to an instance.

Probe Status
Status of the last probe run against an instance.

Last Start Time
Date and time when an instance was last started.

Last Probe Time
Date and time when an instance was last probed.

Tablespace Count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring DBMSs in your environment.

Number of Oracle Instances
Number of Oracle instances monitored by a subordinate server.

Number of SQL/Server Instances
Number of SQL/Server instances monitored by a subordinate server.

Number of Sybase Instances
Number of Sybase instances monitored by a subordinate server.

Number of UDB Instances
Number of UDB instances monitored by a subordinate server.

Tablespace count
Number of tablespaces associated with the instances monitored by a subordinate server.

File Count
Number of data files associated with the instances monitored by a subordinate server.

Log File Count
Number of log files associated with the instances monitored by a subordinate server.

File Free Space
Amount of storage free space available on the data files associated with the instances monitored by a subordinate server.

Total Size
Amount of storage space allocated to the instances monitored by a subordinate server.

Total File Capacity
Amount of storage space allocated to the data files associated with the instances monitored by a subordinate server.

Rollup Reports - TPC Server Storage Space

Use these reports to view storage metrics that are related to the disk space, file system space, used file system space, and available file system space that is gathered by the subordinate servers in your environment.

The information that is shown in these reports is gathered by the master server when it runs "Collect storage statistics from Tivoli Storage Productivity Center servers" on page 331 against associated subordinate servers. If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

Rollup Reports - TPC Server Storage Space - Disk Space:

Use this report to view disk capacity information about the computers and storage systems that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Storage Subsystem: Shows disk capacity aggregated at the storage system level and sorted by storage system name. Storage systems that have been discovered but not probed are not displayed in this report.
- By Storage Subsystem Group: Shows disks capacity aggregated at the storage system monitoring group and storage system reporting group levels and sorted by group name. Storage systems that have been discovered but not probed are not included in the values shown in this report.
- By Computer: Shows disks capacity aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows disks capacity aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By Cluster: Shows disks capacity aggregated at the cluster level and sorted by cluster name.
- By TPC Server (Computers): Shows computer disk capacity aggregated at the subordinate server level and sorted by subordinate server name.
- By TPC Server (Storage Subsystems): Shows storage system disk capacity aggregated at the subordinate server level and sorted by subordinate server name. Storage systems that have been discovered but not probed are not included in the values shown in this report.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

By Storage Subsystem

Subsystem Name
Name of a storage system that is monitored by a subordinate server.

Subsystem Type
Type of storage system, such as the IBM System Storage SAN Volume Controller.

Disk Space (Disk Capacity)
Total raw (unformatted) disk capacity of a storage system.

Available Disk Space (Available Capacity)
Amount of storage space available on the disks associated with a storage system. For Hitachi Data Systems storage systems, internal disk information is not available. This field will be 0.

Number of Disks
Number of disks associated with a storage system. For Hitachi Data Systems storage systems, internal disk information is not available. This field will be set to 1.

Number of Volumes (Number of LUNs)
Number of volumes associated with a storage system.

Storage Subsystem Volume Space (LUN capacity)
Amount of storage space allocated to the volumes associated with the storage system.

Consumable Storage Subsystem Volume Space (Usable LUN Capacity)
Storage capacity of volumes in the created disk array (not including BCVs, VCM database, and Replication targets).

Correlated Volume Space (Visible LUN Capacity)
Capacity of all volumes visible to a Storage Resource agent associated with a storage system.

FlashCopy Target Storage Subsystem Volume Space (FlashCopy Target Capacity)
The total capacity of all FlashCopy target volumes in the system. Note: TotalStorage Enterprise Storage Server systems must have at least the following microcode levels: ess800mincodelevel = 2.4.3.56, essf20mincodelevel = 2.3.3.89. Tivoli Storage Productivity Center does not report FlashCopy information if the TotalStorage Enterprise Storage Server systems do not meet this requirement.

Volume Space Not Correlated (Not Visible LUN Capacity)
Usable volume capacity minus assigned volume capacity.

Unavailable Disk Space (Overhead)
Amount of capacity that is dedicated to redundancy.

Formatted Disk Space (Formatted Capacity)
(IBM TotalStorage Enterprise Storage Server, FAStT, HP StorageWorks, and Hitachi Data Systems only.) Total amount of formatted storage space in the

storage pools associated with the storage system. For example, a RAID volume space has more capacity than a stripe-mirror volume space.

Formatted Disk Space with No Volumes (Total Formatted Space with No Volumes)
(TotalStorage Enterprise Storage Server, FAStT, HP StorageWorks, and Hitachi Data Systems only.) Amount of storage space available to allocate to volumes in the storage pools associated with the storage system.

Unformatted Disk Space (Undefined Disk Group or Array Site Capacity)
Raw disk space in ranks on which no volume space has been defined. For TotalStorage Enterprise Storage Server systems, this value is the total raw capacity of any unformatted disk groups. For IBM XIV Storage System, this value is the total available physical (hard) space that remains for pool creation. For FAStT, HP StorageWorks, and Hitachi Data Systems storage systems, this value is the total capacity of the disks that are not in a storage pool.

Cache Amount of RAM or other cache for a storage system.

Note: This column is blank for SAN Volume Controllers.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage capacity of all volumes that the storage system knows are assigned to a host computer. Note: The column is blank for SAN Volume Controller.

Unassigned Volume Space (Subsystem Unassigned LUN Capacity)
Total storage capacity of all volumes that the storage system knows are not assigned to a host computer. Note: The column is blank for SAN Volume Controller.

zOS Volume Space (zOS LUN Capacity)
Total storage capacity for all the volumes allocated to zOS. Data for this column is calculated for TotalStorage Enterprise Storage Server only. The value for this column is set to N/A for other storage systems.

Open System Volume Space (Open System LUN Capacity)
Total volume storage capacity (excluding the zOS volume space).

Physical Disk Space
Total amount of physical disk space discovered on the monitored storage system arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

Backend Volume Space
Total amount of storage system volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage system virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are back-end disks on a storage system virtualizer that is not monitored.

By Storage Subsystem Group

Subsystem Group Name
Name of a storage system monitoring group or storage system reporting group.

TPC Server
Name of a subordinate server to which the group is associated.

Disk Space (Disk Capacity)
Total raw (unformatted) disk capacity of the storage systems in a group.

Available Disk Space (Available Capacity)
Amount of storage space available on the disks associated with the storage systems in a group.

Number of Disks
Number of disks associated with the storage systems in a group.

Number of Volumes (Number of LUNs)
Number of volumes associated with the storage systems in a group.

Storage Subsystem Volume Space (LUN capacity)
Amount of storage space allocated to the volumes associated with the storage systems in a group.

Consumable Storage Subsystem Volume Space (Usable LUN Capacity)
Storage capacity of volumes in the created disk array (not including BCVs, VCM database, and Replication targets).

FlashCopy Target Storage Subsystem Volume Space (FlashCopy Target Capacity)
The total capacity of all FlashCopy target volumes in the storage systems in a group.

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Capacity of all volumes visible to a Storage Resource agent associated with the storage systems in a group.

Volume Space Not Correlated (LUN Capacity Not Visible to Data Manager)
Usable volume capacity minus assigned volume capacity on the storage systems in a group.

Unavailable Volume Space (Overhead)
Amount of capacity that is dedicated to redundancy in the storage systems in a group.

Formatted Space (Formatted Capacity)
(TotalStorage Enterprise Storage Server, FAStT, HP StorageWorks, and Hitachi Data Systems only.) Total amount of formatted storage space in the storage pools associated with the storage systems in a group.

Total Formatted Space with no Volumes (Total Formatted Space with No Volumes)
(TotalStorage Enterprise Storage Server, FAStT, HP StorageWorks, and Hitachi Data Systems only.) Amount of storage space available to allocate to volumes in the storage pools associated with the storage systems in a group.

Unformatted Disk Space (Total Overhead)
(TotalStorage Enterprise Storage Server only.) Sum of all the space in a disk group minus the formatted capacity on the storage systems in a group.

Cache Amount of RAM or other cache for the storage systems in a group.

Note: This column does not include values for SAN Volume Controllers.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage capacity of all volumes that the storage systems in a group know are assigned to a host computer.

Unassigned Volume Space (Subsystem Unassigned LUN Capacity)
Total storage capacity of all volumes that the storage systems in a group know are not assigned to a host computer.

zOS Volume Space (zOS LUN Capacity)
Total storage capacity for all the volumes on the storage systems in a group allocated to zOS.

Open System Volume Space (Open System LUN Capacity)
Total volume storage capacity (excluding the zOS volume capacity) of the storage systems in a group.

Physical Disk Space
Total amount of physical disk space discovered on the monitored storage system arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

Backend Volume Space
Total amount of storage system volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage system virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are back-end disks on a storage system virtualizer that is not monitored.

By Computer

Note: Two entries for the same host might appear in this report and the values for that host are counted twice in the **TOTAL** column. This scenario occurs under the following circumstances:

- You install two Storage Resource agents on the same host (using the long and short names of the host) that point to different Tivoli Storage Productivity Center servers. For example, you install an agent from serverA using the short name of a host (for example, hostC). You install another agent from serverB using the long name of the same host (for example, hostC.storage.chicago.com). You then set up the servers in a master and subordinate relationship. For example, serverA is the master and serverB is the subordinate. In this example, there are two entries in the report for the host: hostC and hostC.storage.chicago.com. The values for hostC are counted twice in the **TOTAL** column.
- A Storage Resource agent on a host (hostC) points to two different Tivoli Storage Productivity Center servers (serverA and serverB). The host has a dynamic IP address. You set up serverA as the master and serverB as the subordinate. Server serverA probes hostC, but then the IP address of hostC changes before serverB

probes the host. In this example, the report might display two entries for hostC: one entry with the IP address that was probed by serverA; the other entry with the changed IP address that was probed by serverB. The values for hostC are counted twice in the **TOTAL** column.

- A Storage Resource agent on a host (hostC) points to two different Tivoli Storage Productivity Center servers (serverA and serverB). The host has a name that resolves to hostC.storage.chicago.com on serverA and hostC.srm.storage.chicago.com on serverB. You then set up the servers in a master and subordinate relationship. For example, serverA is the master and serverB is the subordinate. In this example, there are two entries in the report for the host: hostC.storage.chicago.com and hostC.srm.storage.chicago.com. The values for hostC are counted twice in the **TOTAL** column.

Computer Name
Name of a computer.

Disk Space (Total Capacity)
Total storage capacity for a computer. If the computer is an IBM Tivoli Storage SAN File Systems client, then this value also includes the capacity of any volumes visible to this computer that are owned by the SAN File System.

Available Disk Space (Total Unallocated Space)
Amount of unused storage space on a computer (not in file systems seen by this operating system). If the computer is an IBM Tivoli Storage SAN File Systems client, then this value excludes the unallocated space on any volumes visible to it that are owned by the SAN File System.

Number of Disks
Number of disks on a computer.

Owned Disk Space (Owned Capacity)
Total storage capacity for a computer that is owned by that computer. If Data Manager can determine that a disk was configured for a different host, the capacity for that disk is counted in the "Owned Capacity".

If the computer is an IBM Tivoli Storage SAN File Systems client, then this value does not include the capacity of any volumes visible to it that are owned by the SAN File System. If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads the host-ID. If this host-ID does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

Other Available Disk Space (Other Unallocated Space)
Total unallocated space on disks owned by the computer. This column does not display the total unallocated disk space on disks not owned by the computer. If the computer is an IBM Tivoli Storage SAN File Systems client, then this value excludes the unallocated space on any volumes visible to it that are owned by the SAN File System.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)
Total disk capacity of all non-Fibre attached disks for a computer.

Fibre Channel Attached Disk Space (Total Other Attached Unallocated Space)
Total unallocated space of all non-Fibre attached disks for a computer.

Non-Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Capacity)
Total disk capacity of all fibre channel attached disks for a computer.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)
Total unallocated space of all fibre channel attached disks for a computer.

By Computer Group

Computer Group Name
Name of a computer monitoring group or computer reporting group.

Disk Space (Total Capacity)
Total storage capacity for the computers in a group. If the computer is an IBM Tivoli Storage SAN File Systems client, then this value also includes the capacity of any volumes visible to this computer that are owned by the SAN File System.

Available Disk Space (Total Unallocated Space)
Amount of unused storage space on the computers in a group.

Number of Disks
Number of disks on the computers in a group.

Owned Disk Space (Owned Capacity)
Total storage capacity for the computers in a group that is owned by those computers. If Data Manager can determine that a disk was configured for a different host, the capacity for that disk is counted in the "Owned Capacity".

If any computers in the group are IBM Tivoli Storage SAN File Systems clients, then this value does not include the capacity of any volumes visible to it that are owned by the SAN File System. If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads the host-ID. If this host-ID does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

Other Available Disk Space (Other Unallocated Space)
Total unallocated space on disks owned by the computers in a group. This column does not display the total unallocated disk space on disks not owned by computers. If any of the computers in the group are IBM Tivoli Storage SAN File Systems clients, then this value excludes the unallocated space on any volumes visible to it that are owned by the SAN File System.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)
Total disk capacity of all non-Fibre attached disks on the computers in a group.

Fibre Channel Attached Disk Space (Total Other Attached Unallocated Space)
Total unallocated space of all non-Fibre attached disks on the computers in a group.

Non-Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Capacity)
Total disk capacity of all fibre channel attached disks on the computers in a group.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)
Total unallocated space of all fibre channel attached disks on the computers in a group.

By Cluster

Cluster Name
Name of a cluster.

For Storwize V7000 Unified or IBM SONAS clusters, the storage system name is included with the cluster name in the Computer or Device column. The format of the column is `storagesystem_name: cluster_name`, where storagesystem_name is the name of the Storwize V7000 Unified or IBM SONAS device that the cluster belongs to. For example, tpcserver1.tpc.example.com: kq98n5d.acct displays storage system tpcserver1.tpc.example.com that cluster kq98n5d.acct belongs to.

Disk Space (Total Capacity)
Total storage capacity for a cluster.

Available Disk Space (Total Unallocated Space)
Amount of unused storage space on a cluster (not in file systems seen by this operating system).

Owned Disk Space (Owned Capacity)
Total storage capacity for a cluster that is owned by that cluster. If Data Manager can determine that a disk was configured for a different host, the capacity for that disk is counted in the "Owned Capacity".

If the computer is an IBM Tivoli Storage SAN File Systems client, then this value does not include the capacity of any volumes visible to it that are owned by the SAN File System. If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads the host-ID. If this host-ID does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

Other Available Disk Space (Other Unallocated Space)
Total unallocated space on disks owned by the cluster. This column does not display the total unallocated disk space on disks not owned by the cluster. If a computer in a cluster is an IBM Tivoli Storage SAN File Systems client, then this value excludes the unallocated space on any volumes visible to it that are owned by the SAN File System.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)
Total disk capacity of all non-Fibre attached disks in a cluster.

Fibre Channel Attached Disk Space (Total Other Attached Unallocated Space)
Total unallocated space of all non-Fibre attached disks in a cluster.

Non-Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Capacity)
Total disk capacity of all fibre channel attached disks in a cluster.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)
Total unallocated space of all fibre channel attached disks in a cluster.

By TPC Server (Computers)

TPC Server
Name of a subordinate server that is monitoring computer disks in your environment.

Disk Space (Total Capacity)
Total storage capacity of all the computer disks monitored by a subordinate server.

Available Disk Space (Total Unallocated Space)
Amount of unused storage space on all the computer disks monitored by a subordinate server.

Number of Disks
Number of computer disks monitored by a subordinate server.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)
Total disk capacity of all non-Fibre attached disks monitored by a subordinate server.

Fibre Channel Attached Disk Space (Total Other Attached Unallocated Space)
Total unallocated space of all non-Fibre attached disks monitored by a subordinate server.

Non-Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Capacity)
Total disk capacity of all fibre channel attached disks monitored by a subordinate server.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)
Total unallocated space of all fibre channel attached disks monitored by a subordinate server.

By TPC Server (Storage Subsystems)

TPC Server
Name of a subordinate server that is monitoring storage system disks in your environment.

Disk Space (Total Capacity)
Total raw (unformatted) disk capacity of the storage systems monitored by a subordinate server.

Available Disk Space (Total Unallocated Space)
Amount of storage space available on the disks associated with the storage systems monitored by a subordinate server.

Number of Disks
Number of storage system disks monitored by a subordinate server.

Physical Disk Space
Total amount of physical disk space discovered on the monitored storage system arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

Rollup Reports - TPC Server Storage Space - File System Space:

Use this report to view file system space information about the file systems that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report by using the following by clauses:

- By Computer: Shows file system capacity that is aggregated at the computer level and sorted by computer name.

- By Computer Group: Shows file system capacity that is aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By Cluster: Shows file system capacity that is aggregated at the cluster level and sorted by cluster name.
- By TPC Server: Shows file system capacity that is aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when you scroll through the other columns on the report.

By Computer

Computer Name
Name of the computer with a file system.

Total Capacity
Total amount (capacity) of storage space in the file systems on a computer. The value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. The extra storage occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**. The difference represents the space that is wasted because of the system overhead.

Total Used Space
Total amount of used storage space in the file systems on a computer. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead). Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column.

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

The value that is displayed in the TOTAL row for this column represents the difference between the TOTAL values in the **Total Capacity** and **Total Free Space** columns. If the calculated value for the TOTAL row is negative, the column displays 0.

The TOTAL row does not add the space that is used by VMWare computers. See the *VMWare capacity reports* topic in the *Planning* chapter of the *Tivoli Storage Productivity Center Installation and Configuration Guide* for information about how file system and disk capacity are calculated in this report for VMWare computers. For related information, see the *Planning* section of the *Tivoli Storage Productivity Center Information Center*.

The value that is displayed in the TOTAL row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent is installed.
- File systems on discovered virtual machine servers where the Storage Resource agent is installed. The hypervisor for the virtual machine

servers is not discovered yet and the total used space exceeds the capacity of the total file system capacity.

Total Free Space
Total amount of unused storage space in the file systems on a computer.

Number of FS/LV
Number of file systems or logical volumes on a computer.

By Computer Group

Computer Group
Name of a computer monitoring group or computer reporting group.

Total Capacity
Total amount (capacity) of storage space on the file systems in a group. The value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. The extra storage occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**. The difference represents the space that is wasted because of the system overhead.

Total Used Space
Total amount of used storage space on the file systems in a group. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space on the file systems in a group.

Number of FS/LV
Number of file systems or logical volumes in a group.

By Cluster

Cluster
Name of a cluster.

Total Capacity
Total storage capacity of the file systems in a cluster. The value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. The extra storage occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**. The difference represents the space that is wasted because of the system overhead.

Total Used Space
Total amount of used storage space on the file systems in a cluster. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space on the file systems in a cluster.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring file systems in your environment.

Total Capacity
Total storage capacity of all the file systems that are monitored by a subordinate server.

Total Used Space
Total amount of used storage space on the file systems that are monitored by a subordinate server. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space on the file systems that are monitored by a subordinate server.

Total Number of FS/LV
Total number of file systems or logical volumes that are monitored by a subordinate server.

Rollup Reports - TPC Server Storage Space - Consumed File System Space:

Use this report to view information about the used storage space on file systems that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report by using the following by clauses.

- By Computer: Shows file system used space that is aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows file system used space that is aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By Cluster: Shows file system used space that is aggregated at the cluster level and sorted by cluster name.
- By TPC Server: Shows file system used space that is aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when you scroll through the other columns on the report.

By Computer

Computer Name
Name of the computer with a file system.

Total Used Space
Total amount of used storage space in the file systems on a computer. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead). Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

The value that displays in the TOTAL row for this column represents the difference between the TOTAL values in the **Total Capacity** and **Total Free Space** columns. If the calculated value for the TOTAL row is negative, the column displays 0.

The TOTAL row does not add the used space that is used by VMWare computers. See the *VMWare capacity reports* topic in the *Planning* chapter of the *Tivoli Storage Productivity Center Installation and Configuration Guide* for information about how file system and disk capacity are calculated in this report for VMWare computers. For related information, see the *Planning* section of the *Tivoli Storage Productivity Center Information Center*.

The value that displays in the TOTAL row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent is installed.
- File systems on discovered virtual machine servers where the Storage Resource agent is installed, but the hypervisor for the virtual machine servers is not discovered yet and the total used space exceeds the capacity of the total file system capacity.

Total Free Space
Total amount of unused storage space in the file systems on a computer.

Total Capacity
Total amount (capacity) of storage space in the file systems on a computer. The value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. The extra space occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**. The difference represents the space that is wasted because of the system overhead.

Number of FS/LV
Number of file systems or logical volumes on a computer.

By Computer Group

Computer Group
Name of a computer monitoring group or computer reporting group.

Total Capacity
Total amount (capacity) of storage space on the file systems in a group. The value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. The extra storage occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**. The difference represents the space that is wasted because of the system overhead.

Total Used Space
Total amount of used storage space on the file systems in a group. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space on the file systems in a group.

Number of FS/LV
Number of file systems or logical volumes in a computer group.

By Cluster

Cluster
Name of a cluster.

Total Used Space
Total amount of used storage space on the file systems in a cluster. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space on the file systems in a cluster.

Total Capacity
Total storage capacity of the file systems in a cluster. The value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. This extra space occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**. The difference represents the space that is wasted because of the system overhead.

Number of FS/LV
Number of file systems or logical volumes in a cluster.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring file systems in your environment.

Total Used Space
Total amount of used storage space on the file systems that are monitored by a subordinate server. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space on the file systems that are monitored by a subordinate server.

Total Capacity
Total storage capacity of all the file systems that are monitored by a subordinate server.

Total Number of FS/LV
Number of file systems or logical volumes that are monitored by a subordinate server.

Rollup Reports - TPC Server Storage Space - Available File System Space:

Use this report to view information about the available storage space on file systems that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report by using the following by-clauses:

- By Computer: Shows file system free space that is aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows file system free space that is aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By Cluster: Shows file system free space that is aggregated at the cluster level and sorted by cluster name.
- By IBM Tivoli Storage Productivity Center Server: Shows file system free space that is aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when you scroll through the other columns on the report.

By Computer

Computer Name
Name of the computer with a file system.

Total Free Space
Total amount of unused storage space in the file systems on a computer.

Total Used Space
Total amount of used storage space in the file systems on a computer. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead). Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

The value that is displayed in the TOTAL row for this column represents the difference between the TOTAL values in the **Total Capacity** and **Total Free Space** columns. If the calculated value for the TOTAL row is negative, the column displays 0.

The TOTAL row does not add the space that is used by VMWare computers. See the *VMWare capacity reports* topic in the *Planning* chapter of the *Tivoli Storage Productivity Center Installation and Configuration Guide* for information about how file system and disk capacity are calculated in this report for VMWare computers. For related information, see the *Planning* section of the *Tivoli Storage Productivity Center Information Center*.

The value that displays in the TOTAL row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent is installed.
- File systems on discovered virtual machine servers where the Storage Resource agent is installed, but the hypervisor for the virtual machine servers is not discovered yet and the total used space exceeds the capacity of the total file system capacity.

Total Capacity
Total amount (capacity) of storage space in the file systems on a computer. The value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. This extra space occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**. The difference represents the space that is wasted because of the system overhead.

Number of FS/LV
Number of file systems or logical volumes on a computer.

By Computer Group

Computer Group
Name of a computer monitoring group or computer reporting group.

Total Free Space
Total amount of unused storage space on the file systems in a group.

Total Used Space
Total amount of used storage space on the file systems in a group. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Total Capacity
Total amount (capacity) of storage space on the file systems in a group. The value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. This extra space occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**. The difference represents the space that is wasted because of the system overhead.

Number of FS/LV
Number of file systems or logical volumes in a computer group.

By Cluster

Cluster
Name of a cluster.

Total Free Space
Total amount of unused storage space on the file systems in a cluster.

Total Used Space
Total amount of used storage space on the file systems in a cluster. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Total Capacity
Total storage capacity of the file systems in a cluster. The value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. This extra space occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents and by the **df** command) does not include some space that is reserved by the operating system (overhead). Therefore, the value in the **Total Capacity**

field is greater than or equal to the value of **Total Used Space** + **Total Free Space**. The difference represents the space that is wasted because of the system overhead.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring file systems in your environment.

Total Free Space
Total amount of unused storage space on the file systems that are monitored by a subordinate server.

Total Used Space
Total amount of used storage space on the file systems that are monitored by a subordinate server. On some UNIX file systems, this value does not include storage space that is reserved by the operating system (overhead).

Total Capacity
Total storage capacity of all the file systems that are monitored by a subordinate server.

Number of FS/LV
Number of file systems or logical volumes that are monitored by a subordinate server.

Rollup Reports - Database Capacity

Use these reports to view storage capacity, used space, and free space at the computer, computer group, database, instance, and IBM Tivoli Storage Productivity Center server level for RDBMSs that are monitored by subordinate servers.

For a specific instance, computer, group of computers, or database that is monitored by a subordinate server, you can use these rollup reports to view:

- how much storage you have for an RDBMS and the objects it contains
- how much storage is being used for an RDBMS and the objects it contains
- what percentage of RDBMS storage is free for expansion

The information shown in these reports is gathered by the master server when it runs Tivoli Storage Productivity Center server probes against associated subordinate servers. If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

Rollup Reports - Database Capacity - All DBMSs:

Use this report to view storage capacity information for the DBMSs (including Oracle, Sybase, SQL Server, and UDB/DB2) that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Instance: Shows database capacity aggregated at the computer level and sorted by instance name.
- By Computer: Shows database capacity aggregated at the computer level and sorted by computer name.

- By Computer Group: Shows database capacity aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By TPC Server: Shows database capacity aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Instance

Instance
Name of an instance.

RDBMS Type
RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

Computer Name
Name of the computer on which an Instance resides.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Tablespace Count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

By Computer

Computer Name
Name of the computer on which an Instance resides.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Tablespace Count
Number of Tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

By Computer Group

Computer Group
Name of a computer monitoring group or reporting group.

Total Size
Amount of storage space allocated to instances in a group.

File Capacity
Amount of storage space allocated to the data files associated with instances in a group.

File Free Space
Amount of storage free space available on the data files associated with instances in a group.

Tablespace Count
Number of databases or tablespaces associated with instances in a group.

File Count
Number of data files associated with instances in a group.

Log File Count
Number of log files associated with instances in a group.

TPC Server
Name of a subordinate server that is monitoring the computer group.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring DBMSs in your environment.

Total Size
Amount of storage space allocated to instances monitored by a subordinate server.

Total File Capacity
Amount of storage space allocated to the data files associated with instances monitored by a subordinate server.

File Free Space
Amount of storage free space available on the data files associated with instances monitored by a subordinate server.

Tablespace Count
Number of tablespaces associated with instances monitored by a subordinate server.

File Count
Number of data files associated with instances monitored by a subordinate server.

Log File Count
Number of log files associated with instances monitored by a subordinate server.

Rollup Reports - Database Capacity - Oracle, Sybase, SQL Server, and UDB:

Use this report to view storage capacity information for the Oracle, Sybase, SQL Server, and UDB/DB2 databases that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Database: Shows database capacity aggregated at the computer level and sorted by instance name.
- By Computer: Shows database capacity aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows database capacity aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By TPC Server: Shows database capacity aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Database

Database Instance
Name of a database instance.

Computer Name
Name of the computer on which an instance resides.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Tablespace Count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

By Computer

Computer Name
Name of the computer on which an instance resides.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Tablespace Count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

By Computer Group

Computer Group
Name of a computer monitoring group or reporting group.

Total Size
Amount of storage space allocated to instances in a group.

File Capacity
Amount of storage space allocated to the data files associated with instances in a group.

File Free Space
Amount of storage free space available on the data files associated with instances in a group.

Tablespace Count
Number of tablespaces associated with instances in a group.

File Count
Number of data files associated with instances in a group.

Log File Count
Number of log files associated with instances in a group.

TPC Server
Name of a subordinate server that is monitoring the group.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring DBMSs in your environment.

Total Size
Amount of storage space allocated to instances monitored by a subordinate server.

Total File Capacity
Amount of storage space allocated to the data files associated with instances monitored by a subordinate server.

File Free Space
Amount of storage free space available on the data files associated with instances monitored by a subordinate server.

Tablespace Count
Number of tablespaces associated with instances monitored by a subordinate server.

File Count
Number of data files associated with instances monitored by a subordinate server.

Log File Count
Number of log files associated with instances monitored by a subordinate server.

Data source reports - overview

Use data source reports to view detailed information about the agents from which IBM Tivoli Storage Productivity Center collects information about storage resources. Data sources include CIMOM agents, Storage Resource agents, Fabric agents, VMWare, and subordinate Tivoli Storage Productivity Center servers.

Data source reports - CIMOM agents

Use these reports to view detailed information about the CIMOM agents that have been associated with IBM Tivoli Storage Productivity Center.

Related Topics

- Selecting resources upon which to report
- Filtering the resources that appear in a report
- Generating a report
- Drilling down a report row (By Managed Device report only)
- Navigating tabs

Data source report - By CIMOM Agent

Use this report to view detailed information about the CIMOM agents that are used as data sources for IBM Tivoli Storage Productivity Center.

Each CIMOM that has been associated with IBM Tivoli Storage Productivity Center appears as its own row in this report.

Service URL
The service URL of the CIMOM containing the IP address of the CIMOM, the port on which the CIMOM is listening, and the protocol used for communication. This URL has a protocol [http | https], an IP address (IPv4 or IPv6), and a port number. For example:

- **IPv4:** https://127.0.0.1:5989
- **IPv6:** https://[2001:DB8::1234:0000:0000:5678:ABCD]:5989

Display Name
The name of the CIMOM as specified by the CIMOM provider that will appear in the Tivoli Storage Productivity Center user interface.

Description
The optional description that was entered on the Add CIMOM panel.

Managed Device Category
The type of device that is managed by a CIMOM (for example, storage subsystem or switch). The value **NA** appears in this column if you have not run a CIMOM discovery against a CIMOM or a CIMOM does not manage any devices.

Connection Status
The status of this CIMOM with respect to Tivoli Storage Productivity Center. Possible values are: SUCCESS, UNCONFIGURED, UNKNOWN,

INVALID_NAMESPACE, TIMEOUT, REFUSED, LOGIN_FAILED, SSL_HANDSHAKE_ERROR, SSL_REGISTRATION_INVALID, CIMCLIENT_ERROR

Status Timestamp
The date and time when the **Connection Status** information was last collected.

Number of Managed Devices
The number of devices managed by a CIMOM.

Interoperability Namespace
The interoperability namespace of a CIMOM. This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information. For a NetApp CIMOM, enter /interop.

Username
The CIMOM user name used for authentication.

Truststore Location
The location (path on this computer) of a certificate file for certificate based authentication in the https protocol.

User Interface Description
The name of the Human Interface Service (if any) supported by this CIMOM.

Software Level
The software version level of the CIMOM agent.

Protocol Version
The version of the cim-xml protocol.

Authentication Mechanism
The authentication mechanism supported by the CIMOM. This column can contain the following values: **Unknown**, **None**, **Other**, **Basic**, **Digest**.

Alias The alias of the CIMOM.

Protocol
The protocol used to communicate with the CIMOM. Possible values are **http** and **https**.

SLP Attributes
The standard set of attributes for this CIMOM. The attributes are retrieved via SLP.

Data source report - By Managed Device

Use this report to view information about the devices managed by CIMOMs in your environment. Only devices that are managed by the CIMOMs used as data sources for IBM Tivoli Storage Productivity Center are displayed.

You must run a CIMOM discovery job before you can view the devices managed by CIMOMs in this report. If a device is managed by multiple CIMOMs, the report displays a separate row for each of those CIMOMs. Therefore, the number of rows for a device in the report is equal to the number of CIMOMs managing that device.

Magnifying glass icon
Click this button to view a more detailed report about a device.

Device Name
The display name (label) of the managed storage subsystem or switch.

Device Category
Indicates whether a device that is managed by a CIMOM is a storage subsystem or switch.

Device Type
Indicates the type of a managed device based on its category. For example, a device that is a storage subsystem might have a value of **DS8000** in this column.

Device Status
The consolidated status of the managed device.

Device Operational Status
The operational status of a storage subsystem.

Device Last Probe Time
The date and time when a probe was last run on a managed device.

Device Vendor
The manufacturer of the managed device.

Device Model
The model number of the managed device.

Device Serial Number
The serial number of the managed device.

Device Microcode Level
The microcode or firmware level of the managed device.

Device IP Address
The IP address of the managed device.

Device Element Manager
The URL or the fully-qualified name of the executable file that launches the element manager.

CIMOM Service URL
The service URL of the CIMOM. This URL includes the IP address or host name of the CIMOM, the port on which the CIMOM is listening, and the protocol used for communication (http|https). For example: https://127.0.0.1:9551.

CIMOM Display Name
The name of the CIMOM as specified by the CIMOM provider that will appear in the Tivoli Storage Productivity Center user interface.

CIMOM Description
The optional description that was entered on the Add CIMOM panel.

CIMOM Connection Status
The status of this CIMOM with respect to Tivoli Storage Productivity Center. Possible values are: SUCCESS, UNCONFIGURED, UNKNOWN, INVALID_NAMESPACE, TIMEOUT, REFUSED, LOGIN_FAILED, SSL_HANDSHAKE_ERROR, SSL_REGISTRATION_INVALID, CIMCLIENT_ERROR

CIMOM Status Timestamp
The date and time when the **Connection Status** information was last collected.

CIMOM Interoperability Namespace
The interoperability namespace of a CIMOM. This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information. For a NetApp CIMOM, enter /interop.

CIMOM Username
The CIMOM user name used for authentication.

CIMOM Trustore Location
The location (path on this computer) of a certificate file for certificate based authentication in the https protocol.

CIMOM User Interface Description
The name of the Human Interface Service (if any) supported by this CIMOM.

CIMOM Software Level
The software version level of the CIMOM agent.

CIMOM Protocol Version
The version of the cim-xml protocol.

CIMOM Authentication Mechanism
The authentication mechanism supported by the CIMOM. This column can contain the following values: **Unknown**, **None**, **Other**, **Basic**, **Digest**.

CIMOM Alias
The alias of the CIMOM.

CIMOM Protocol
The protocol used to communicate with the CIMOM. Possible values are **http** and **https**.

CIMOM SLP Attributes
The standard set of attributes for this CIMOM. The attributes are retrieved via SLP.

Tivoli Storage Productivity Center Monitoring Agent reports

After you have installed and configured the Tivoli Storage Productivity Center Monitoring Agent, you can view reports and information from Tivoli Storage Productivity Center.

There are four predefined workspaces which are shipped with Tivoli Storage Productivity Center:

TPC Server Overview
Provides an over view of the key parameters and infrastructure metrics about the Data server and Device server managed entities.

Managed SAN Overview
Provides an overview about several metrics concerning the managed SAN environment.

Managed SAN Health
Monitors the health and availability of the SAN attached devices managed by Tivoli Storage Productivity Center.

TPC Data Collection Overview
Displays the status of the latest probe jobs and performance monitors.

Managed SAN Overview workspace

The Managed SAN Overview workspace provides an overview about several metrics about the managed SAN environment. You can access this workspace by clicking **Managed SAN Overview** in the navigation tree of the Navigator window.

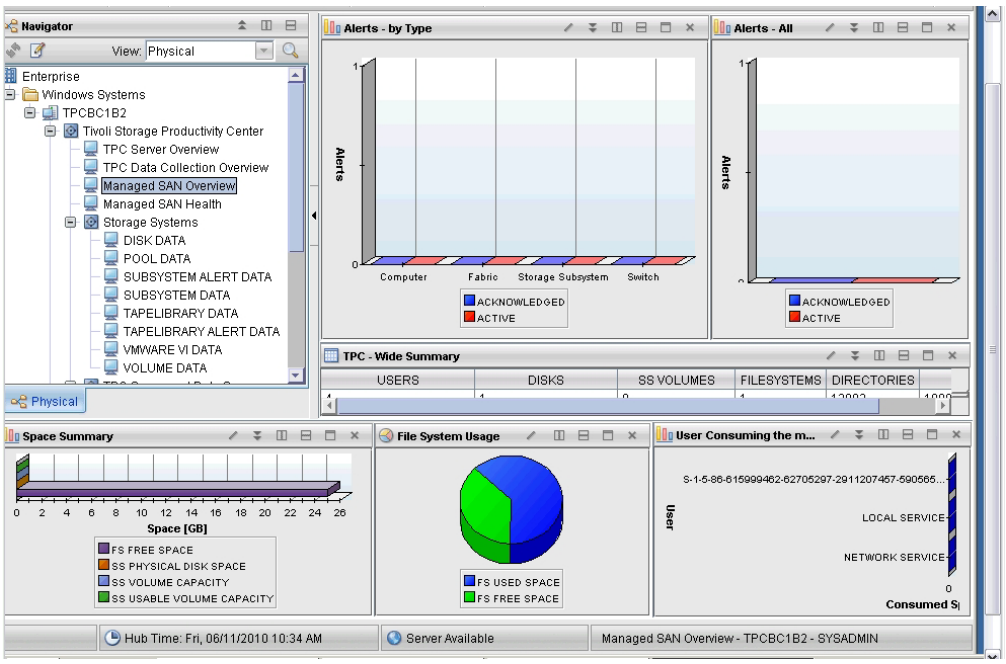

Figure 42. Managed SAN overview report

You can see the following views:

Alerts - by Type
Displays a bar chart showing the number of all computers, switches, storage systems, and fabric alerts including their current state (acknowledged or active).

Alerts - All
Displays a bar chart showing the number of all alerts of all devices including their current state (acknowledged or active).

TPC-Wide Summary
Displays a table showing the total amount of users, disks, storage system volumes, file systems, files, and folders.

Space Summary
Displays a bar chart showing the total amount of file system capacity, physical disk space of all storage systems, total amount of the available volume capacity, and total amount of usable volume capacity in the managed storage environment.

User Consuming the most Space
Displays a chart with the total amount of space allocated by single users. Only the ten users using the most space is shown.

File System Usage
Displays a pie chart showing the total amount of free and used file system space in the managed storage environment.

Managed SAN Health workspace

The Managed SAN Health workspace provides an overview about the objects in the managed SAN environment. You can access this workspace by clicking **Managed SAN Health** in the navigation tree of the Navigator window.

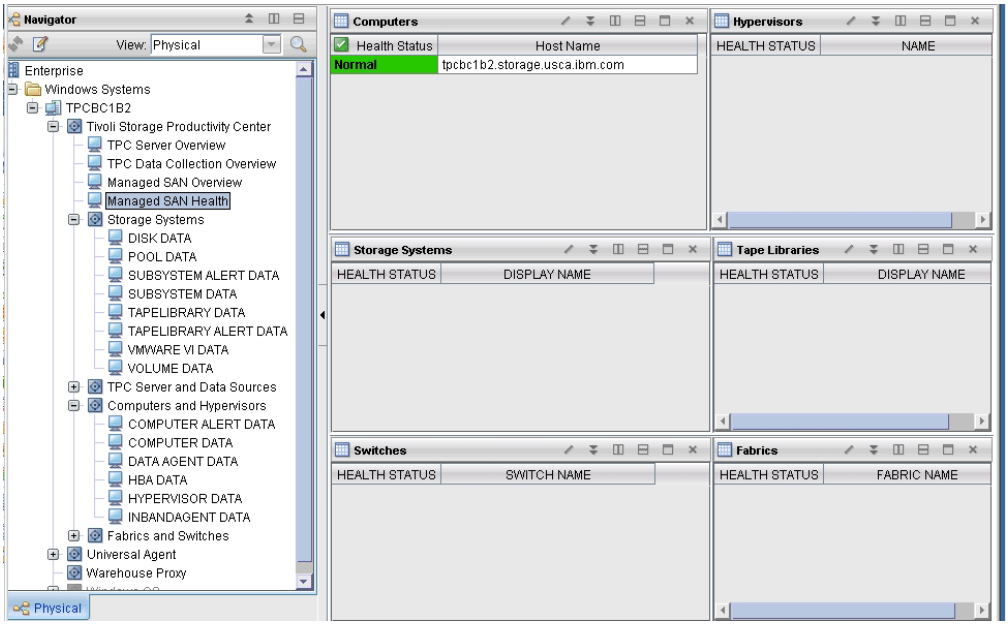

Figure 43. Managed SAN Health report

This workspace shows the following information:

Computers
Displays a list of all managed computers including their health status.

Hypervisors
Displays a list of all managed VMware hypervisors including their health status.

Storage Systems
Displays a list of all managed storage systems including their health status.

Fabrics
Displays a list of all managed fabrics including their health status.

Switches
Displays a list of all managed switches including their health status.

TPC Server Overview workspace

The TPC Server Overview workspace provides an overview of several metrics about the Tivoli Storage Productivity Center server. You can access this workspace by clicking **TPC Server Overview** in the navigation tree of the Navigator window.

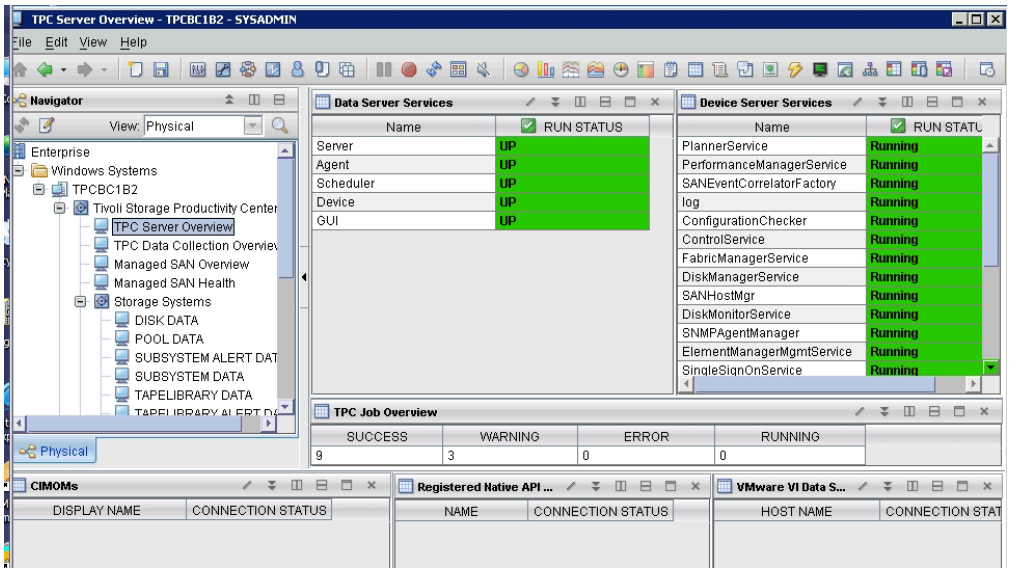

Figure 44. TPC Server overview report

This workspace provides the following views:

Data Server Services
Displays a table showing all the Data server services (Server, Agent, Scheduler, Device, and GUI) including their run status.

Device Server Services
Displays a table showing all Device server services including their run status.

TPC Job Overview
Displays a table showing statistics about the total amount of jobs and their corresponding status.

CIMOMs
Displays a table showing all the managed CIMOMs including their connection status.

Registered Native API Devices
Displays a table showing all managed native interface devices including their corresponding connection status.

VMware VI Data Sources
Displays a report showing all attached VMware VI Data Sources and their corresponding connection status.

TPC Data Collection Overview workspace

The TPC Data Collection Overview workspace provides the status of the latest probe jobs and performance monitor jobs. You can access this workspace by clicking **TPC Data Collection Overview** in the navigation tree of the Navigator window.

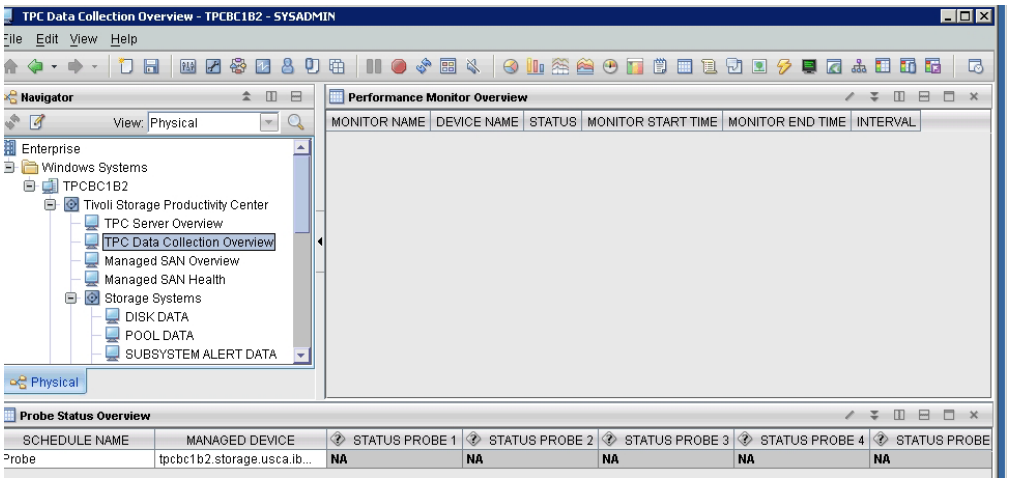

Figure 45. TPC Server overview report

This workspace provides the following information:

Probe status overview
Displays a table showing all probe jobs and their status with the last five runs.

Performance monitor overview
Displays a table showing all performance monitor jobs and their corresponding status.

Detailed Tivoli Storage Productivity Center reports

There are several reports which provide detailed information about the monitored entities. You can access these reports by clicking the appropriate entity in the navigation tree of the Navigator window. These entities are:

Storage Systems
Under **Storage Systems**, you can view these detailed reports:

- DISK DATA
- POOL DATA
- SUBSYSTEM ALERT DATA
- SUBSYSTEM DATA
- VMWARE VI DATA
- VOLUME DATA

TPC Server and Data Sources
Under TPC Server and Data Sources, you can view these detailed reports:

- CIMOM DATA
- DATASERVER API
- DEVICESERVER API
- JOB STATUS DATA
- REGISTERED NATIVE API DATA
- OTHER DEVICE DATA
- PERFORMANCE MONITOR DATA
- PROBE STATUS DATA
- SCHEDULED JOB DATA

- TPC SERVER DATA

Computers and Hypervisors

Under **Computers and Hypervisors**, you can view these detailed reports:

- COMPUTER ALERT DATA
- COMPUTER DATA
- HBA DATA
- HYPERVISOR DATA

Fabrics and Switches

Under **Fabrics and Switches**, you can view these detailed reports:

- FABRIC DATA
- FABRIC ALERT DATA
- OUTBANDAGENT DATA
- SWITCH DATA
- SWITCH ALERT DATA

See the Tivoli Storage Productivity Center online help for information about these entities.

Sharing Tivoli Storage Productivity Center alerts

Alerts can be shared with external applications such as VMware and viewed by administrators. You can share Tivoli Storage Productivity Center alerts that are mapped against VMware alerts.

To be available to VMware vSphere, a Tivoli Storage Productivity Center alert must be enabled for external reporting.

In the Alert Definition dialog in the stand-alone GUI, select the **Share with external applications** check box for each of the alerts from the following table that you want to share.

This table correlates Tivoli Storage Productivity Center alert names with the alert messages in vCenter:

Table 127. Tivoli Storage Productivity Center alerts in vCenter

Alert	vCenter name	Details
Export State Change	`Alarm.ExportStateChanged`	The export state {name} changed from {old_value} to {new_value}.
Export Is Missing	`Alarm.ExportIsMissing`	Export {name} is missing
Export Was Rediscovered	`Alarm.ExportWasRediscovered`	Export {name} was rediscovered.
Volume Status Change Offline	`Alarm.VolumeStatusChangeOffline`	Volume {name} is now offline.
Volume Status Change Online	`Alarm.VolumeStatusChangeOnline`	Volume {name} is now online.
Volume Not Found	`Alarm.VolumeIsMissing`	Volume {name} is missing.
Pool Shortfall Too High	`Alarm.ThinProvisionedLUNCapacityExceeded`	Thin provisioned LUN {name} capacity threshold exceeded.
Volume Capacity Change	`Event.VolumeCapacityChanged`	The volume capacity changed from {old_value} to {new_value} for volume {name}.

Note: Tivoli Storage Productivity Center VASA provider currently supports the sharing of alerts only with VMware vCenter.

Tivoli Storage Productivity Center can share only alerts that relate to volumes or file systems that are mapped to ESX hosts connected to vCenter.

Viewing Tivoli Storage Productivity Center VASA provider reports

Use the Tivoli Storage Productivity Center VASA provider to view Tivoli Storage Productivity Center data in vCenter reports.

When a Tivoli Storage Productivity Center VASA provider is online and inventory details are loaded, you can view Tivoli Storage Productivity Center data in vCenter reports. You can use the vSphere Client or the vSphere Web Client.

1. **vSphere Client**: Under **Home**, click **Hosts and Clusters**. Select **ESX** and click the **Storage Views** tab.
2. **vSphere Web Client**: Under **Home**, click **Hosts and Clusters**. Select **ESX** and in the **Monitor** tab, click **Storage Reports**.
3. From the **Report On** list, select **Show all SCSI Volumes (LUNs)**, **Show all Datastores**, or **Show all NAS Mounts**.

vCenter periodically updates storage data from the Tivoli Storage Productivity Center VASA provider.

To view current storage data in storage views and reports on vCenter, trigger the most recent vCenter Server updates from the Tivoli Storage Productivity Center VASA provider:

1. **vSphere Client**: Under **Home**, click **Storage Providers**.
2. Select the provider and click **Sync**.
3. Allow some minutes for the sync to take effect. Click **Refresh All**. Verify that the **Last Sync Time** value for the VASA provider is updated to the latest time stamp.

or

1. **vSphere Web Client**: Click **Hosts and Clusters**.
2. In the **Manage** tab, click **Storage Providers**.
3. Select the provider and click the **Rescan** icon.
4. Allow some minutes for the sync to take effect. Refresh the page. Verify that the **Last Rescan Time** value for the VASA provider is updated to the latest time stamp.

Now update the storage views and reports on vCenter:

1. Under **Home**, click **Hosts and Clusters**. Select **ESX** and click the **Storage Views** tab.
2. Click the **Update...** link to update the storage views.

or

1. Under **Home**, click **Hosts and Clusters**. Select **ESX** and in the **Monitor** tab, click **Storage Reports**.
2. Click the **Rescan vCenter Server for new inventory and storage devices** icon to update the storage reports.

Reporting in a vSphere environment

You can use the vSphere Web Client extension for Tivoli Storage Productivity Center to view reports on your virtual environment and storage devices.

You can view reports that are customized to use information from Tivoli Storage Productivity Center. The reports include fabric connections, storage mapping information, and performance metrics for storage systems.

To view Tivoli Storage Productivity Center storage information in block and file storage reports, you must register Tivoli Storage Productivity Center as a VASA provider.

For VASA provider-specific reports, ensure that the storage systems from which hypervisors have volumes or assigned shares are managed by Tivoli Storage Productivity Center.

For **Fabric Connections**, ensure that both the hypervisor and the fabric that is part of are managed by Tivoli Storage Productivity Center.

For **Storage Mappings**, ensure that both the hypervisor and the storage system are managed by Tivoli Storage Productivity Center and that there is an SRA deployed on the virtual machine.

For **Storage System Metrics**, ensure that both the hypervisor and the storage system are managed by Tivoli Storage Productivity Center and that performance is running on the storage system.

Use the following vSphere reports to view information about your virtual resources and the back-end storage systems:

Fabric Connections
This storage report displays fabric information that includes zone and switch details in the vSphere Web Client.

Storage Mapping
This storage mapping report displays end-to-end mappings between back-end storage that is monitored by Tivoli Storage Productivity Center and the virtual resources that are monitored by vSphere.

Storage System Metrics
This performance report displays performance metrics that include the total I/O rate and response time for the back-end storage systems that are monitored by Tivoli Storage Productivity Center, and that performance is running on the storage system.

SCSI Volumes (LUNs)
This volume report displays block storage information that is provided by Tivoli Storage Productivity Center. The information includes the following details:

- Space that is committed to the volume
- Thin provisioning status
- System capability
- Storage array name
- Volume identifier on the array
- Namespace of the VASA provider

Datastores
The Datastores report includes file system information that is provided by Tivoli Storage Productivity Center, including the system capability and the namespace of the VASA provider.

NAS Mounts
The NAS Mounts report includes file system information that is provided by Tivoli Storage Productivity Center, including the system capability and the namespace of the VASA provider.

Viewing vSphere storage mapping to Tivoli Storage Productivity Center

Use the vSphere Web Client extension to view the end-to-end mappings between back-end storage that is monitored by Tivoli Storage Productivity Center and virtual resources that are monitored by vSphere.

Start the vSphere Web Client, and log in to the vCenter Server system.

1. From the vSphere Web Client Inventories view, select the appropriate inventory object. For example, select a virtual machine.
2. Click the **Monitor** tab to view the monitoring options for the object.
3. Click the **Storage Mapping** tab to view a table that contains the mapping information from the virtual resource to Tivoli Storage Productivity Center.
4. Optional: To change the column that determines the sort order of resources, click the name of that column. For example, click the **Virtual Disk Size** column to sort the list of virtual disks according to size.

   Tip: To view descriptions of the columns in the **Storage Mapping** table, click the **Help** icon beside the **Hypervisor Host Name**.
5. Optional: To retrieve new data from Tivoli Storage Productivity Center, click the **Refresh** icon beside the **Guest OS Host Name**. The **Last updated** information in the toolbar is also updated.
6. Optional: To export storage mapping data to a CSV file or to copy the data to the clipboard, click the **Export** icon or select the relevant option from the **Export menu** in the status bar.

vSphere storage mapping to Tivoli Storage Productivity Center

You can use the vSphere Web Client extension to view the end-to-end mappings between back-end storage that is monitored by Tivoli Storage Productivity Center and virtual resources that are monitored by vSphere.

In the vSphere Web Client, use the **Monitor** tab to access the **Storage Mapping** tab for a virtual resource. You can view the attributes of each virtual disk, including the information that is provided by Tivoli Storage Productivity Center.

Use the **Storage Mapping** tab to help you to understand the relationships between an inventory object and all the virtual and physical storage resources that are related to this object. You can view details of storage usage inside the Guest Operating System (OS) and the back-end storage details such as volume, pool, and storage system information.

The Storage attribute column in the following table displays all of the attributes that are displayed in the **Storage Mapping** tab for a virtual disk. The Information source column shows whether the attribute information is from Tivoli Storage Productivity Center or vSphere.

Table 128. vSphere Storage mapping to Tivoli Storage Productivity Center

Storage attribute	Information source
Guest OS Host Name	vSphere
Virtual Disk	vSphere
Virtual Disk Size	vSphere
Virtual Disk Maximum Size	vSphere
Thin Provisioned	vSphere
Raw Device Mapping	vSphere
File	vSphere
Mount Point on Guest OS	Tivoli Storage Productivity Center This value is not available and "--" is displayed in any of the following situations: • The Guest OS on the virtual machine is not monitored by Tivoli Storage Productivity Center using a Storage Resource agent. • The ESX Server is not monitored by Tivoli Storage Productivity Center. • There is no file system on the virtual disk.
Used Space on Guest OS	Tivoli Storage Productivity Center This value is not available and "--" is displayed in any of the following situations: • The Guest OS on the virtual machine is not monitored by Tivoli Storage Productivity Center using a Storage Resource agent. • The ESX Server is not monitored by Tivoli Storage Productivity Center. • There is no file system on the virtual disk.
Disk on Guest OS	Tivoli Storage Productivity Center This value is not available and "--" is displayed in any of the following situations: • The Guest OS on the virtual machine is not monitored by Tivoli Storage Productivity Center using a Storage Resource agent. • The ESX Server is not monitored by Tivoli Storage Productivity Center. • There is no file system on the virtual disk.
Disk on Hypervisor	vSphere

Table 128. vSphere Storage mapping to Tivoli Storage Productivity Center (continued)

Storage attribute	Information source
Storage System	Tivoli Storage Productivity Center This value is not available and "--" is displayed if the storage system or ESX Server is not monitored by Tivoli Storage Productivity Center. This value is not applicable and "N/A" is displayed if the back-end storage is a NAS device.
Storage Pool	Tivoli Storage Productivity Center This value is not available and "--" is displayed if the storage system or ESX Server is not monitored by Tivoli Storage Productivity Center. This value is not applicable and "N/A" is displayed if the back-end storage is a NAS device.
Storage Volume	Tivoli Storage Productivity Center This value is not available and "--" is displayed if the storage system or ESX Server is not monitored by Tivoli Storage Productivity Center. This value is not applicable and "N/A" is displayed if the back-end storage is a NAS device.
NAS Server	vSphere This value is not available and "--" is displayed if the NAS server or ESX Server is not monitored by Tivoli Storage Productivity Center. This value is not applicable and "N/A" is displayed if the back-end storage is a storage system.
NAS Share Path	vSphere This value is not available and "--" is displayed if the NAS server or ESX Server is not monitored by Tivoli Storage Productivity Center. This value is not applicable and "N/A" is displayed if the back-end storage is a storage system.

Related tasks:

"Viewing vSphere storage mapping to Tivoli Storage Productivity Center" on page 1644

Use the vSphere Web Client extension to view the end-to-end mappings between back-end storage that is monitored by Tivoli Storage Productivity Center and virtual resources that are monitored by vSphere.

CSV file exports

You can use the export actions to export the information from the reports tables in the vSphere Web Client extension for Tivoli Storage Productivity Center to a CSV file. You can also copy the information to the clipboard.

Columns that are hidden on a page are included in the generated file. Depending on your web browser, you can determine the name and location of the file. The default file name is in the format `export_yyyy_mm_dd_hh_mm_ss.csv`.

To export information, click the **Export** icon or select the relevant option from the **Export menu** in the status bar. If no resources are selected, the information for all resources is exported.

Viewing storage system metrics in vSphere Web Client extension

Use the vSphere Web Client extension to view the key performance data for virtual resources that are monitored by vSphere. The data is retrieved from the back-end storage systems that are monitored by IBM Tivoli Storage Productivity Center.

Start the vSphere Web Client, and log in to the vCenter Server system.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters** and browse to the host for which you want to view performance data.
2. Click **Monitor** > **Performance** > **Storage System Metrics** to navigate to the performance data that was collected by Tivoli Storage Productivity Center. The performance information is displayed in the following formats:

   Performance Chart
   The chart shows a visual representation of how storage resource performance trends over time. Each line on the chart represents a metric and a resource. If the corresponding volumes have the same metric value, lines are overlaid on the chart.

   Performance Chart Legend
   The Performance Chart Legend shows performance information that is formatted into rows and columns.
3. To show the performance chart for a volume, in the **Performance Chart Legend**, click the volume. To select multiple volumes, press Ctrl and click the volumes. To select a series of volumes, select the first volume, and then press Shift and click the last volume. You can also press Shift and the Up/Down Arrow. If you select more than 10 volumes, only the first 10 volumes are shown in the performance chart. When you select volumes, the performance chart for those volumes is shown.
4. To select the time period for which you want to view performance data, select a value in the **Time** field. New data is retrieved from Tivoli Storage Productivity Center and a new chart is generated.
5. To select the metric that you want to view, select a value in the **Metric** field. New data is retrieved from Tivoli Storage Productivity Center and a new chart is generated.
6. To view the default chart for the 10 volumes with the worst performance, press Ctrl and click the selected volumes to deselect them. The chart is generated again.

Tip: To view descriptions of the metrics in the performance chart and the columns in the **Performance Chart Legend**, click the **Help** icon beside the **Refresh** icon .

7. Optional: To retrieve new data from Tivoli Storage Productivity Center, click the **Refresh** icon . The **Last updated** information in the toolbar is also updated.
8. Optional: To customize the column display in the **Performance Chart Legend**, use the pop-up menu for the table.
9. Optional: To export storage system metrics to a CSV file or to copy the data to the clipboard, click the **Export** icon or select the relevant option from the **Export menu** in the status bar.

Related reference:

"Report Displays" on page 1653
You can change the appearance of the **Storage System Metrics** table in the vSphere Web Client extension for Tivoli Storage Productivity Center. You can customize the appearance of table columns, including showing columns that are hidden, and hiding columns that are shown.

"CSV file exports" on page 1647
You can use the export actions to export the information from the reports tables in the vSphere Web Client extension for Tivoli Storage Productivity Center to a CSV file. You can also copy the information to the clipboard.

Storage System Metrics

Use the **Storage System Metrics** tab to view key performance data about the storage system that was collected from IBM Tivoli Storage Productivity Center during the last hour, day, or week. You can view the data in chart or table format, and export data to a CSV file. If no performance data was collected for the resource, the message `No data available` is displayed.

Performance Chart

You can view the key performance data for storage systems in a chart layout. The chart shows a visual representation of how storage resource performance trends over time. The x-axis shows the time period in which the data was collected. The y-axis shows the unit of measurement for a metric. Each line on the chart represents a metric and a resource. By default, the chart displays lines that represent the read I/O rate for the 10 volumes with the worst performance in the last hour.

Tip: Lines are overlaid on the chart if the corresponding volumes have the same metric value.

Click a row in the **Performance Chart Legend** to show the performance chart for that volume. To select multiple volumes, press Ctrl and click the volumes. To select a series of volumes, select the first volume, and then press Shift and click the last volume. You can also press Shift and the Up/Down Arrow. To deselect a volume, press Ctrl and click the selected volume.

Restriction: If you select more than 10 volumes, only the first 10 volumes are shown in the performance chart.

When you deselect all volumes, the default chart for the 10 volumes with the worst performance is generated again.

Note: When you select or deselect a volume, the chart is updated, but new data is not retrieved from Tivoli Storage Productivity Center. To retrieve new data, click the **Refresh** icon .

You can view the following information for volumes in the performance chart:

Time The time period that is used to select performance data. When you change the time period, new data is retrieved from Tivoli Storage Productivity Center and a new chart is generated.

You can select from the following options:

Last 1 Hour
The performance data that is selected on a sample basis from the data that was collected over the last hour. This sample data is collected at the same frequency as the data collection by the performance monitor for the volume. For example, if a performance monitor is scheduled to collect sample data every 15 minutes, each point on the chart is shown in 15-minute increments for the last hour. The default time period is **1 hour**.

Last 1 Day
The performance data that is selected on an hourly basis from the data that was collected over the last day.

Last 1 Week
The performance data that is selected on an hourly basis from the data that was collected over the last week.

Metric The performance metric that is displayed for the selected volume or volumes. When you change the **Metric**, new data is retrieved from Tivoli Storage Productivity Center and a new chart is generated.

You can select from the following options:

Read I/O rate (ops/s)
The average number of read operations per second over the specified time period.

Write I/O rate (ops/s)
The average number of write operations per second over the specified time period.

Total I/O rate (ops/s)
The average number of read and write operations per second over the specified time period.

Read data rate (KiB/s)
The average number of kibibytes per second that were transferred for read operations over the specified time period.

Write data rate (KiB/s)
The average number of kibibytes per second that were transferred for write operations over the specified time period.

Total data rate (KiB/s)
The average number of kibibytes per second that were transferred for read and write operations over the specified time period.

Read response time (ms/op)
The average number of milliseconds to complete a read operation over the specified time period.

Write response time (ms/op)
The average number of milliseconds to complete a write operation over the specified time period.

Overall response time (ms/op)
The average number of milliseconds to complete read and write operations over the specified time period.

Read transfer size (KiB/op)
The average number of kibibytes that are transferred per read operation over the specified time period.

Write transfer size (KiB/op)
The average number of kibibytes that are transferred per write operation over the specified time period.

Overall transfer size (KiB/op)
The average number of kibibytes that are transferred for read operations and write operations over the specified time period.

Refresh

Click the **Refresh** icon to retrieve new data from Tivoli Storage Productivity Center and generate a new chart.

Hover the mouse pointer over a point on a volume line to view the following information:

Volume WWN
The worldwide name of the volume.

Time The date and time that the performance data was collected.

The details for the selected metric are also displayed. If you hover the mouse pointer over the highest point on a volume line, the highest (worst) value for the **I/O rate**, **Data rate**, or **Response time** metric, is displayed. If the **Transfer size** metric is selected, the worst value is the lowest point on the volume line.

Performance Chart Legend

You can view performance data for storage systems in table format. The table shows performance information that is formatted into rows and columns. Each row represents a resource and a time stamp. Each metric is shown as a column in the table. The type of resource determines the metrics and information that is available. Click a column name to change the column that is used to sort the table. Click the link to open the associated page in the Tivoli Storage Productivity Center web-based GUI.

You can view information for each volume that is monitored by Tivoli Storage Productivity Center. The following information is displayed in the table by default:

Key The color key for the volume. The chart line for a volume is displayed in this color.

Tip: The chart color palette uses 24 colors. If the number of volumes is greater than 24, the colors repeat.

Name The name of the volume. Click the link to open the **Hypervisors** > **View Details** > **Disks** > **View Properties** page in the Tivoli Storage Productivity Center web-based GUI. If you are not authenticated, you must first log in to Tivoli Storage Productivity Center, and then the page is displayed.

Tip: Right-click on any column heading in the volumes table and select the **Unique ID** column. Use the information in this column to filter the search to locate the volume in the table.

Location
The location of the volume.

Snapshot time
The date and time that the performance data for the volume was collected.

Interval (sec.)
The size of the sample interval (in seconds) for the collection of performance data for the volume.

Storage system
The name of the storage system as displayed in Tivoli Storage Productivity Center. Click the link to open the **Storage System** > **Block Storage** > **View Details** page for the storage system in the Tivoli Storage Productivity Center web-based GUI. If you are not authenticated, you must first log in to Tivoli Storage Productivity Center, and then the page is displayed.

Storage volume
The name or identifier of the storage volume as displayed in Tivoli Storage Productivity Center. Click the link to open the **Storage System** > **Block Storage** > **View Details** > **Volume Properties** page for the storage volume in the Tivoli Storage Productivity Center web-based GUI. If you are not authenticated, you must first log in to Tivoli Storage Productivity Center, and then the page is displayed.

Storage pool
The name of the storage pool as displayed in Tivoli Storage Productivity Center. Click the link to open the **Storage System** > **Block Storage** > **View Details** > **Pool Properties** page for the storage pool in the Tivoli Storage Productivity Center web-based GUI. If you are not authenticated, you must first log in to Tivoli Storage Productivity Center, and then the page is displayed.

Read I/O rate sample (ops/s)
The average number of read operations per second for the last hour of sample performance data.

Read I/O rate hourly (ops/s)
The hourly average number of read operations per second for the selected last day or week of performance data.

Total I/O rate sample (ops/s)
The average number of read and write operations per second for the last hour of sample performance data.

Total I/O rate hourly (ops/s)
The hourly average number of read and write operations per second for the selected last day or week of performance data.

Overall response time sample (ms/op)
The average number of milliseconds that it took to complete each read and write operation for the latest sample of performance data.

Overall response time hourly (ms/op)
The hourly average number of milliseconds that it took to complete each read and write operation for the selected last day or week of performance data.

When you select a **Metric** in the performance chart, the corresponding columns for both sample and hourly averages are also displayed in the table:

Write I/O rate sample (ops/s)
The average number of write operations per second for the last hour of sample performance data.

Write I/O rate hourly (ops/s)
The hourly average number of write operations per second for the selected last day or week of performance data.

Read data rate sample (KiB/s)
The average number of kibibytes per second that were transferred for read operations for the latest sample of performance data.

Read data rate hourly (KiB/s)
The hourly average number of kibibytes per second that were transferred for read operations for the selected last day or week of performance data.

Write data rate sample (KiB/s)
The average number of kibibytes per second that were transferred for write operations for the latest sample of performance data.

Write data rate hourly (KiB/s)
The hourly average number of kibibytes per second that were transferred for write operations for the selected last day or week of performance data.

Total data rate sample (KiB/s)
The average number of kibibytes per second that were transferred for read and write operations for the latest sample of performance data.

Total data rate hourly (KiB/s)
The hourly average number of kibibytes per second that were transferred for read and write operations for the selected last day or week of performance data.

Read response time sample (ms/op)
The average number of milliseconds that it took to service each read operation for the latest sample of performance data.

Read response time hourly (ms/op)
The hourly average number of milliseconds that it took to service each read operation for the selected last day or week of performance data.

Write response time sample (ms/op)
The average number of milliseconds that it took to service each write operation for the latest sample of performance data.

Write response time hourly (ms/op)
The hourly average number of milliseconds that it took to service each write operation for the selected last day or week of performance data.

Read transfer size sample (KiB/op)
The average number of kibibytes that were transferred per read operation for the latest sample of performance data.

Read transfer size hourly (KiB/op)
The hourly average number of kibibytes that were transferred per read operation for the selected last day or week of performance data.

Write transfer size sample (KiB/op)
The average number of kibibytes that are transferred per write operation for the latest sample of performance data.

Write transfer size hourly (KiB/op)
The hourly average number of kibibytes that are transferred per write operation for the selected last day or week of performance data.

Overall transfer size sample (KiB/op)
The average number of kibibytes that were transferred per read and write operation for the latest sample of performance data.

Overall transfer size hourly (KiB/op)
The hourly average number of kibibytes that were transferred per read and write operation for the selected last day or week of performance data.

Related reference:

"Report Displays"
You can change the appearance of the **Storage System Metrics** table in the vSphere Web Client extension for Tivoli Storage Productivity Center. You can customize the appearance of table columns, including showing columns that are hidden, and hiding columns that are shown.

"CSV file exports" on page 1647
You can use the export actions to export the information from the reports tables in the vSphere Web Client extension for Tivoli Storage Productivity Center to a CSV file. You can also copy the information to the clipboard.

Report Displays:

You can change the appearance of the **Storage System Metrics** table in the vSphere Web Client extension for Tivoli Storage Productivity Center. You can customize the appearance of table columns, including showing columns that are hidden, and hiding columns that are shown.

Use the pop-up menu for the table to show, hide, size, or lock columns in the table:

Hide [Column Name]
Hides the selected column in the table.

Restriction: If only one column is visible, this column cannot be hidden.

Size [Column Name] to Fit
Resizes the selected column to fit the data in the column.

Size All Columns to Fit
Resizes all columns in the table to fit in the view and removes the scrollbar.

Lock First Column
Locks the first available column in the table and keeps it visible when you are scrolling.

Show/Hide Columns
Shows or hides the columns in the table.

Show Toolbar

Changes the display space for the table by showing or hiding the performance chart toolbar. If you hide the toolbar, you cannot change the metric or time that is displayed in the chart. You also cannot refresh the data and generate a new chart. Show the toolbar again to use these options.

Viewing Tivoli Storage Productivity Center fabric connections in vSphere Web Client extension

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to view information about the fabric that is connected to a storage adapter.

Start the vSphere Web Client, and log in to the vCenter Server system.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters**.
2. Select a host and click the **Manage** tab.
3. Click the **Storage** tab, and then click **Storage Adapters**.
4. In the **Adapter Details** section of the **Storage Adapters** page, select a Fibre Channel host bus adaptor.
5. Click the **Fabric Connections** tab to view the details of the fabric that the host is connected to.

   Tip: To view descriptions of the fields in the **Fabric Connections** tab, click the **Help** icon in the upper-right part of the tab.
6. Optional: Click the links in the Fabric name, Switch name, and Switch port fields to open the associated pages in the Tivoli Storage Productivity Center web-based GUI. If you are not authenticated, you must first log in to Tivoli Storage Productivity Center, and then the pages are displayed.

Viewing block storage reports in the vSphere Web Client

Use the vSphere Web Client to view information for volumes. When you configure Tivoli Storage Productivity Center as a VASA provider, you can view detailed block storage information in the SCSI Volumes (LUNs) report. When you use the Provision Block Storage wizard to create a data store, you can view the details in the Datastores report.

Tip: You can also view detailed block storage information from Tivoli Storage Productivity Center in the vSphere Client.

Start the vSphere Web Client, and log in to the vCenter Server system.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters**, and browse to the entity that contains the volumes or data stores that you want to view. You can select from the following entities:
   - Data center
   - Host
   - Cluster
   - Virtual machine

   The SCSI Volumes (LUNs) report is not available for data store clusters.
2. Click **Monitor** > **Storage Reports**.
3. To view the details for all volumes, in the **Report On** list, select the **SCSI Volumes (LUNs)** category.

4. To view the details for all data stores, in the **Report On** list, select the **Datastores** category.
5. Optional: Move the cursor over the column headings to see the description of each column. An asterisk indicates that the information in a report column is provided by Tivoli Storage Productivity Center.

Block storage reports in vSphere Web Client

You can view block storage reports in the vSphere Web Client. When you configure Tivoli Storage Productivity Center as a VASA provider, you can then view detailed block storage information in the storage reports. When you use the Provision Block Storage wizard to create a data store, you can view the details in the Datastores report.

Use the **Monitor** > **Storage Reports** page to view the reports. In the vSphere Web Client, an asterisk indicates that the information in a report column is provided by Tivoli Storage Productivity Center.

SCSI Volumes (LUNs) report

You can view the following information, which is provided by Tivoli Storage Productivity Center, in the SCSI Volumes (LUNs) report:

Committed
The amount of storage space that is allocated to the volume by the back-end storage system, if the volume is thin-provisioned.

Thin Provisioned
Shows whether the volume is thin provisioned.

System Capability
The storage capability of the volume. This capability describes the physical capabilities and services that the back-end storage offers, such as performance and redundancy of the volume.

Storage Array
The identifier of the storage array.

Identifier on Array
The identifier of the volume on the storage array.

Storage Provider Namespace
The namespace of the storage provider.

Datastores report

You can view the following information, which is provided by Tivoli Storage Productivity Center, in the Datastores report:

System Capability
The storage capability of the data store. This capability describes the physical capabilities and services that the back-end storage offers, such as performance and redundancy of the data store.

User Defined Capability
The namespace of the storage provider.

Storage Provider Namespace
The namespace of the storage provider.

Viewing file system reports in the vSphere Web Client

Use the vSphere Web Client to view detailed information for file systems. When you configure Tivoli Storage Productivity Center as a VASA provider, you can then view detailed file system information in the vSphere Web Client Datastores report and in the NAS Mounts report.

Tip: You can also view detailed file system information from Tivoli Storage Productivity Center in the vSphere Client.

Start the vSphere Web Client, and log in to the vCenter Server system.

1. From the vSphere Web Client Inventories view, select **Hosts and Clusters**, and browse to the entity that contains the NAS mounts that you want to view. You can select from the following entities:
   - Data center
   - Host
   - Cluster
   - Virtual machine

   The NAS Mounts report is not available for data store clusters.
2. Click **Monitor** > **Storage Reports**.
3. To view the details for all data stores, in the **Report On** list, select the **Datastores** category.
4. To view the details for all NAS mounts, in the **Report On** list, select the **NAS Mounts** category.
5. Optional: Move the cursor over the column headings to see the description of each column. An asterisk indicates that the information in a report column is provided by Tivoli Storage Productivity Center VASA provider.

File system reports in vSphere Web Client

You can view storage reports in the vSphere Web Client. When you configure Tivoli Storage Productivity Center as a VASA provider, you can then view detailed file system information in the storage reports.

Use the **Monitor** > **Storage Reports** page to view the **Datastores** report and the **NAS Mounts** report. In the vSphere Web Client, an asterisk indicates that the information in a report column is provided by Tivoli Storage Productivity Center.

You can view the following information that is provided by Tivoli Storage Productivity Center, in both the **Datastores** report and the **NAS Mounts** report:

System Capability
: The storage capability of the NAS mount. For file system storage, this column displays either the service class that is associated with the NAS mount or the value File.

Storage Provider Namespace
: The namespace of the storage provider.

Viewing Tivoli Storage Productivity Center VASA provider reports

Use the Tivoli Storage Productivity Center VASA provider to view Tivoli Storage Productivity Center data in vCenter reports.

When a Tivoli Storage Productivity Center VASA provider is online and inventory details are loaded, you can view Tivoli Storage Productivity Center data in vCenter reports. You can use the vSphere Client or the vSphere Web Client.

1. **vSphere Client**: Under **Home**, click **Hosts and Clusters**. Select **ESX** and click the **Storage Views** tab.
2. **vSphere Web Client**: Under **Home**, click **Hosts and Clusters**. Select **ESX** and in the **Monitor** tab, click **Storage Reports**.
3. From the **Report On** list, select **Show all SCSI Volumes (LUNs)**, **Show all Datastores**, or **Show all NAS Mounts**.

vCenter periodically updates storage data from the Tivoli Storage Productivity Center VASA provider.

To view current storage data in storage views and reports on vCenter, trigger the most recent vCenter Server updates from the Tivoli Storage Productivity Center VASA provider:

1. **vSphere Client**: Under **Home**, click **Storage Providers**.
2. Select the provider and click **Sync**.
3. Allow some minutes for the sync to take effect. Click **Refresh All**. Verify that the **Last Sync Time** value for the VASA provider is updated to the latest time stamp.

or

1. **vSphere Web Client**: Click **Hosts and Clusters**.
2. In the **Manage** tab, click **Storage Providers**.
3. Select the provider and click the **Rescan** icon.
4. Allow some minutes for the sync to take effect. Refresh the page. Verify that the **Last Rescan Time** value for the VASA provider is updated to the latest time stamp.

Now update the storage views and reports on vCenter:

1. Under **Home**, click **Hosts and Clusters**. Select **ESX** and click the **Storage Views** tab.
2. Click the **Update...** link to update the storage views.

or

1. Under **Home**, click **Hosts and Clusters**. Select **ESX** and in the **Monitor** tab, click **Storage Reports**.
2. Click the **Rescan vCenter Server for new inventory and storage devices** icon to update the storage reports.

Appendix A. Reference

View reference information that is related to Tivoli Storage Productivity Center. Topics include information about alerts, fabrics, commands, configuration and log files, performance metrics, protocols, standards, and accessibility features.

Return codes used by Storage Resource agent

This topic lists the return codes used by the Storage Resource agent.

The following table lists the return codes used by the Storage Resource agent during installation, uninstallation, and upgrade.

Table 129. Storage resource agent return codes

Return code	Explanation
1	There is a problem uninstalling the Fabric agent. For more information about what caused the uninstallation to fail, check the uninstallation logs for the Fabric agent.
2	Command not valid.
3	Option provided is not valid.
5	Argument is not valid.
6	Missing value for argument (e.g. -installLoc <Value>, where <Value> is missing).
7	Missing localized string in message file.
8	Probe is running.
9	Failed to open file for write.
10	Failed to close file.
11	Logfile not specified.
19	Failed in tracing.
21	Cannot spawn a probe because it is busy.
30	This is an internal error in initializing tracing. Save the error message and error log and contact your service representative for assistance in resolving the error.
32	Invalid socket.
33	Start service failed.
34	Registry entry not found.
35	Deployment of the agent failed, error creating startup scripts.
36	File does not exist.
40	Missing upgrade files.
41	Failed to extract files in upgrade process.
42	Failed to stop probe in upgrade process.
43	Failed to stop Agent.
44	Failed to start Agent.
45	Agent Registration to server failed.

© Copyright IBM Corp. 2001, 2013

Table 129. Storage resource agent return codes (continued)

Return code	Explanation
46	File extraction needs more space.
47	Failed to open archive file.
48	Agent did not start after upgrade.
49	Installation directory not valid at upgrade time.
50	Probe is running.
51	Data file not found.
52	Exit code not in the output file.
53	Failed to send job status.
54	Failed to copy certificate files.
55	Failed to create directory.
56	Failed to remove directory.
57	Exec command failed.
58	Conversion of wide character failed.
59	Installation directory not valid.
60	Server name not defined.
61	Error in removing entries from configuration file.
62	Failed to stop probe at uninstall time.
63	Failed to remove registry entry at uninstall time.
64	Failed to remove service entry at uninstall time.
65	Failed to stop service at uninstall time.
66	Specified server name is not valid.
67	There is an error installing the Storage Resource agent as part of the migration process. For more information about what caused the installation to fail, check the installation logs for the Storage Resource agent.
71	Failed to spawn process.
73	Failed to spawn probe process.
101	Failed to create lock at installation time.
102	Failed to stop probe at re-installation time.
103	Failed to stop agent at re-installation time.
104	Failed to create registry at installation time.
105	Failed to extract files at installation time.
106	Failed to create entries in configuration file at installation time.
107	Failed to stop service at re-installation time.
108	Service already exists.
109	Failed to create service.
110	Failed to start service.
111	Probe failed at installation time.
112	Creation of daemon failed.
113	Installation of GUID failed.
114	Commtype parameter is not valid.

Table 129. Storage resource agent return codes (continued)

Return code	Explanation
115	Specified port is in use.
116	Installation/Upgrade does not have enough space.
117	Installation in progress.
118	Cannot get server name from Server.
119	Installation location is not empty.
120	Missing parameter Server Name.
121	Missing parameter Server Port.
123	Missing parameter Server IP.
124	Missing parameter Agent Port.
125	Missing parameter Installation location.
126	A value has not been specified for parameter userID. A value for this parameter is required when using RXA-based communication to deploy a Storage Resource agent as a non-daemon service. Tivoli Storage Productivity Center uses this user ID when connecting to the computer on which the agent will be installed.
127	Deployment from Windows to Linux failed.
130	Failed to send probe results.
131	Failed to initialize Agent.
133	Missing port number for Service.
134	Get data file stat failed.
135	Get data file read failed.
137	Failed to send data to server.
138	Failed to receive data from server.
139	Full path not specified for copy file.
140	Create file failed in copy file function.
141	Write file failed in copy file function.
142	Open file failed in copy file function.
143	Read file failed in copy file function.
145	UCS conversion failed.
146	Server connection failed.
148	Failed to create zip file.
149	Failed to unzip file.
160	Failed to send scan results.
161	Failed to send TSM status results.
164	Failed to validate user.
165	Job file was not found.
166	Job was not stopped.
168	Not enough free space available while copying file from server.
169	Multipath driver not found.
170	Multipath device not found.
171	Multipath policy is not supported.

Table 129. Storage resource agent return codes (continued)

Return code	Explanation
172	Only Round Robin policy is supported for Multipath DM driver.

Agent types for monitoring fabrics and switches

Depending on the monitoring functions that you want to use for switches and fabrics, you can use a CIM agent, SNMP agent, or Storage Resource agent to manage switches. The vendors of the switches can help you determine which agents to use.

Table 130. Supported agent types for Switch and Fabric Functions

Function > Switch	Brocade	Cisco [1]	QLogic
Switch Performance Monitoring	CIMOM agent	CIMOM agent	CIMOM agent
Switch and switch port information collected	Recommended: CIMOM agent Also supported: Out-of-band Fabric agent or Storage Resource agent	Recommended: CIMOM agent Also supported: Out-of-band Fabric agent or Storage Resource agent	Out-of-band Fabric agent or Storage Resource agent
Topology connectivity information collected	Recommended: CIMOM agent Also supported: Out-of-band Fabric agent or Storage Resource agent	Out-of-band Fabric agent or Storage Resource agent	Out-of-band Fabric agent or Storage Resource agent
Zoning information collected [2]	Recommended: CIMOM agent	Storage Resource agent required in each VSAN [3] [4] [5]	Storage Resource agent
Tivoli Storage Productivity Center Alerts	Recommended: CIMOM agent Also supported: Out-of-band Fabric agent or Storage Resource agent	Out-of-band Fabric agent or Storage Resource agent	Out-of-band Fabric agent or Storage Resource agent
Hosts, endpoint devices, device-centric and host-centric information collected	Storage Resource agent	Storage Resource agent	Storage Resource agent
Switch “"Sensors and Events”	Out-of-band Fabric agent	Out-of-band Fabric agent	Out-of-band Fabric agent

Table 130. Supported agent types for Switch and Fabric Functions (continued)

Function > Switch	Brocade	Cisco [1]	QLogic

Notes

[1] The Storage Resource agent gets the information about a VSAN basis. Each VSAN is viewed as an individual SAN. The Out-of-band Fabric agent and CIMOM agent get the physical fabric information and can correlate the VSAN information to a physical infrastructure. The Out-of-band Fabric agent also collects some VSAN information.

[2] "Zone information collected" indicates that the active zone set and the inactive zoning library information is collected during a fabric probe.

[3] For many Cisco proprietary zone members, the values are displayed as zeros.

[4] If a zoning configuration has Cisco proprietary zone members, then you will not be able to modify the zoning for that configuration through Tivoli Storage Productivity Center. You need to use the element manager to remove proprietary zone members or do zoning.

[5] Orphan Cisco zones are not reported.

Supported storage systems providing full disk encryption and solid-state drives

IBM Tivoli Storage Productivity Center supports full disk encryption and solid-state drives in the IBM System Storage DS8000 and IBM System Storage DS5000 Storage Manager systems described in this topic.

DS8000 4.2 and later

Tivoli Storage Productivity Center supports full disk encryption and solid-state drives in DS8000 4.2 and later.

IBM System Storage DS5000 Storage Manager series

IBM Tivoli Storage Productivity Center supports full disk encryption and solid-state drives in the following DS5000 series systems:

- DS5100 and DS5300 - full disk encryption, solid-state drives, 1 TB SATA drives
- DS5020 - full disk encryption

To support full disk encryption and solid-state drives in DS5000, the LSI SMI-S Provider is required. For more information about the LSI SMI-S Provider, see the LSI Web site at http://www.lsi.com.

Alerts

Determining when and how you are alerted to conditions and violations of performance thresholds within a storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines the data about your resources for the conditions and performance thresholds that you specify.

Alerts are triggered by conditions and violations of performance thresholds that are detected during data collection and event processing. For storage systems such as SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, or the XIV system, events are polled every minute from the resource. For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to Tivoli Storage Productivity Center when conditions change on the resource.

Triggering conditions for alerts

The conditions that trigger alert notifications depend on the type of resource that you are monitoring. Some triggering conditions, like performance thresholds, require you to enter values for triggering alerts. In general, the following types of conditions can trigger alerts:

- A performance threshold was violated
- A data collection job did not complete
- A change occurred in the storage infrastructure

For example, you can use performance thresholds to be notified when the total I/O rate for storage systems falls outside a specified range. This information can help you identify areas in your storage infrastructure that are over utilized or under utilized. Tivoli Storage Productivity Center provides many metrics for measuring performance and determining violations of the thresholds that you specify.

Alert notifications and triggered actions

When an event occurs and triggers an alert, the alert is written to a log. You can also select one or more other ways to be notified of the event. These alert notifications include SNMP traps, Tivoli Enterprise Console events, login notifications, entries in Windows event log or UNIX syslog, and emails. Additionally, if a Storage Resource agent is deployed on a monitored server, you can run a script or start an IBM Tivoli Storage Manager job in response to the alert.

Acknowledging alerts: Some alerts are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify other alerts that must be reviewed and resolved.

Prerequisites for using alerts

The following conditions must be met to successfully use alerts:

- Data collection schedules are configured and scheduled to run regularly. For example, to detect violations of performance thresholds, you must run performance monitors to collect performance data about resources. Running performance monitors regularly also helps to establish a history of performance for trending analysis.
- If you want to be notified about an alert in some way other than an entry in the log file, such as using SNMP traps, Tivoli Enterprise Console events, or email, you must configure those alert destinations before you use the alert.
- If an alert is triggered based on an SNMP trap from the monitored resource, you must properly configure the SNMP server of the monitor resource to enable Tivoli Storage Productivity Center to listen to SNMP traps. The default port number is 162, and the default community is public.

Related concepts:

"Alerting" on page 341
Determining when and how you are alerted to conditions and violations of performance thresholds within your storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines the data about your resources for the conditions and performance thresholds that you specify.

Related tasks:

"Configuring alert destinations" on page 355
In the stand-alone GUI, configure the control information that is used to generate alerts when a corresponding event is discovered. You can specify default configuration settings for SNMP, TEC/OMNIbus, and email-directed alerts.

"Acknowledging alerts" on page 359
Some alerts in Tivoli Storage Productivity Center are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify the other alerts in lists that must still be reviewed and resolved.

Triggering conditions for alerts

You can define alerts so that IBM Tivoli Storage Productivity Center notifies you when a specified event occurs. Such events are the *triggering conditions* for the alert. The specific triggering conditions that you can use to define an alert depend on the sort of storage resource that you are monitoring.

Some triggering conditions, like thresholds, require you to enter values. If input is not required for the triggering condition you select, the value fields will be grayed out. For information about thresholds, see "Performance thresholds" on page 1690.

Related concepts:

"Alerting" on page 341
Determining when and how you are alerted to conditions and violations of performance thresholds within your storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines the data about your resources for the conditions and performance thresholds that you specify.

Related tasks:

"Adding resources" on page 257
You must add resources to Tivoli Storage Productivity Center before you can collect data, generate reports, and manage storage that is related to those resources.

Triggering conditions for computer alerts

Triggering conditions are events on a computer that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected on a computer.

Remember: The term "computer" is used in the stand-alone GUI to refer to a server.

Table 131. Triggering conditions and related error messages for computer alerts

Triggering Condition	Explanation	Related Error Message*
RAM Increased	The amount of RAM on a monitored computer is increased.	ALR0001I
RAM Decreased	The amount of RAM on a monitored computer is decreased.	ALR0001W
Virtual Memory Increased	The amount of virtual memory on a monitored computer is increased.	ALR0002I
Virtual Memory Decreased	The amount of virtual memory on a monitored computer is decreased.	ALR0002I

Table 131. Triggering conditions and related error messages for computer alerts (continued)

Triggering Condition	Explanation	Related Error Message*
New Disk Detected	A new disk is discovered on a monitored computer.	ALR0003I
Disk Not Found	A disk is removed from a monitored computer.	ALR0004E
New Disk Defect Found	A new grown disk defect is detected on a monitored disk.	ALR0011W
Grown Disk Defects Exceed *n* value	The grown disk defect threshold is exceeded. Enter a value for *n*.	ALR0012W
Disk Failure Predicted	A managed disk predicts that a disk failure is imminent.	ALR0007E
Cluster Resource Group Added	A new cluster resource group is added.	ALR0055I
Cluster Resource Group Removed	A cluster resource group is removed.	ALR0056I
Cluster Resource Group Moved	A cluster resource group is moved.	ALR0057I
Computer Discovered	A new, unmanaged computer is discovered. This alert condition is not triggered when hypervisor virtual machines are discovered.	ALR0022I
Computer Status Change Offline	A monitored computer goes offline.	ALR4273E, ALR4044W
Computer Status Change Online	A monitored computer comes online.	ALR4274I, ALR4045I
Computer Property Change	An HBA or HBA to a computer node is newly discovered, is missing, or is rediscovered.	ALR4275W
HBA Driver Version Change	The driver version of host bus adapter changes.	ALR4034W
HBA Firmware Version Change	The firmware version of host bus adapter changes.	ALR4035W
Note: * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.		

Triggering conditions for file system alerts

Triggering conditions are events on a file system that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is shown when the condition is detected on a file system.

Table 132. Triggering conditions and related error messages for file system alerts

Triggering Condition	Explanation	Related Error Message*
New File System Detected	A file system was detected for the first time. This alert applies to a file system on a Storwize V7000 Unified or a server managed by a Storage Resource agent.	ALR0005I, ALR0005W
File System Was Not Found	A file system is removed or unmounted from a cluster. This alert applies to a file system on a Storwize V7000 Unified or a server managed by a Storage Resource agent.	ALR0006E

Table 132. Triggering conditions and related error messages for file system alerts (continued)

Triggering Condition	Explanation	Related Error Message*
File System Was Reconfigured	The physical space definition of a managed file system is reconfigured. This alert applies to a file system on a Storwize V7000 Unified or a server managed by a Storage Resource agent.	ALR0008I
File System Free Space Less Than	A file system is low on available storage space. This value is usually entered as a percent.	ALR0009W
Available UNIX File System I-Nodes Less Than	A UNIX file system is low on free inodes. This value is usually as a percent.	ALR0010W
New File System Was Detected	A file system was detected for the first time.	ALR4427I, ALR0005W
File System Is Low on Relative Free Space	A file system is low on free space based on a percentage value threshold. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 Unified.	ALR4433W
File System Is Low on Absolute Free Space	A file system is low on free space based on an absolute value threshold. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 Unified.	ALR4433W
File System Mount State Changed to Normal	The mount state for a file system reached normal level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 Unified.	ALR4358I
File System Mount State Changed to Error	The mount state for a file system reached error level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 Unified.	ALR4356E
File System Is Low On Relative Free I-Nodes	A file system is low on free i-nodes based on a percentage value threshold. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 Unified.	ALR4458W
File System Is Low On Absolute Free I-Nodes	A file system is low on free i-nodes based on an absolute value threshold. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 Unified.	ALR4458W
New File Set Was Detected	A file set was detected for the first time. This alert applies only to Storwize V7000 Unified.	ALR4447I
File Set Is Missing	A file set that was found during a previous probe was not found in the current probe. This alert applies only to Storwize V7000 Unified.	ALR4474W
File Set Is Linked	A file set was linked to a file system. This alert applies only to Storwize V7000 Unified.	ALR4448I
File Set Is Unlinked	A file set was unlinked from a file system. This alert applies only to Storwize V7000 Unified.	ALR4455I
File Set Was Rediscovered	A file set that was missing was found by the current probe. This alert applies only to Storwize V7000 Unified.	ALR4475I
Quota Was Exceeded	A quota was exceeded. This alert applies only to Storwize V7000 Unified.	ALR4482W
New Quota Was Detected	A quota was detected for the first time. This alert applies only to Storwize V7000 Unified.	ALR4496I

Table 132. Triggering conditions and related error messages for file system alerts (continued)

Triggering Condition	Explanation	Related Error Message*
New Home System Was Detected	A home system was detected on a file set.	ALR4393I
Home System Is Removed	A home system is removed from a file set.	ALR4394I
Home System Is Missing	A home system that was found during a previous probe was not found in the current probe.	ALR4395W
Home System Was Rediscovered	A home system that was missing was found by the current probe.	ALR4396I
New Cache Client Was Detected	A cache client was detected for the first time.	ALR4403I
Cache Client Is Removed	A cache client that was found during a previous probe was not found in the current probe.	ALR4404I
New Cache File Set Was Detected	A cache file set was detected for the first time.	ALR4397I
Cache File Set Is Missing	A cache file set that was found during a previous probe was not found in the current probe.	ALR4398W
Cache File Set Was Rediscovered	A cache file set that was missing was found by the current probe.	ALR4399I
Cache File Set Name Has Changed	A cache file set name was changed.	ALR4400I
Cache File Set Mode Changed	A cache file set mode has changed.	ALR4402I
Note: * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.		

Triggering conditions for directory alerts

Triggering conditions are events in a monitored directory that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is shown when the condition is detected in a directory.

Table 133. Triggering conditions and related error messages for directory alerts

Triggering Condition	Explanation	Related Error Message*
Directory Not Found	A directory cannot be found.	ALR0014E
Directory Consumes More Than	A user or directory storage quota is exceeded. Enter a value, usually in megabytes.	ALR0015W
Note: * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.		

Triggering conditions for instance alerts

Triggering conditions are events on a database instance that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected on an instance.

Table 134. Triggering conditions and related error messages for instance alerts

Database Application	Triggering Condition	Explanation	Related Error Message*
All	New Database - Tablespace Discovered	A new database-tablespace is discovered.	ALR0028I
Oracle	Archive Log Directory Contains More Than	An archived log directory quota is exceeded. Enter a value for quota.	ALR0027W
	New Tablespace Discovered	A new tablespace is discovered.	ALR0028I
SQL Server	New Database Discovered	A new database is discovered.	ALR0028I
Sybase	New Device Discovered	A new device is discovered.	ALR0042I
	Device Dropped	A device is dropped.	ALR0043E
	New Database Discovered	A new database is discovered.	ALR0028I
DB2	New Tablespace Discovered	A new database-tablespace is discovered.	ALR0028I

Note: * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.

Trigger conditions for database-table space alerts

Triggering conditions are events on a database-table space that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected on a database-table space.

Table 135. Triggering conditions and related error messages for database-table space alerts

Database application	Triggering Condition	Explanation	Related Error Message*
All	Database - Table Space Free Space Less Than	The free space on a database-table space is less than a specified amount.	ALR0031W
	Database - Table Space Offline	A database-table space is taken offline.	ALR0030E
	Database - Table Space Dropped	A database-table space is dropped.	ALR0029E
Oracle	Table Space Free Space Less Than	The free space on a table space is less than a specified amount.	ALR0031W
	Free Space Fragmented into More Than	The free space on a table space is fragmented into more than the specified number of extents.	ALR0032W
	Largest Free Extent Less Than	The largest free extent is less than the specified size.	ALR0033W
	Table Space Offline	A table space is taken offline	ALR0030E
	Table Space Dropped	A table space is dropped.	ALR0029E
SQL Server	Database Data Free Space Less Than	The free space on a database is less than a specified amount.	ALR0031W
	Database Log Free Space Less Than	The free space on a database log is less than the specified percent.	ALR0041W
	Database Offline	A database is taken offline.	ALR0030E
	Database Dropped	A database is dropped.	ALR0029E

Table 135. Triggering conditions and related error messages for database-table space alerts (continued)

Database application	Triggering Condition	Explanation	Related Error Message*
Sybase	Database Data Free Space less Than	The free space on a database is less than a specified amount.	ALR0031W
	Database Log Free Space Less Than	The free space on a database log is less than the specified percent.	ALR0027W
	Database Dropped	A database is dropped	ALR0029E
DB2	Table Space Free Space Less Than	The free space on a table space is less than a specified amount.	ALR0031W
	Table Space Offline	A table space is taken offline.	ALR0030E
	Table Space Dropped	A table space is dropped	ALR0029E
Note: * Use the error message identifiers in this column to locate any additional information about the error messages that are related to triggering conditions in the Messages section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.			

Triggering conditions for table alerts

Triggering conditions are events on a database table that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected on a table.

Table 136. Triggering conditions and related error messages for table alerts

Database application	Triggering Condition	Explanation	Related Error Message*
All	Total Table Size Greater Than	The total size of a table is greater than a specified size.	ALR0036W
	Table Dropped	A table is dropped.	ALR0039E
Oracle	Max Extents - Allocated Extents	A table segment is near its maximum allowed extents.	ALR0035W
	Segment Has More Than	A table segment has taken too many extents.	ALR0034W
	Chained Row Count Greater Than	A table has too many chained rows.	ALR0037W
	Empty Used Segment Space Exceeds	An empty used segment space exceeds the specified size.	ALR0038W
	Total Table Size Greater Than	The total size of a table is greater than a specified size.	ALR0036W
	Table Dropped	A table is dropped.	ALR0039E
SQL Server	Total Table Size Greater Than	The total size of a table is greater than a specified size.	ALR0036W
	Table Dropped	A table is dropped.	ALR0039E
Sybase	Forwarded Row Count Greater Than	A forwarded row count is greater than the specified percent.	ALR0037W
	Total Table Size Greater Than	The total size of a table is greater than a specified size.	ALR0036W
	Table Dropped	A table is dropped.	ALR0039E

Table 136. Triggering conditions and related error messages for table alerts (continued)

Database application	Triggering Condition	Explanation	Related Error Message*
DB2	Overflow Row Count Greater Than	An overflow row count is greater than the specified percent.	ALR0037W
	Total Table Size Greater Than	The total size of a table is greater than a specified size.	ALR0036W
	Table Dropped	A table is dropped.	ALR0039E

Note: * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.

Triggering conditions for storage system alerts

Triggering conditions are events in a storage system that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected on a storage system.

Table 137. Triggering conditions and related error messages for storage-system alerts

Triggering Condition	Explanation	Related Error Message[1]
New Storage System discovered	A new storage system is discovered.	ALR1114I
Cache Increased	A storage system cache increases.	ALR0050W
Cache Decreased	A storage system cache decreases.	ALR0050W
New Disk Detected[2]	A new disk is detected.	ALR4323I
Disk Not Found[2]	A previously detected disk is not found.	ALR4324W
Storage System Not Found	A previously detected storage system is not found.	ALR0048W
Storage System Status Change Offline	A storage system goes offline.	ALR4241E
Storage System Status Change Online	A storage system comes online.	ALR4242W
Storage System Version Change	A storage system version changes.	ALR4243W
Storage System Property Change	A port on a storage system is newly discovered, missing, or rediscovered.	ALR4278W
Storage System Allocated Capacity Change	The allocated capacity of a storage system changes.	ALR4244W
Storage System Available Capacity Change	The available capacity of a storage system changes.	ALR4245W
Storage System Backend Capacity Change	The allocated capacity of a storage system changes.	ALR4246W
Backend Storage System Status Change Offline	A back-end storage system goes offline.	ALR4247W
Backend Storage System Status Change Online	A back-end storage system comes online.	ALR4248I

Table 137. Triggering conditions and related error messages for storage-system alerts (continued)

Triggering Condition	Explanation	Related Error Message[1]
Volume Status Change Offline	A volume goes offline.	ALR4249W
Volume Capacity Change	The capacity of a volume changes.	ALR4251W
Pool Status Change	A pool is missing or rediscovered.	ALR4252W
Pool Discovered	A pool is discovered.	ALR4253I
Pool Status Change Offline	A pool goes offline.	ALR4254W
Pool Status Change Online	A pool comes online.	ALR4255I
Pool Capacity Change	The capacity of a pool changes.	ALR4256W
Pool Available Space Change	The available space in a pool changes.	ALR4257W
Pool Available Space Lower Than Critical Boundary	The available space is too low for a pool on a storage system. The measured value violates the critical boundary.	ALR4541E
Pool Available Space Lower Than Warning Boundary	The available space is too low for a pool on a storage system. The measured value violates the warning boundary.	ALR4542W
Pool Allocation Percentage Higher Than Critical Boundary	The allocation percentage is too high for a pool on a storage system. The measured value violates the critical boundary.	ALR4543E
Pool Allocation Percentage Higher Than Warning Boundary	The allocation percentage is too high for a pool on a storage system. The measured value violates the critical boundary.	ALR4544W
Realizable Allocation Percentage Below Critical Boundary	The realizable allocation percentage is too low for a pool on a storage system. The measured value violates the critical boundary.	ALR4545E
Realizable Allocation Percentage Below Warning Boundary	The realizable allocation percentage is too low for a pool on a storage system. The measured value violates the warning boundary.	ALR4546W
Pool Available Space Too Low	Two threshold boundary values can be specified, a warning and a critical boundary. You can specify a value in Gigabytes or as a %.	ALR4541E, ALR4542W
Pool Virtual Allocation Too High	Two threshold boundary values can be specified, a warning and a critical boundary. You can specify a value in Gigabytes or as a %.	ALR4543E, ALR4544W
Pool Shortfall too High	Two threshold boundary values can be specified, a warning and a critical boundary. You can specify a value as a %.	ALR4545E, ALR4546W
Node Status Change Offline	A node goes offline.	ALR4226W
Node Status Change Online	A node comes online.	ALR4227I
Node State Change	An IBM System Storage SAN Volume Controller node is missing or rediscovered.	ALR4224W
Node Discovered	A SAN Volume Controller node is discovered.	ALR4225I

Table 137. Triggering conditions and related error messages for storage-system alerts (continued)

Triggering Condition	Explanation	Related Error Message[1]
Datapath State Change	The state of a data path changes.	ALR4353W
Datapath Discovered	A data path is discovered.	ALR4354I
New Volume discovered	A new volume is discovered during a probe of a storage system.	ALR4325W
Volume not found	A volume which had existed before is not found during a probe of a storage system.	ALR4326W
Link Failure Rate Threshold	The average number of miscellaneous fibre channel link errors (for example, an unexpected Not Operational (NOS) received or a link state machine failure detected) per second that were encountered by a port over a period of time.	ALR0548E, ALR0549W, ALR0550W, ALR0551E
Error Frame Rate Threshold	The threshold for the number of frames per second that violated Fibre Channel protocol for a port over a period of time.	ALR0544E, ALR0545W. ALR0546W, ALR0547E
Overall Port Response Time Threshold	The threshold for the average number of milliseconds that it took to service each operation (send and receive), by a particular port over a particular time interval.	ALR0540E, ALR0541W, ALR0542W, ALR0543E
Total Port Data Rate Threshold	The threshold for the average number of megabytes transferred (send and receive) for a particular port over a particular time interval	ALR0536E, ALR0537W, ALR0538W, ALR0539E
Total Port IO Rate Threshold	The threshold for the average number of send and receive operations processed per second, for a particular port over a particular time interval	ALR0532E, ALR0533W, ALR0534W, ALR0535E
CPU Utilization Threshold	The threshold is based on the average utilization percentage of the CPUs in the SAN Volume Controller nodes. This value for each node is checked against the threshold boundaries for each collection interval.	ALR0556E, ALR0557W, ALR0558W, ALR0559E
Cache Holding Time Threshold	The threshold is reached for the average cache holding time, in seconds, for I/O data in a particular component over a particular time interval. Shorter time periods indicate adverse performance	ALR0528E, ALR0529W, ALR0530W, ALR0531E
Write-cache Delay Percentage Threshold	The threshold is based on the percentage of I/O operations that were delayed due to write-cache space constraints. This value for each node is checked against the threshold boundaries for each collection interval	ALR0524E, ALR0525W, ALR0526W, ALR0527E
Total Data Rate Threshold	For the storage system controllers or I/O groups, the threshold for the average number of MB per second for read and write operations is reached	ALR0520E, ALR0521W, ALR0522W, ALR0523E
Total I/O Rate Threshold	For the storage system controllers or I/O groups, the threshold for the average number of I/O operations per second is reached	ALR0516E, ALR0517W, ALR0518W, ALR0519E
Overall Backend Response Time Threshold	The threshold for the average number of milliseconds that it takes to respond to MDisk I/O operations is reached	ALR0512E, ALR0513W, ALR0514W, ALR0515E
Total Backend Data Rate Threshold	For MDisk groups, the threshold for the average number of megabytes per second that are transferred for MDisk read and write operations is reached	ALR0508E, ALR0509W, ALR0510W, ALR0511E
Zero Buffer-Buffer Credit Timer Threshold	The number of microseconds that the port was not able to send frames because of insufficient buffer credit since the last node reset.	ALR0624E, ALR0625W, ALR0626W, ALR0627E

Table 137. Triggering conditions and related error messages for storage-system alerts (continued)

Triggering Condition	Explanation	Related Error Message[1]
Total Backend I/O Rate Threshold	For MDisk groups, the threshold for the average number of MDisk read and write operations per second is reached.	ALR0504E , ALR0505W, ALR0506W, ALR0507E
Invalid Transmission Word Rate Threshold	The threshold for the average number of times per second that "bit" errors are detected on a port over a period of time.	ALR0620E, ALR0621W, ALR0622W, ALR0623E
CRC Error Rate Threshold	The threshold for the average number of frames received per second where the CRC in the frame did not match the CRC computed by the receiver for a component over a period of time.	ALR0616E, ALR0617W, ALR0618W, ALR0619E
Disk Utilization Percentage Threshold	The threshold is reached for the approximate utilization percentage of a particular array over a particular time interval. For example, the average percentage of time that the disk associated with the array was busy	ALR0500E, ALR0501W, ALR0502W, ALR0503E
Port Receive Bandwidth Percentage Threshold[3]	The threshold for the average port bandwidth utilization percentage for receive operations.	ALR0612E, ALR0613W, ALR0614W, ALR0615E
Port Send Bandwidth Percentage Threshold[3]	The threshold for the average port bandwidth utilization percentage for send operations.	ALR0608E, ALR0609W, ALR0610W, ALR0611E
Port Receive Utilization Percentage Threshold	The threshold for the average amount of time that ports are busy receiving data.	ALR0604E, ALR0605W, ALR0606W, ALR0607E
Port Send Utilization Percentage Threshold	The threshold for the average port bandwidth utilization percentage for send operations.	ALR0600E, ALR0601W, ALR0602W, ALR0603E
Peak Back-end Write Response Time Threshold	The threshold for the peak (worst) response time among all MDisk write operations by a node.	ALR0596E, ALR0597W, ALR0598W, ALR0599E
Non-preferred Node Usage Percentage Threshold	The threshold for the non-preferred node usage percentage of an I/O group.	ALR0592E, ALR0593W, ALR0594W, ALR0595E
Port to Local Node Receive Queue Time Threshold	The threshold for the average number of milliseconds that each receive operation issued to another node in the local SVC cluster spent on the queue before being issued.	ALR0588E, ALR0589W, ALR0590W, ALR0591E
Port to Local Node Send Queue Time Threshold	The threshold for the average number of milliseconds that each send operation issued to another node in the local SVC cluster spent on the queue before being issued.	ALR0584E, ALR0585W, ALR0586W, ALR0587E
Port to Local Node Receive Response Time Threshold	The threshold for the average number of milliseconds it took to service each receive operation from another node in the local SVC cluster.	ALR0580E, ALR0581W, ALR0582W, ALR0583E
Port to Local Node Send Response Time Threshold	The threshold for the average number of milliseconds it took to service each send operation from another node in the local SVC cluster.	ALR0576E, ALR0577W, ALR0578W, ALR0579E
Back-end Write Queue Time Threshold	The threshold for the average number of milliseconds that each write operation spent on the queue before being issued to the back-end device.	ALR0572E, ALR0573W, ALR0574W, ALR0575E
Back-end Read Queue Time Threshold	The threshold for the average number of milliseconds that each read operation spent on the queue before being issued to the back-end device.	ALR0568E, ALR0569W, ALR0570W, ALR0571E

Table 137. Triggering conditions and related error messages for storage-system alerts (continued)

Triggering Condition	Explanation	Related Error Message[1]
Back-end Write Response Time Threshold	The threshold for the average number of milliseconds that it took to service each array and MDisk write operation.	ALR0564E, ALR0565W, ALR0566W, ALR0567E
Back-end Read Response Time Threshold	The threshold for the average number of milliseconds that it took to service each array and MDisk read operation.	ALR0560E, ALR0561W, ALR0562W, ALR0563E

Notes:

1. Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.
2. This condition does not trigger alerts for the following storage systems:
   - Storwize V3500
   - Storwize V3700
   - Storwize V7000
   - Storwize V7000 Unified
   - SAN Volume Controller
3. This condition does not trigger alerts for the DS6000.

Related tasks:

"Creating a performance threshold and viewing violations" on page 100
You can monitor the performance of a storage environment by defining performance thresholds for storage systems and SAN switches. A threshold violation is detected when a performance monitor discovers that the measured value of a metric falls outside of the specified threshold range for a resource. By defining alerts that are triggered by performance thresholds, you can get timely information about performance issues in your environment.

Related reference:

"Performance thresholds" on page 1690
Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

Triggering conditions for fabric alerts

Triggering conditions are events in a fabric that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected in a fabric.

Table 138. Triggering conditions and related error messages for fabric alerts

Triggering Condition	Explanation	Related Error Message*
Fabric Discovered	A fabric is discovered.	ALR4046I
Fabric State Change	A fabric is missing or rediscovered.	ALR4047W, ALR4048I
Fabric to Switch Change	A switch to fabric association is discovered, rediscovered, or missing.	ALR4066I, ALR4067W, ALR4068I
Zone State Change	A zone is missing or rediscovered.	ALR4052W, ALR4053I, ALR4090W, ALR4094I
Zone Discovered	A zone is discovered.	ALR4051I, ALR4091I
Zone Set State Change	A zone set in a fabric was activated, deactivated, or has gone missing.	ALR4055W, ALR4089W, ALR4093E, ALR4056I, ALR4095I
Zone Set Discovered	A zone set is discovered.	ALR4054I, ALR4092I

Table 138. Triggering conditions and related error messages for fabric alerts (continued)

Triggering Condition	Explanation	Related Error Message*
Zone to Zone Alias Change	A zone to zone alias association is discovered, missing, or rediscovered.	ALR4079W, ALR4080I, ALR4110W, ALR4111I, ALR4078I, ALR4109I
Zone to Zone Member Change	A zone to zone member association is discovered, missing, or rediscovered.	ALR4081I, ALR4096I, ALR4082I, ALR4083I, ALR4097I, ALR4098I
Zone Set to Zone Change	A zone set to zone association is discovered, missing, or rediscovered.	ALR4084I, ALR4099I, ALR4085I, ALR4086I, ALR4107I, ALR4108I
New Connection	A new connection is detected.	ALR4063I
Connection State Change	A connection is missing or rediscovered.	ALR4064W, ALR4065I
Zone Alias to Member Change	A zone alias to member association is discovered, rediscovered, or missing.	ALR4329I, ALR4330I, ALR4331I
Endpoint Discovered	An endpoint is discovered.	ALR4000I
Endpoint State Change	An endpoint is missing or rediscovered.	ALR4001W, ALR4002I
Note: * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.		

Triggering conditions for switch alerts

Triggering conditions are events on a switch that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected on a switch.

Important: The following triggering conditions are performance thresholds. A performance monitor must collect data about a switch before Tivoli Storage Productivity Center can determine if a performance threshold has generated an alert.

- Total Port Data Rate Threshold
- Port Receive Bandwidth Percentage Threshold
- Port Send Bandwidth Percentage Threshold
- Link Failure Rate Threshold
- Error Frame Rate Threshold
- Total Port Packet Rate Threshold

Table 139. Triggering conditions and related error messages for switch alerts

Triggering Condition	Explanation	Related Error Message*
Switch Discovered	A switch is discovered.	ALR4020I
Switch State Change	A switch is missing or is rediscovered.	ALR4021E, ALR4022I
Switch Property Changes	A port on the switch has changed to online or offline, is discovered, is missing, or is rediscovered.	ALR4015I, ALR4016W, ALR4017I, ALR4018W, ALR4019I,
Switch Status Degraded	The status of the switch has changed from OK to a reduced level. This condition does not necessarily mean that the switch is missing.	ALR4024W
Switch Status Improved	The status of the switch has been upgraded to OK.	ALR4025I
Switch Version Change	The version of a switch changes.	ALR4023W

Table 139. Triggering conditions and related error messages for switch alerts (continued)

Triggering Condition	Explanation	Related Error Message*
Switch to Port Change	A switch to switch port association is discovered, is missing, or is rediscovered.	ALR4069I, ALR4070W, ALR4071I
Switch Blade Change	A switch module is discovered, is missing, or is rediscovered.	ALR4026I, ALR4027W, ALR4028I
Switch Blade Change Offline	A switch module goes offline.	ALR4029E
Switch Blade Change Online	A switch module comes online.	ALR4030I
Total Port Data Rate Threshold	A total port data rate threshold is reached.	ALR0536E, ALR0537W, ALR0538W, ALR0539E
Port Receive Bandwidth Percentage Threshold	The threshold on the average port bandwidth utilization percentage for receive operations. The Port Send Utilization Percentage metric is checked against the threshold boundaries for each collection interval. This threshold uses the following default values: 85,75,-1,-1.	ALR0612E, ALR0613W, ALR0614W, ALR0615E
Port Send Bandwidth Percentage Threshold	The threshold on the average port bandwidth utilization percentage for send operations. The Port Send Utilization Percentage metric is checked against the threshold boundaries for each collection interval. This threshold uses the following default values: 85,75,-1,-1.	ALR0608E, ALR0609W, ALR0610W, ALR0611E
Link Failure Rate Threshold	A link failure rate threshold is reached.	ALR0548E, ALR0549W, ALR0550W, ALR0551E
Error Frame Rate Threshold	An error frame rate threshold is reached.	ALR0544E, ALR0545W, ALR0546W, ALR0547E
Total Port Packet Rate Threshold	A total port packet rate threshold is reached.	ALR0552E, ALR0553W, ALR0554W, ALR0555E
Note: * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.		

Triggering conditions for endpoint device alerts

Use this topic to view the events that can trigger endpoint device alerts and their related error message identifiers.

Table 140. Triggering conditions and related error messages for endpoint device alerts

Triggering Condition	Explanation	Related Error Message*
Endpoint Discovered	An endpoint device is discovered.	ALR4239W
Endpoint State Change	An endpoint device is missing or is rediscovered.	ALR4231W
Endpoint Version Change	The version for an endpoint device changes.	ALR4313W
Note: * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.		

Triggering conditions for hypervisor alerts

Triggering conditions are events on a hypervisor that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected on a hypervisor.

Table 141. Triggering conditions and related error messages for hypervisor alerts

Triggering Condition	Explanation	Related Error Message*
Hypervisor Discovered	A hypervisor is discovered.	ALR4315I
Hypervisor Missing	A hypervisor that was previously discovered is missing.	ALR4321W
Virtual Machine Added	A new virtual machine was added.	ALR4316I
Virtual Machine Removed	A virtual machine was removed.	ALR4317W
Note: * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.		

Triggering conditions for Other NAS alerts

Triggering conditions are events on network-attached storage (NAS) that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when the condition is detected on a NAS device.

Table 142. Triggering conditions and related error messages for NAS alerts

Triggering Condition	Explanation	Related Error Message*
Node Status Changed to Normal	The clustered Common Internet File System (CIFS) for a NAS device is active. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4365I
Node Status Changed to Warning	The clustered CIFS for a NAS device is disabled. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4366W
Node Status Changed to Error	The status of a clustered CIFS for a NAS device reached the error level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4367E
GPFS Status Reached Normal Level	The General Parallel File System (GPFS) for a NAS device is active. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4368I
GPFS Status Reached Warning Level	The GPFS state for a NAS device reached the warning level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4369W
GPFS Status Reached Error Level	The GPFS for a NAS device is down. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4370E
Node State Changed	The state of a node changed.	ALR1245W
Node Was Discovered	A node was discovered.	ALR1246W
Node Status Changed to Online	A node status changed to online.	ALR4227W
Node Status Changed to Offline	A node status changed to offline.	ALR4226W
Processor Usage Reached Normal Level	Processor use for a network-attached storage (NAS) device reached the normal level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4361I

Table 142. Triggering conditions and related error messages for NAS alerts (continued)

Triggering Condition	Explanation	Related Error Message*
Processor Usage Reached Warning Level	Processor use for a NAS device reached the warning level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4360W
Processor Usage Reached Error Level	Processor use for a NAS device reached the error level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4359E
Memory Usage Reached Normal Level	Memory use for a NAS device reached the normal level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4364I
Memory Usage Reached Error Level	Memory use for a NAS device reached the error level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.	ALR4362E
Cluster Was Discovered	A cluster was discovered.	ALR4528I
Cluster Was Removed	A cluster that was found during a probe was not found in the current probe.	ALR4529I
Cluster Was Rediscovered	A cluster that was missing was found by the current probe.	ALR4530I
New Export Was Detected	An export was detected for the first time.	ALR4460I
Export Is Missing	An export that was found during a previous probe was not found in the current probe.	ALR4470W
Export Was Rediscovered	An export that was missing was found by the current probe.	ALR4471I
Export State Changed	The state of an export changed.	ALR4461W
Export Path Changed	An export was linked to a new path.	ALR4462W
Export Protocol Changed	The protocols that were configured for an export changed.	ALR4463W
New NSD Detected	A Network Shared Disk (NSD) was detected for the first time.	ALR4503I
NSD Is Missing	An NSD that was found during a previous probe was not found in the current probe.	ALR4505W
NSD State Changed	The state of an NSD changed.	ALR4507W
NSD Status Reached Normal Level	NSD use for a NAS device reached the normal level.	ALR4387I
NSD Status Reached Warning Level	NSD use for a NAS device reached the warning level.	ALR4386W
NSD Status Reached Error Level	NSD use for a NAS device reached the error level.	ALR4385E
Node Is Selected As Cache Gateway Node	A node was selected as a cache gateway node.	ALR4391I
Node Is Unselected As Cache Gateway Node	A node was unselected as a cache gateway node.	ALR4392I
Note: * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.		

Triggering conditions for replication alerts

Triggering conditions are events in Tivoli Storage Productivity Center for Replication that generate alerts in Tivoli Storage Productivity Center. Each condition has a related error message identifier that is displayed when it is detected.

Table 143. Triggering conditions and related error messages for replication alerts

Triggering Condition	Explanation	Related Error Message*
Replication Session Change	An alert is generated if a replication session changes to the following states: • Terminating • Target Available • Suspended, Suspended H2H3, Suspended H1H3 • Recovering • Preparing • Prepared • Flashing • Defined An alert is generated if the following conditions are detected for a role pair: • The role pair in a replication session passed the Recovery Point Objective (RPO) warning threshold. • The role pair in a replication session passed the RPO severe threshold.	ALR4339W
Replication Suspending Event Notification	The state of a replication session changed to suspending.	ALR4339W
Replication Communication Failure	An alert is generated when the following communication errors are detected: • The Replication standby server encountered communication errors with the Active server. • The Replication server timed out when attempting to communicate with the storage system. • The Replication server encountered communication errors with the storage system. • The Replication active server stopped communicating with the standby server.	ALR4339W
Replication Management Server State Change	An alert is generated when the replication HA connection changes to the following states: • Unknown • Synchronized • Synchronization Pending • Disconnected • Disconnected Consistent	ALR4339W
Replication Configuration Change	An alert is generated when the following replication configuration changes are detected: • One or more path definitions were added or removed. • One or more copy sets were added to a session. • One or more copy sets were deleted from a session.	ALR4339W

Table 143. Triggering conditions and related error messages for replication alerts (continued)

Triggering Condition	Explanation	Related Error Message*
Note: * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.		

Alert suppression actions

If you selected a threshold as a triggering condition for alerts, you can specify the conditions under which alerts are triggered and choose whether you want to suppress repeating alerts. Alerts can be suppressed to avoid generating too many alert log entries or too many actions when the triggering condition occurs often. You can view suppressed alerts in the constraint violation reports.

You can define the following options that enable you to specify conditions that trigger and suppress alerts. If a threshold is not selected as a triggering condition, these options are not available.

Trigger alerts for both critical and warning conditions
Generates an alert upon the violation of either critical or warning threshold boundaries. This is the default.

Trigger alerts for critical conditions only
Generates alerts only upon violation of one of the critical threshold boundaries. Violation of a warning boundary creates an entry in the constraint violation report, but does not result in an entry in the alert log or an action being triggered.

Trigger no alerts
Does not generate an alert upon violation of any threshold boundaries. Creates entries only in the constraint violation report.

Do not suppress repeating alerts
Does not suppress any repeating alerts. This is the default.

Suppress alerts unless the Triggering Condition has been violated continuously for a specified length of time
Generates alerts only if the triggering condition has occurred continuously within the length of time specified in the **Length of time** field. Alerts for the first and any subsequent occurrences of the triggering condition within the specified time in minutes will be suppressed. At the point that there have been consecutive occurrences with the specified number of minutes, an alert is generated. When the specified suppression period has expired, the cycle starts again. Note that the timing for this feature is based on the IBM Tivoli Storage Productivity Center server clock rather than the various system clocks.

This option is useful for cases when a single occurrence of the triggering condition may be insignificant, but repeated occurrences can signal a potential problem.

Suppress alerts if a repeat violation has occurred within a specified length of time after the initial violation of the Triggering Condition
Generates alerts only for the first occurrence of the triggering condition. Alerts for repeated occurrences of the triggering condition within the length of time specified in the **Length of time** field are suppressed. When the specified suppression period has expired, the cycle starts again. Note

that the timing for this feature is based on the IBM Tivoli Storage Productivity Center server clock rather than the various system clocks.

This option is useful for avoiding e-mail messages or similar disruptive alerts if the same triggering condition occurs repeatedly in successive sample passes. This option generally useful for all threshold types.

Triggered actions for alerts

You define triggered actions to occur as a result of the alert.

An alert is always entered in the error log file each time it occurs. In addition, you can define the following notifications or triggered actions:

SNMP Trap
Generates an SNMP trap message to any network-management system, console, or terminal. You must set up the SNMP trap ringer with the `fabric.mib` and `tivoliSRM.mib` files.

TEC / OMNIbus Event
Sends an alert to the IBM Tivoli Enterprise Console or OMNIbus EIF probe server in your environment. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze theIBM Tivoli Storage Productivity Center events. Tivoli Enterprise Console also performs responses such as sending further notifications, creating or updating trouble tickets, and running programs. You must load the `tivoliSRM.baroc` and `fabric.baroc` files into the active rulebase of the Tivoli Enterprise Console server.

Login Notification
Tivoli Storage Productivity Center lists the alerts for the specified users when they log in. Use the **Preferences** > **Edit General** option on the Tivoli Storage Productivity Center menu bar to control which alerts a user views upon logging in to the product.

Windows Event Log, UNIX Syslog
Writes the alerts to the operating-system log. You can select from the following levels of operating-system events:

- UNIX: Informational, Notice, Warning, Error, Critical, or Alert
- Microsoft Windows: Error, Warning, and Information.

For UNIX, you also can select User or Local.

Run Script
Runs a script in response to the alert. Use this option to run third-party tools for such actions as recovery and provisioning.

1. Click **Define** to specify the script that is run when the triggering condition occurs. The **Specify Script** window opens.
2. Type the name of the script in the **Script Name** field.
3. Select where to run the script.

If parameters are given for the triggering condition, they are listed in the **Specify Script** window. These parameters are used when the alert is triggered.

Note: Scripts can only be run against systems on which a Storage Resource agent is deployed.

Archive/Backup
(Constraints only) Runs an IBM System Storage Productivity Center backup or archive job in response to an alert associated with a constraint. Click the Define button.

Email Sends an email. Use this section to view, add, and delete email addresses and edit the text of the message in the email.

For **Directory Alerts**, if the **Directory Consumes More Than** triggering condition is selected, additional fields are available to send email to the quota violator. The email address can be based on **Directory Owner** or **Directory Name**.

Alert display colors

In the stand-alone GUI, status symbols for alerts are shown in the colors that are listed for the alert action.

When you have active alerts that are shown within the nodes under the **IBM Tivoli Storage Productivity Center** > **Alerting** > **Alert Log** node in the navigation tree, the indicator next to those nodes are displayed as green or red depending on their corresponding alerts. The following table shows the colors that are generated for the subnodes by their different alerts:

Table 144. Alert colors

Alert	Color
Probe failed	Red
RAM increase	Green
RAM decrease	Red
Memory increase	Green
Memory decrease	Red
New disk	Green
Missing disk	Red
Disk defect	Red
Disk failure predicted	Red
New file system	Green
File system reconfigured	Red
Scan failed	Red
Missing file system	Red
File system free space	Red
File system free Inodes	Red
Computer unreachable	Red
New directory	Red
Missing directory	Red
Directory quota violation	Red
File system quota violation	Red
Computer quota violation	Red
Network quota violation	Red
Constraint violation	Red

Table 144. Alert colors (continued)

Alert	Color
Failed job	Red
Disk defect quota	Red
Ping failed	Red
Discovery failed	Red
Aggregator failed	Red
File system quota failed	Red
Computer quota failed	Red
Network quota failed	Red
Computer discovered	Red
Report failed	Red
DB probe failed	Red
DB table space quota violation	Red
DB instance quota violation	Red
DB net quota violation	Red
DB table space quota failed	Red
DB instance quota failed	Red
DB net quota failed	Red
Scan failed	Red
Archive log overflow	Red
New table space	Red
Table space free space	Red
Table space offline	Red
Table space fragmented	Red
Table space extent size	Red
Dropped table space	Red
Table fragmented	Red
Table maximum extents	Red
Table size	Red
Table chained rows	Red
Table wasted space	Red
Dropped table	Red
NAS discovery failed	Red
NAS computer discovered	Green
Script failed	Red
NetApp quota job failed	Red
Device free space less	Red
Device free space greater	Red
Dropped device	Red
New device	Red
DB log free space	Red

Table 144. Alert colors (continued)

Alert	Color
DB dump time	Red
Disk array discovered	Green
Disk array missing	Red
Filer missing	Red
Cache increase	Green
Cache decrease	Red
Automatic file system extension	Red
Stopped automatic file system extension	Red
Extend job failed	Red
Upgrade failed	Red
Added cluster resource group	Green
Removed cluster resource group	Green
Moved cluster resource group	Green
DB report failed	Red
File backup failed	Red
Hypervisor discovered	Green
Hypervisor missing	Red
Virtual machine discovered	Green
Virtual machine missing	Red

Alerts and monitoring groups

Use alerts to notify you when conditions on monitored resources are detected by Tivoli Storage Productivity Center. Include resources in predefined monitoring groups to automatically associate a set of alerts with those resources.

Alerts are triggered by conditions that are detected during data collection and event processing. For storage systems such as SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, or the XIV system, events are polled every minute from the resource. For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to Tivoli Storage Productivity Center when conditions change on the resource.

IBM Tivoli Storage Productivity Center includes predefined groups that are associated with sets of alerts. The group that you select for a resource determines what conditions are checked for those resources.

Tips:

- If you want to include a resource in a group, you can select a user-defined group rather than a predefined group. Any alerts that are associated with a user-defined group are automatically applied to the resource.
- You are not required to include a resource in a group to associate alerts with it. You can associate alerts with a resource at any time. Monitoring groups are intended to help you assign a set of alerts to a number of related resources at the same time.

- You can modify predefined monitoring groups and their associated alerts at any time.

Storage Subsystems

Table 145. Predefined monitoring groups and alerts for storage subsystems

Group name	Alert conditions that are checked for resources in the group during data collection[1, 2]
Subsystem Basic Group	Subsystem Discovered Disk Not Found Storage Subsystem Not Found Subsystem Status Change Offline Subsystem Status Change Online Volume Status Change Offline Volume Status Change Online Pool Status Change Pool Discovered Pool Status Change Offline Pool Status Change Online Volume Not Found File System Is Missing Cluster Was Rediscovered Cluster Was Removed NSD Is Missing NSD State Changed
Subsystem Standard Group	Subsystem Discovered Disk Not Found Storage Subsystem Not Found Subsystem Status Change Offline Subsystem Status Change Online Backend Controller Status Change Offline Backend Controller Status Change Online Volume Status Change Offline Volume Status Change Online Volume Capacity Change Pool Status Change Pool Discovered Pool Status Change Offline Pool Status Change Online Pool Capacity Change Volume Not Found File System Is Missing File System Reconfigured Cluster Was Rediscovered Cluster Was Removed Export Is Missing Export State Changed NSD Is Missing NSD State Changed

Table 145. Predefined monitoring groups and alerts for storage subsystems (continued)

Group name	Alert conditions that are checked for resources in the group during data collection[1, 2]
Subsystem Advanced Group	Subsystem Discovered Cache Increased Cache Decreased Disk Not Found Storage Subsystem Not Found Subsystem Status Change Offline Subsystem Status Change Online Subsystem Version Change Subsystem Property Change Subsystem Allocated Capacity Change Subsystem Available Capacity Change Subsystem Backend Capacity Change Backend Controller Status Change Offline Backend Controller Status Change Online Volume Status Change Offline Volume Status Change Online Volume Capacity Change Pool Status Change Pool Discovered Pool Status Change Offline Pool Status Change Online Pool Capacity Change Node Status Change Offline Node Status Change Online Node Status Change Node Discovered Datapath State Change Volume Not Found File System Is Missing File System Reconfigured Cluster Was Rediscovered Cluster Was Removed Export Is Missing Export Path Changed Export Protocol Changed Export State Changed NSD Is Missing NSD State Changed Node State Changed Node Status Changed to Error Node Status Changed to Offline Node Status Changed to Online Node Status Changed to Warning
Default Storage Subsystem Group[3]	No alert conditions are associated with this group

Table 145. Predefined monitoring groups and alerts for storage subsystems (continued)

Group name	Alert conditions that are checked for resources in the group during data collection[1, 2]
Notes: 1. For descriptions of the alert triggering conditions for storage systems, see Triggering conditions for storage subsystem alerts. 2. Performance alerts for storage subsystems are not associated with groups. Performance alerts are alerts that contain a "threshold" triggering condition. If you want to associate a performance alert with a storage subsystem, you must create an alert with a threshold triggering condition and then individually include the storage subsystem in that alert. For information about how to create storage subsystem alerts, see Creating storage subsystem alerts. 3. The Default Storage Subsystem Group is not automatically associated with alerting conditions. However, you can manually associate alerts with it at any time.	

Fabrics

Table 146. Predefined monitoring groups and alerts for fabrics

Group name	Alert conditions that are checked for resources in the group during data collection[1]
Fabric Basic Group	Fabric State Change Fabric to Switch Change Switch discovered Switch State Change Switch Status Change Offline Switch Status Change Online
Fabric Standard Group	Fabric State Change Fabric to Switch Change New Connection Connection State Change Switch discovered Switch State Change Switch Status Change Offline Switch Status Change Online Switch Blade Change Switch Blade Change Offline Switch Blade Change Online
Fabric Advanced Group	Fabric State Change Fabric to Switch Change Zone State Change Zone Set State Change Zone Set Discovered Zone Set to Zone Change Zone Alias to Member Change New Connection Connection State Change Switch discovered Switch State Change Switch Status Change Offline Switch Status Change Online Switch Version Change Switch Blade Change Switch Blade Change Offline Switch Blade Change Online
Default Fabric Group[2]	No alert conditions are associated with this group

Table 146. Predefined monitoring groups and alerts for fabrics (continued)

Group name	Alert conditions that are checked for resources in the group during data collection[1]
Notes: 1. For descriptions of the alert triggering conditions for fabrics, see Triggering conditions for fabric alerts. For descriptions of the alert triggering conditions for switches, see Triggering conditions for switch alerts. 2. The Default Fabric Group is not automatically associated with a schedule or alerting conditions. However, you can manually include this group in a schedule or associate it with alerts at any time.	

Computers

Table 147. Predefined monitoring groups alerts for computers

Group name / Template name	Alert conditions that are checked for resources in the group during data collection[1]
Computer Basic Group	Disk Not Found New Disk Defect Found Disk Failure Predicted Computer Unreachable Computer Discovered Computer Status Change Offline Computer Status Change Online
Computer Standard Group	New Disk Detected Disk Not Found New Disk Defect Found Grown Disk Defects Exceed 5 Disk Failure Predicted New Filesystem Detected Computer Unreachable Computer Discovered Computer Status Change Offline Computer Status Change Online HBA Driver Version Change HBA Firmware Version Change
Computer Advanced Group	RAM Increased RAM Decreased Virtual Memory Increased Virtual Memory Decreased New Disk Detected Disk Not Found New Disk Defect Found Grown Disk Defects Exceed 5 Disk Failure Predicted New Filesystem Detected Cluster Resource Group Added Cluster Resource Group Removed Cluster Resource Group Moved Computer Unreachable Computer Discovered Computer Status Change Offline Computer Status Change Online Computer Property Change HBA Driver Version Change HBA Firmware Version Change
Default Computer Group[2]	No alert conditions are associated with this group

Table 147. Predefined monitoring groups alerts for computers (continued)

Group name / Template name	Alert conditions that are checked for resources in the group during data collection[1]
Notes: 1. For descriptions of the alert triggering conditions for computers, see Triggering conditions for computer alerts. 2. The Default Computer Group is not automatically associated with alerting conditions. However, you can manually associate it with alerts at any time.	

Performance data

Use the IBM Tivoli Storage Productivity Center to view the performance data of devices in your environment.

Performance thresholds

Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

You can monitor the performance of your enterprise by creating alerts on performance thresholds for switches and storage subsystems. By creating alerts that are triggered by performance thresholds, you can be informed about performance issues in your enterprise. Threshold events tell you when a component has fallen outside of the user-defined values. For example, when a threshold value has reached critical stress.

Threshold boundaries

You can establish your boundaries for the normal expected subsystem performance when defining storage subsystem alerts for performance threshold events. When the collected performance data samples fall outside the range you have set, you are notified of this threshold violation so you are aware of the potential problem. The upper-threshold boundaries are **Critical Stress** and **Warning Stress**. The lower-threshold boundaries are **Warning Idle** and **Critical Idle**. Usually you will want the stress boundaries to be high numbers and the idle to be low numbers. The exception to this rule is **Cache Holding Time Threshold**, where you want the stress numbers to be low and the idle numbers to be high.

If you do not want to be notified of threshold violations for any boundaries, you can leave the boundary field blank and the performance data will not be checked against any value. For example, if you leave the **Critical Idle** and **Warning Idle** fields blank, no alerts will be sent for any idle conditions.

The **Ignore triggering condition when the sequential I/O percentage exceeds** check box is active only for the triggering condition **Disk Utilization Percentage Threshold**. It is a filter condition. The default is 80%.

The **Ignore triggering condition when the Back-end Read I/O Rate is less than** check box only applies to the **Back-end Read Response Time** and **Back-end Read Queue Time** thresholds.

The **Ignore triggering condition when the Back-end Write I/O Rate is less than** check box only applies to the **Back-end Write Response Time** and **Back-end Write Queue Time** thresholds.

The **Ignore triggering condition when the Total Back-end I/O Rate is less than** check box only applies to the **Overall Back-end Response Time** threshold.

The **Ignore triggering condition when the Total I/O Rate is less than** check box only applies to the **Non-preferred Node Usage Percentage** threshold.

The **Ignore triggering condition when the Write-cache Delay I/O Rate is less than** check box only applies to the **Write-cache Delay Percentage** threshold.

Setting the thresholds

Only a few thresholds have defaults and on the other thresholds you will have to determine the best values for stress, idle, critical, and warning values so you can derive the maximum benefit without generating too many false alerts. Since suitable stress thresholds are highly dependant on the type of workload you are running, your exact hardware configuration, the number of physical disks, exact model numbers, and so forth, there are no easy or standard default rules.

One of the best approaches is to monitor your performance for a number of weeks and, using this historical data, determine reasonable values for each threshold setting. Once that is done you can fine tune these settings to minimize the number of false alerts.

Related concepts:

"Alerting" on page 341
Determining when and how you are alerted to conditions and violations of performance thresholds within your storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. You can set up Tivoli Storage Productivity Center so that it examines the data about your resources for the conditions and performance thresholds that you specify.

"Monitoring the performance of resources" on page 383
Tivoli Storage Productivity Center can collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations. Metrics can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage. Threshold violations alert you when the performance of a monitored resource falls outside of a specified range.

Related tasks:

"Archiving and deleting files" on page 103
In this scenario, you are shown how to increase storage space by archiving and deleting files that meet the criteria, such as file size and file usage, that you set.

"Adding resources" on page 257
You must add resources to Tivoli Storage Productivity Center before you can collect data, generate reports, and manage storage that is related to those resources.

Related reference:

"Controller thresholds" on page 1696
The following table lists and describes the Controller thresholds.

"Port thresholds" on page 1697
Port thresholds are used to set limits for such things as bandwidth utilization, data rates, and I/O operations.

"Array thresholds" on page 1692
The following table lists and describes the Array thresholds.

Array thresholds

The following table lists and describes the Array thresholds.

Table 148. Array thresholds

Threshold (Metric)	Device/Component Type	Description
Array Thresholds		
Disk Utilization Percentage	TotalStorage Enterprise Storage Server, DS6000, DS8000 array	Sets thresholds on the approximate utilization percentage of the arrays in a particular subsystem; for example, the average percentage of time that the disks associated with the array were busy. The Disk Utilization metric for each array is checked against the threshold boundaries for each collection interval. This threshold is enabled by default for TotalStorage Enterprise Storage Server systems and disabled by default for others. The default threshold boundaries are 80%, 50%, -1, -1. For DS6000 and DS8000 subsystems, this threshold applies only to those ranks which are the only ranks in their associated extent pool.
Total Back-end I/O Rate	SAN Volume Controller MDisk group and MDisk, TotalStorage Enterprise Storage Server, DS6000, DS8000 array	Sets thresholds on the average number of I/O operations per second for array and MDisk read and write operations. The Total I/O Rate metric for each array or MDisk is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Total Back-end Data Rate	SAN Volume Controller MDisk group and MDisk, TotalStorage Enterprise Storage Server, DS6000, DS8000 array	Sets thresholds on the average number of MB per second that were transferred for array and MDisk read and write operations. The Total Data Rate metric for each array or MDisk is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Back-end Read Response Time	SAN Volume Controller MDisk, TotalStorage Enterprise Storage Server, DS6000, DS8000 array	Sets thresholds on the average number of milliseconds that it took to service each array and MDisk read operation. The Back-end Read Response Time metric for each array or MDisk is checked against the threshold boundaries for each collection interval. Though this threshold is disabled by default, suggested boundary values of 35,25,-1,-1 are pre-populated. A filter is available for this threshold which will ignore any boundary violations if the Back-end Read I/O Rate is less than a specified filter value. The pre-populated filter value is 5.

Table 148. Array thresholds (continued)

Threshold (Metric)	Device/Component Type	Description
Back-end Write Response Time	SAN Volume Controller MDisk, TotalStorage Enterprise Storage Server, DS6000, DS8000 array	Sets thresholds on the average number of milliseconds that it took to service each array and MDisk write operation. The Back-end Write Response Time metric for each array or MDisk is checked against the threshold boundaries for each collection interval. Though this threshold is disabled by default, suggested boundary values of 120,80,-1,-1 are pre-populated. A filter is available for this threshold which will ignore any boundary violations if the Back-end Write I/O Rate is less than a specified filter value. The pre-populated filter value is 5.
Overall Back-end Response Time	SAN Volume Controller MDisk	Sets thresholds on the average number of milliseconds that it took to service each MDisk I/O operation, measured at the MDisk level. The Total Response Time (external) metric for each MDisk is checked against the threshold boundaries for each collection interval. This threshold is disabled by default. A filter is available for this threshold which will ignore any boundary violations if the Total Back-end I/O Rate is less than a specified filter value. The pre-populated filter value is 10.
Back-end Read Queue Time	SAN Volume Controller MDisk	Sets thresholds on the average number of milliseconds that each read operation spent on the queue before being issued to the back-end device. The Back-end Read Queue Time metric for each MDisk is checked against the threshold boundaries for each collection interval. Though this threshold is disabled by default, suggested boundary values of 5,3,-1,-1 are pre-populated. A filter is available for this threshold which will ignore any boundary violations if the Back-end Read I/O Rate is less than a specified filter value. The pre-populated filter value is 5. Violation of these threshold boundaries means that the SVC deems the MDisk to be overloaded. There is a queue algorithm that determines the number of concurrent I/O operations that the SVC will send to a given MDisk. If there is any queuing (other than during a backup process) then this suggests performance can be improved by resolving the queuing issue.

Table 148. Array thresholds (continued)

Threshold (Metric)	Device/Component Type	Description
Back-end Write Queue Time	SAN Volume Controller MDisk	Sets thresholds on the average number of milliseconds that each write operation spent on the queue before being issued to the back-end device. The Back-end Write Queue Time metric for each MDisk is checked against the threshold boundaries for each collection interval. Though this threshold is disabled by default, suggested boundary values of 5,3,-1,-1 are pre-populated. A filter is available for this threshold which will ignore any boundary violations if the Back-end Read I/O Rate is less than a specified filter value. The pre-populated filter value is 5. Violation of these threshold boundaries means that the SVC deems the MDisk to be overloaded. There is a queue algorithm that determines the number of concurrent I/O operations that the SVC will send to a given MDisk. If there is any queuing (other than during a backup process) then this suggests performance can be improved by resolving the queuing issue.
Peak Back-end Write Response Time	SAN Volume Controller Node	Sets thresholds on the peak (worst) response time among all MDisk write operations by a node. The Back-end Peak Write Response Time metric for each node is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundary values of 30000,10000,-1,-1. Violation of these threshold boundaries means that the SVC cache is having to “partition-limit” for a given MDisk group. The de-staged data from the SVC cache for this MDisk group is causing the cache to fill up (writes are being received faster than they can be de-staged to disk). If delays reach 30 seconds or more, then the SVC will switch into “short-term mode” where writes are no longer cached for the MDisk Group.

Table 148. Array thresholds (continued)

Threshold (Metric)	Device/Component Type	Description
Port to Local Node Send Response Time	SAN Volume Controller Node	Sets thresholds on the average number of milliseconds it took to service each send operation to another node in the local SVC cluster. The Port to Local Node Send Response Time metric for each node is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundary values of 3,1.5,-1,-1. Violation of these threshold boundaries means that it is taking too long to send data between nodes (on the fabric), and suggests that there is either congestion around these FC ports, or an internal SVC microcode problem.
Port to Local Node Receive Response Time	SAN Volume Controller Node	Sets thresholds on the average number of milliseconds it took to service each receive operation from another node in the local SVC cluster. The Port to Local Node Receive Response Time metric for each node is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundary values of 1,0.5,-1,-1. Violation of these threshold boundaries means that it is taking too long to send data between nodes (on the fabric), and suggests that there is either congestion around these FC ports, or an internal SVC microcode problem.
Port to Local Node Send Queue Time	SAN Volume Controller Node	Sets thresholds on the average number of milliseconds that each send operation issued to another node in the local SVC cluster spent on the queue before being issued. The Port to Local Node Send Queue Time metric for each node is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundary values of 2,1,-1,-1. Violation of these threshold boundaries means that the node has to wait too long to send data to other nodes (on the fabric), and suggests congestion on the fabric.

Table 148. Array thresholds (continued)

Threshold (Metric)	Device/Component Type	Description
Port to Local Node Receive Queue Time	SAN Volume Controller Node	Sets thresholds on the average number of milliseconds that each receive operation issued to another node in the local SVC cluster spent on the queue before being issued. The Port to Local Node Receive Queue Time metric for each node is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundary values of 1,0.5,-1,-1. Violation of these threshold boundaries means that the node has to wait too long to receive data from other nodes (on the fabric), and suggests congestion on the fabric.

Controller thresholds

The following table lists and describes the Controller thresholds.

Table 149. Controller thresholds

Threshold (Metric)	Device/Component Type	Description
Controller Thresholds		
Total I/O Rate (overall)	TotalStorage Enterprise Storage Server, DS6000, DS8000 controller SAN Volume Controller I/O group	Sets threshold on the average number of I/O operations per second for read and write operations, for the subsystem controllers (clusters) or I/O groups. The Total I/O Rate metric for each controller or I/O group is checked against the threshold boundaries for each collection interval. These thresholds are disabled by default.
Total Data Rate	TotalStorage Enterprise Storage Server, DS6000, DS8000 controller SAN Volume Controller I/O group	Sets threshold on the average number of MB per second for read and write operations for the subsystem controllers (clusters) or I/O groups. The Total Data Rate metric for each controller or I/O group is checked against the threshold boundaries for each collection interval. These thresholds are disabled by default.
NVS Full Percentage	TotalStorage Enterprise Storage Server, DS6000, DS8000 controller	Sets thresholds on the percentage of time that NVS space constraints caused I/O operations to be delayed, for the subsystem controllers (clusters). The NVS Full Percentage metric for each controller is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries of 10, 3, -1, -1.

Table 149. Controller thresholds (continued)

Threshold (Metric)	Device/Component Type	Description
Cache Holding Time	TotalStorage Enterprise Storage Server, DS6000, DS8000 controller	Sets thresholds on the average cache holding time, in seconds, for I/O data in the subsystem controllers (clusters). Shorter time periods indicate adverse performance. The Cache Holding Time metric for each controller is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries of 30, 60, -1, -1.
Write-cache Delay Percentage	TotalStorage Enterprise Storage Server, DS6000, DS8000 controller, SAN Volume Controller node	Sets thresholds on the percentage of I/O operations that were delayed due to write-cache space constraints. This metric for each controller or node is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries of 10, 3, -1, -1. In addition, a filter is available for this threshold which will ignore any boundary violations if the Write-cache Delay I/O Rate is less than a specified filter value. The pre-populated filter value is 10 I/Os per second.
Non-Preferred Node Usage Percentage	SAN Volume Controller I/O group	Sets thresholds on the Non-Preferred Node Usage Percentage of an I/O group. This metric of each I/O group is checked against the threshold boundaries at each collection interval. his threshold is disabled by default. In addition, a filter is available for this threshold which will ignore any boundary violations if the Total I/O Rate of the I/O group is less than a specified filter value.

Port thresholds

Port thresholds are used to set limits for such things as bandwidth utilization, data rates, and I/O operations.

Table 150. Port thresholds

Threshold (Metric)	Device/Component Type	Description
Port Thresholds		
Total Port I/O Rate	TotalStorage Enterprise Storage Server, DS6000, DS8000 switch port, XIV	Sets thresholds for ports on the average number of I/O operations or packets per second for send and receive operations. The Total I/O Rate metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.

Table 150. Port thresholds (continued)

Threshold (Metric)	Device/Component Type	Description
Total Port Data Rate	TotalStorage Enterprise Storage Server, DS6000, DS8000 switch port, XIV	Sets thresholds for ports on the average number of MB per second for send and receive operations. The Total Data Rate metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Overall Port Response Time	TotalStorage Enterprise Storage Server, DS6000, DS8000	Sets thresholds for ports on the average number of milliseconds that it takes to service each send and receive I/O operation. The Total Response Time metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Error Frame Rate	DS8000	Sets thresholds on the average number of frames per second received in error by ports. The Error Frame Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default.
Link Failure Rate	DS8000 SAN Volume Controller	Sets thresholds on the average number of link errors per second for ports. The Link Failure Rate metric for each port is checked against the threshold boundary for each collection interval. The default boundaries are 0.0030,-1,-1,-1 cnt/s. This threshold is disabled by default.
Total Port Data Rate	Switch port	Sets thresholds on the critical and warning data rates for stress and idle in MB per second. The Total Port Data Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default.

Table 150. Port thresholds (continued)

Threshold (Metric)	Device/Component Type	Description
Total Port Packet Rate	Switch port	Sets thresholds on critical and warning data rates for stress and idle conditions in packets per second. For example, a critical stress or warning stress condition occurs when the upper-threshold boundary for the packet rate of a switch is detected. A critical idle or warning idle condition occurs when the lower-threshold boundary for the packet rate of a switch is detected. The Total Port Packet Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default.
Port Send Utilization Percentage	TotalStorage Enterprise Storage Server, DS6000, DS8000	Sets thresholds on the average amount of time that ports are busy sending data. The metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Port Receive Utilization Percentage	TotalStorage Enterprise Storage Server, DS6000, DS8000	Sets thresholds on the average amount of time that ports are busy receiving data. The metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Port Send Bandwidth Percentage	TotalStorage Enterprise Storage Server, DS8000, SAN Volume Controller switch port, XIV	Sets thresholds on the average port bandwidth utilization percentage for send operations. The Port Send Utilization Percentage metric is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries 85,75,-1,-1.

Table 150. Port thresholds (continued)

Threshold (Metric)	Device/Component Type	Description
Port Receive Bandwidth Percentage	TotalStorage Enterprise Storage Server, DS8000, SAN Volume Controller switch port, XIV	Sets thresholds on the average port bandwidth utilization percentage for receive operations. The Port Send Utilization Percentage metric is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries 85,75,-1,-1.
CRC Error Rate	DS8000 port, SAN Volume Controller switch port	Sets thresholds on the average number of frames that are received in which the cyclic redundancy check (CRC) in a frame does not match the CRC computed by the receiver. The CRC Error Rate metric for each port is checked against the threshold boundary for each collection interval. The default boundaries are 0.033, 0.01, -1, -1 cnt/s. This threshold is disabled by default.
Invalid Transmission Word Rate	DS8000, SAN Volume Controller switch port	Sets thresholds on the average number of bit errors that are detected on a port. The Invalid Transmission Word Rate metric for each port is checked against the threshold boundary for each collection interval. The default boundaries are 0.033, 0.01, -1, -1 cnt/s. This threshold is disabled by default.
Zero Buffer - Buffer Credit Timer	SAN Volume Controller	Sets thresholds on the number of microseconds that a port was unable to send frames because of a lack of buffer credit since the last node reset. The Zero Buffer-Buffer Credit Timer metric for each port is checked against the threshold boundary for each collection interval. The default boundaries are 198000000,75000000,-1,-1 us. This threshold is disabled by default.

Performance metrics

Tivoli Storage Productivity Center can collect information about the performance of monitored storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks.

You can view performance metrics in the web-based GUI and stand-alone GUI at the following locations:

Resource pages in the web-based GUI
To access a view of performance metrics for a resource, go to the following pages in the web-based GUI:

- Server details page > Related Resources
- Storage systems page, Storage system details page
- Switches page, Switch details page

Reports in the web-based GUI
To view detailed performance reports for multiple resources in the web-based GUI, hover the mouse pointer over **Reporting** in the navigation pane and select an option.

Reports in the stand-alone GUI
To view detailed performance reports for multiple resources in the stand-alone GUI, run the reports that are available under the following nodes in the navigation tree:

- **Disk Manager** > **Reporting**
- **Fabric Manager** > **Reporting**
- **Tivoli Storage Productivity Center** > **System Reports**

Common information

The following table contains information that is common to performance reports in the stand-alone GUI.

Column	Description
Time	Date and time that the data was collected
Interval	Size of the sample interval in seconds. You can specify a minimum interval length of 5 minutes and a maximum interval length of 60 minutes for the following storage systems: • TotalStorage Enterprise Storage Server • DS6000 • DS8000 • XIV system For SAN Volume Controller models earlier than V4.1, you can specify a minimum interval length of 15 minutes and a maximum interval length of 60 minutes. For SAN Volume Controller models V4.1 and later, you can specify a minimum interval length of 5 minutes, and a maximum interval length of 60 minutes.

Metrics for ESS and DS6000 and DS8000 storage systems

When you view metrics for the ESS and DS6000 and DS8000 storage systems, you must take into account the following differences between IBM Tivoli Storage Productivity Center reports and the native reports of those systems:

- Tivoli Storage Productivity Center reports display port performance metrics as send and receive metrics, such as Send Data Rate and Receive Data Rate.
- Reports that are included with storage systems (for example, reports based on data that is collected by the DS CLI) display port performance metrics as read and write metrics, such as Byteread and Bytewrite.

When a host runs a read operation, the DS port sends data to the host. Therefore, "read" metrics in DS6000 and DS8000 reports correspond to "send" metrics in Tivoli Storage Productivity Center reports. When a host runs a write operation, DS6000 and DS8000 ports receive data from the host. Therefore, "write" metrics in DS6000 and DS8000 reports correspond to "receive" metrics in Tivoli Storage Productivity Center reports. When you view port Peer-to-Peer Remote Copy (PPRC) performance metrics, you must take into account the following additional differences between Tivoli Storage Productivity Center reports and native reports for storage systems:

- Metrics for PPRC reads in storage system native reports are represented as PPRC receives in Tivoli Storage Productivity Center (reads = receives).
- Metrics for PPRC writes in storage system native reports are represented as PPRC sends in Tivoli Storage Productivity Center (writes = sends).

Restriction: Tivoli Storage Productivity Center does not calculate volume-based metrics if thin-provisioned volumes are allocated in an extent pool that consists of multiple ranks. In this case, the columns for volume-based metrics display the value **N/A** in the Storage Subsystem Performance > By Array report for the arrays that are associated with that extent pool. However, if no thin-provisioned volumes are allocated in a multi-rank extent pool, or if the space efficient volumes are allocated in an extent pool that consists of a single rank, then this limitation does not apply and all volume-based metrics are displayed in the By Array report.

Metrics for XIV systems

To distinguish between metrics that were introduced in different versions of the XIV systems, the following conventions are used:

- [2] is displayed next to metrics that are available in the XIV system version 10.2.2 or later.
- [3] is displayed next to metrics that are available in the XIV system version 10.2.4 or later.

For example:

- The Read I/O Rate (overall) metric is available for the XIV systems version 10.2.2 and later. In the Resources: components column of the list of metrics, the entry for Read I/O Rate (overall) is displayed like this example: XIV[2]: volume, module, storage system
- The Small Transfers Response Time metric is available for the XIV systems version 10.2.4 and later. In the Resources: components column of the list of metrics, the entry for Small Transfers Response Time is displayed in the following format: XIV[3]: volume, module, storage system

Restriction: XIV systems do not track performance statistics for volumes that were never used. Because there are no performance statistics, Tivoli Storage Productivity Center does not display metrics for these volumes and their related components.

Metrics for SAN Volume Controller

Metrics that apply to SAN Volume Controller also apply to Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified. For example, the Read I/O Rate (overall) metric is available for SAN Volume Controller and Storwize storage system volumes, even though only SAN Volume Controller is shown in the Resources: components column of the following lists.

Volume metrics

The following table describes the metrics for the volumes in storage systems.

Metric	Resources: components[1]	Description
I/O Rates		
Read I/O Rate (normal)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of I/O operations per second for nonsequential read operations for a component over a specified time interval.
Read I/O rate (sequential)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of I/O operations per second for sequential read operations for a component over a specified time interval.
Read I/O Rate (overall)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller[1]: volume, pool, node, I/O group, host connection, storage system SMI-S BSP: volume, controller, storage system XIV[2]: volume, pool, module, host connection, storage system	Average number of I/O operations per second for both sequential and nonsequential read operations for a component over a specified time interval.
Write I/O Rate (normal)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of I/O operations per second for nonsequential write operations for a component over a specified time interval.
Write I/O Rate (sequential)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of I/O operations per second for sequential write operations for a component over a specified time interval.

Metric	Resources: components[1]	Description
Write I/O Rate (overall)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system SMI-S BSP: volume, controller, storage system XIV[2]: volume, pool, module, host connection, storage system	Average number of I/O operations per second for both sequential and nonsequential write operations for a component over a specified time interval.
Total I/O Rate (normal)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of I/O operations per second for nonsequential read and write operations for a component over a specified time interval.
Total I/O Rate (sequential)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of I/O operations per second for sequential read and write operations for a component over a specified time interval.
Total I/O Rate (overall)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system SMI-S BSP: volume, controller, storage system XIV[2]: volume, pool, module, host connection, storage system	Average number of I/O operations per second for both sequential and nonsequential read and write operations for a component over a specified time interval.
Global Mirror Write I/O Rate	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of write operations per second issued to the Global Mirror secondary site for a component over a specified time interval.

Metric	Resources: components[1]	Description
Global Mirror Overlapping Write Percentage	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average percentage of write operations that are issued by the Global Mirror primary site and are serialized overlapping writes for a component over a specified time interval. For SAN Volume Controller 4.3.1 and later, some overlapping writes are processed in parallel (are not serialized) and are excluded. For earlier versions of SAN Volume Controller, all overlapping writes were serialized.
Global Mirror Overlapping Write I/O Rate	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of serialized overlapping write operations per second encountered by the Global Mirror primary site for a component over a specified time interval. For SAN Volume Controller 4.3.1 and later, some overlapping writes are processed in parallel (are not serialized) and are excluded. For earlier versions of SAN Volume Controller, all overlapping writes are serialized.
HPF Read I/O Rate	DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of read operations per second that were issued by the High Performance FICON (HPF) feature of the storage system for a component over a specified time interval.
HPF Write I/O Rate	DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of write operations per second that were issued by the High Performance FICON (HPF) feature of the storage system for a component over a specified time interval.
Total HPF I/O Rate	DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of read and write operations per second that were issued by the High Performance FICON (HPF) feature of the storage system for a component over a specified time interval.
HPF I/O Percentage	DS8000: volume, RAID array, pool, controller, host connection, storage system	The percentage of all I/O operations that were issued by the High Performance FICON (HPF) feature of the storage system for a component over a specified time interval.
PPRC Transfer Rate	ESS/DS6000/DS8000: volume, RAID array, pool, controller, host connection, storage system	Average number of track transfer operations per second for PPRC usage for a component over a specified time interval. This metric shows the activity for the source of the PPRC relationship, but shows no activity for the target.
Small Transfers I/O Percentage	XIV[3]: module, pool, volume, host connection, storage system	Percentage of I/O operations over a specified interval. Applies to data transfer sizes that are <= 8 KB.
Medium Transfers I/O Percentage	XIV[3]: module, pool, volume, host connection, storage system	Percentage of I/O operations over a specified interval. Applies to data transfer sizes that are > 8 KB and <= 64 KB.
Large Transfers I/O Percentage	XIV[3]: module, pool, volume, host connection, storage system	Percentage of I/O operations over a specified interval. Applies to data transfer sizes that are > 64 KB and <= 512 KB.
Very Large Transfers I/O Percentage	XIV[3]: module, pool, volume, host connection, storage system	Percentage of I/O operations over a specified interval. Applies to data transfer sizes that are > 512 KB.
Cache hit percentages		

Metric	Resources: components[1]	Description
Read Cache Hits Percentage (normal)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Percentage of cache hits for nonsequential read operations for a component over a specified time interval.
Read Cache Hits Percentage (sequential)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Percentage of cache hits for sequential read operations for a component over a specified time interval.
Read Cache Hits Percentage (overall)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, node, I/O group, host connection, storage system SMI-S BSP: volume, controller, storage system XIV[1]: volume, pool, module, host connection, storage system	Percentage of cache hits for both sequential and nonsequential read operations for a component over a specified time interval.
Write Cache Hits Percentage (normal)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Percentage of cache hits for nonsequential write operations for a component over a specified time interval.
Write Cache Hits Percentage (sequential)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Percentage of cache hits for sequential write operations for a component over a specified time interval.
Write Cache Hits Percentage (overall)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, node, I/O group, host connection, storage system SMI-S BSP: volume, controller, storage system XIV[2]: volume, pool, module, host connection, storage system	Percentage of cache hits for both sequential and nonsequential write operations for a component over a specified time interval.

Metric	Resources: components[1]	Description
Total Cache Hits Percentage (normal)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Percentage of cache hits for nonsequential read and write operations for a component over a specified time interval.
Total Cache Hits Percentage (sequential)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system	Percentage of cache hits for sequential read and write operations for a component over a specified time interval.
Total Cache Hits Percentage (overall)	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, node, I/O group, host connection, storage system SMI-S BSP: volume, controller, storage system XIV: volume, pool, module, host connection, storage system	Percentage of cache hits for both sequential and nonsequential read and write operations for a component over a specified time interval.
Readahead Percentage of Cache Hits	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Percentage of all read cache hits, which occurred on prestaged data.
Dirty Write Percentage of Cache Hits	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Percentage of all write cache hits, which occurred on already dirty data in the cache.
Read Data Cache Hit Percentage	XIV[3]: module, pool, volume, host connection, storage system	Percentage of read data that was read from the cache over a specified time interval.
Write Data Cache Hit Percentage	XIV[3]: module, pool, volume, host connection, storage system	Percentage of write data that was written to the cache over a specified time interval.
Total Data Cache Hit Percentage	XIV[3]: module, pool, volume, host connection, storage system	Percentage of all data that was read from or written to the cache for a component over a specified time interval.

Metric	Resources: components[1]	Description
SSD Read Cache Hit Percentage	XIV[3]: module, pool, volume, host connection, storage system	The percentage of all read operations for a component that accessed SSD cache memory during a specified time interval. This metric is only available for XIV systems that use SSD drives as an additional layer of caching. The value for this metric is also included in the value for the Read Cache Hit Percentage (overall) metric.
SSD Read Data Cache Hit Percentage	XIV[3]: module, pool, volume, host connection, storage system	The percentage of all read data for a component that was read from SSD cache memory during a specified time interval. This metric is only available for XIV systems that use SSD drives as an additional layer of caching. The value for this metric is also included in the value for the Read Data Cache Hit Percentage (overall) metric.
Data rates		
Read Data Rate	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system SMI-S BSP: volume, controller, storage system XIV[2]: volume, pool, module, host connection, storage system	Average number of megabytes (2^20 bytes) per second that were transferred for read operations for a component over a specified time interval.
Write Data Rate	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system SMI-S BSP: volume, controller, storage system XIV[2]: volume, pool, module, host connection, storage system	Average number of megabytes (2^20 bytes) per second that were transferred for write operations for a component over a specified time interval.

Metric	Resources: components[1]	Description
Total Data Rate	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system SMI-S BSP: controller, volume, storage system XIV[2]: module, pool, volume, host connection, storage system	Average number of megabytes (2^20 bytes) per second that were transferred for read and write operations for a component over a specified time interval.
Small Transfers Data Percentage	XIV[3]: module, volume, host connection, storage system	Percentage of data that was transferred over a specified interval. Applies to I/O operations with data transfer sizes that are <= 8 KB.
Medium Transfers Data Percentage	XIV[3]: module, volume, host connection, storage system	Percentage of data that was transferred over a specified interval. Applies to I/O operations with data transfer sizes that are > 8 KB and <= 64 KB.
Large Transfers Data Percentage	XIV[3]: module, volume, host connection, storage system	Percentage of data that was transferred over a specified interval. Applies to I/O operations with data transfer sizes that are > 64 KB and <= 512 KB.
Very Large Transfers Data	XIV[3]: module, volume, host connection, storage system	Percentage of data that was transferred over a specified interval. Applies to I/O operations with data transfer sizes that are > 512 KB.
Response times		
Read Response Time	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system SMI-S BSP: volume, controller, storage system XIV[2]: volume, pool, module, host connection, storage system	Average number of milliseconds that it took to service each read operation for a component over a specified time interval.

Metric	Resources: components[1]	Description
Write Response Time	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system SMI-S BSP volume, controller, storage system XIV[2]: volume, pool, module, host connection, storage system	Average number of milliseconds that it took to service each write operation for a component over a specified time interval.
Overall Response Time	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system SMI-S BSP: volume, controller, storage system XIV[2]: volume, pool, module, host connection, storage system	Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.
Peak Read Response Time	SAN Volume Controller: volume, node, I/O group, host connection, storage system	The peak (worst) response time among all read operations.
Peak Write Response Time	SAN Volume Controller: volume, node, I/O group, host connection, storage system	The peak (worst) response time among all write operations.
Global Mirror Write Secondary Lag	SAN Volume Controller: volume, node, I/O group, host connection, storage system	The average number of additional milliseconds it takes to service each secondary write operation for Global Mirror, beyond the time that is required to service primary writes.

Metric	Resources: components[1]	Description
Overall Host Attributed Response Time Percentage	SAN Volume Controller: volume, node, I/O group, host connection, storage system	The percentage of the average response time, both read response time and write response time, that can be attributed to delays from host systems. This metric is provided to help diagnose slow hosts and poorly performing fabrics. The value is based on the time that is taken for hosts to respond to transfer-ready notifications from the SAN Volume Controller nodes (for read) and the time that is taken for hosts to send the write data after the node responds to a transfer-ready notification (for write).
Read Cache Hit Response Time	XIV[3]: volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each read cache hit operation over a specified time interval.
Write Cache Hit Response Time	XIV[3]: volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each write cache hit operation over a specified time interval.
Overall Cache Hit Response Time	XIV[3]: volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each read cache hit operation and each write cache hit operation over a specified time interval.
Read Cache Miss Response Time	XIV[3]: volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each read cache miss operation over a specified time interval.
Write Cache Miss Response Time	XIV[3]: volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each write cache miss operation over a specified time interval.
Overall Cache Miss Response Time	XIV[3]: volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each read cache miss operation and each write cache miss operation over a specified time interval.
Small Transfers Response Time	XIV[3]: volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each I/O operation. Applies to data transfer sizes that are <= 8 KB.
Medium Transfers Response Time	XIV[3]: volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each I/O operation. Applies to data transfer sizes that are > 8 KB and <= 64 KB.
Large Transfers Response Time	XIV[3]: volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each I/O operation. Applies to data transfer sizes that are > 64 KB and <= 512 KB.
Very Large Transfers Response Time	XIV[3]: volume, pool, module, host connection, storage system	Average number of milliseconds that it takes to service each I/O operation. Applies to data transfer sizes that are > 512 KB.

Metric	Resources: components[1]	Description
SSD Read Cache Hit Response Time	XIV[3]: volume, pool, module, host connection, storage system	Average number of milliseconds that was required to service each hit operation of SSD read memory for a component over a specified time interval. This metric is only available for XIV systems that use SSD drives as an additional layer of caching. The value for this metric is also included in the value for the Read Cache Hit Response Time metric.
Transfer sizes		
Read Transfer Size	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system SMI-S BSP: volume, controller, storage system XIV[2]: volume, pool, module, host connection, storage system	Average number of KB per I/O for read operations.
Write Transfer Size	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system SMI-S BSP: volume, controller, storage system XIV[2]: volume, pool, module, host connection, storage system	Average number of KB per I/O for write operations.

Metric	Resources: components[1]	Description
Overall Transfer Size	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, pool, node, I/O group, host connection, storage system SMI-S BSP: volume, controller, storage system XIV[2]: volume, pool, module, host connection, storage system	Average number of KB per I/O for read and write operations.
Write-cache constraints		
Write-Cache Delay Percentage	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, node, I/O group, host connection, storage system	Percentage of I/O operations that were delayed because of write-cache space constraints or other conditions for a component over a specified time interval. (The ratio of delayed operations to total I/Os.)
Write-Cache Delay I/O Rate	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection, storage system SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of I/O operations per second that were delayed because of write-cache space constraints or other conditions for a component over a specified time interval.
Write-Cache Overflow Percentage	SAN Volume Controller: volume, node, I/O group, host connection, storage system	The percentage of write operations that are delayed because there is not enough space in the write cache for a component over a specified time interval.
Write-Cache Overflow I/O Rate	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of write operations per second that are delayed because there is not enough space in the write cache for a component over a specified time interval.
Write-Cache Flush-through Percentage	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Percentage of tracks that is written to disk in flush-through mode for a component over a specified time interval.
Write-Cache Flush-through I/O Rate	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of tracks per second that are written to disk in flush-through mode for a component over a specified time interval.

Metric	Resources: components[1]	Description
Write-Cache Write-through Percentage	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Percentage of tracks that is written to disk in write-through mode for a component over a specified time interval.
Write-Cache Write-through I/O Rate	SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of tracks per second that are written to disk in write-through mode for a component over a specified time interval.
Record mode reads		
Record Mode Read I/O Rate	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection	Average number of I/O operations per second for record mode read operations for a component over a specified time interval.
Record Mode Read Cache Hit Percentage	ESS, DS6000, DS8000: volume, RAID array, pool, controller, host connection	Percentage of cache hits for record mode read operations for a component over a specified time interval.
Cache transfers		
Disk to Cache I/O Rate	ESS/DS6000/DS8000: volume, RAID array, pool, controller, host connection SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of I/O operations (track transfers) per second for disk to cache transfers for a component over a specified time interval.
Cache to Disk I/O Rate	ESS/DS6000/DS8000: volume, RAID array, pool, controller, host connection SAN Volume Controller: volume, node, I/O group, host connection, storage system	Average number of I/O operations (track transfers) per second for cache to disk transfers for a component over a specified time interval.
Miscellaneous computed values		
Cache Holding Time	ESS, DS6000, DS8000: controller, storage system	Average cache holding time, in seconds, for I/O data in this storage system controller (cluster). Shorter time periods indicate adverse performance.
CPU Utilization	SAN Volume Controller: node, I/O group, storage system	Average utilization percentage of the processors.
Non-Preferred Node Usage Percentage	SAN Volume Controller: volume, I/O group, host connection	The overall percentage of I/O performed or data transferred through the non-preferred nodes of the volumes, for a component over a specified time interval.

Metric	Resources: components[1]	Description
Volume Utilization	ESS, DS6000, DS8000: volume SAN Volume Controller: volume SMI-S BSP: volume XIV[2]: volume	The approximate utilization percentage of a volume over a specified time interval (the average percent of time that the volume was busy).
Important: • [1] The "Resource: components" column shows the resources and components that are measured by a metric. Metrics that apply to SAN Volume Controller also apply to Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified. • [2] This metric is available in the XIV system version 10.2.2 or later. • [3] This metric is available in the XIV system version 10.2.4 or later.		

Disk metrics

The following table describes the metrics for the disks on storage systems.

Column	Resources: components[1]	Description
I/O rates		
Backend Read I/O Rate	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system SAN Volume Controller[1]: volume, pool, node, I/O group, storage system	Average number of I/O operations per second for read operations.
Backend Write I/O Rate	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system SAN Volume Controller: volume, pool, node, I/O group, storage system	Average number of I/O operations per second for write operations.
Total Backend I/O Rate	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system SAN Volume Controller: volume, pool, node, I/O group, storage system	Average number of I/O operations per second for read and write operations.
Data rates		

Column	Resources: components[1]	Description
Backend Read Data Rate	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of megabytes (2^20 bytes) that were transferred for read operations.
Backend Write Data Rate	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of megabytes (2^20 bytes) that were transferred for write operations.
Total Backend Data Rate	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of megabytes (2^20 bytes) that were transferred for read and write operations.
Response times		
Backend Read Response Time	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of milliseconds that it took to respond to each read operation. For SAN Volume Controller models, this time is the external response time of the managed disks (MDisks).
Backend Write Response Time	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of milliseconds that it took to respond to each write operation. For SAN Volume Controller models, this time is the external response time of the managed disks (MDisks).
Overall Backend Response Time	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of milliseconds that it took to respond to each I/O operation (read and write). For SAN Volume Controller models, this time is the external response time of the managed disks (MDisks).

Column	Resources: components[1]	Description
Backend Read Queue Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of milliseconds that each read operation was in the queue before it was issued to the back-end resource.
Backend Write Queue Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of milliseconds that each write operation was in the queue before it was issued to the back-end resource.
Overall Backend Queue Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of milliseconds that read and write operations were in the queue before they were issued to the back-end resource.
Peak Back-end Read Response Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	The peak (worst) response time among all read operations for a component over a specified time interval. For SAN Volume Controller, it represents the external response time of the managed disks.
Peak Back-end Write Response Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	The peak (worst) response time among all write operations for a component over a specified time interval. For SAN Volume Controller, it represents the external response time of the managed disks.
Peak Back-end Read Queue Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	The lower bound on the peak (worst) queue time for read operations for a component over a specified time interval. The queue time is the amount of time that the read operation was in the queue before it was issued to the back-end resource.
Peak Back-end Write Queue Time	SAN Volume Controller: managed disk, pool, node, I/O group, storage system	The lower bound on the peak (worst) queue time for write operations for a component over a specified time interval. The queue time is the amount of time that the write operation was in the queue before it was issued to the back-end resource.
Transfer sizes		
Back-end Read Transfer Size	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of KB per I/O for read operations for a component over a specified time interval.
Back-end Write Transfer Size	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of KB per I/O for write operations for a component over a specified time interval.

Column	Resources: components[1]	Description
Overall Back-end Transfer Size	ESS, DS6000, DS8000: rank, RAID array, pool, controller, storage system SAN Volume Controller: managed disk, pool, node, I/O group, storage system	Average number of KB per I/O for read and write operations for a component over a specified time interval.
Disk utilization		
Disk Utilization Percentage	ESS, DS6000, DS8000: RAID array	The approximate utilization percentage of a rank over a specified time interval (the average percent of time that the disks associated with the array were busy). **Note:** Tivoli Storage Productivity Center does not calculate a value for this column if there are multiple ranks in the extent pool where thin-provisioned volumes are allocated. This column displays value of **N/A** for the reports in which it is displayed. However, if there is only a single rank in the extent pool, Tivoli Storage Productivity Center does calculate the value for this column regardless of the thin-provisioned volumes.
Sequential I/O Percentage	ESS, DS6000, DS8000: RAID array	Percentage of all I/O operations that were completed for a RAID array over a specified time interval that were sequential operations.
Important: [1] The "Resource: components" column shows the resources and components that are measured by a metric. Metrics that apply to Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified.		

Port metrics

The following table describes the metrics for the ports on storage systems and switches.

Column	Resources: components[1]	Description
I/O or frame rates		
Port Send I/O Rate	ESS, DS6000, DS8000: port, storage system SAN Volume Controller[1]: port, node, I/O group, storage system SMI-S BSP: port XIV[3]: port	Average number of I/O operations per second for send operations for a port over a specified time interval.

Column	Resources: components[1]	Description
Port Receive I/O Rate	ESS, DS6000, DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system SMI-S BSP: port XIV[3]: port	Average number of I/O operations per second for receive operations for a port over a specified time interval.
Total Port I/O Rate	ESS, DS6000, DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system SMI-S BSP: port XIV[3]: port	Average number of I/O operations per second for send and receive operations for a port over a specified time interval.
Port Send Frame Rate	switch port, switch	Average number of frames per second that are sent through the port.
Port Receive Frame Rate	switch port, switch	Average number of frames per second that are received by the port.
Total Port Frame Rate	switch port, switch	Average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.
Port to Host Send I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second sent to host computers by a component over a specified time interval.
Port to Host Receive I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second received from host computers by a component over a specified time interval.
Total Port to Host I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second transmitted between host computers and a component over a specified time interval.
Port to Disk Send I/O Rate A	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second sent to storage systems by a component over a specified time interval.
Port to Disk Receive I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second received from storage systems by a component over a specified time interval.
Total Port to Disk I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second transmitted between storage systems and a component over a specified time interval.

Column	Resources: components[1]	Description
Port to Local Node Send I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second sent to other nodes in the local SAN Volume Controller cluster by a component over a specified time interval.
Port to Local Node Receive I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second received from other nodes in the local SAN Volume Controller cluster by a component over a specified time interval.
Total Port to Local Node I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second transmitted between other nodes in the local SAN Volume Controller cluster and a component over a specified time interval.
Port to Remote Node Send I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second sent to nodes in the remote SAN Volume Controller cluster by a component over a specified time interval.
Port to Remote Node Receive I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second received from nodes in the remote SAN Volume Controller cluster.
Total Port to Remote Node I/O Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of exchanges (I/Os) per second transmitted between nodes in the remote SAN Volume Controller cluster and a component over a specified time interval.
Port FCP Send I/O Rate	ESS, DS6000, DS8000: port	Average number of send operations per second using the FCP protocol, for a port over a specified time interval.
Port FCP Receive I/O Rate	ESS, DS6000, DS8000: port	Average number of receive operations per second using the FCP protocol for a port over a specified time interval.
Total Port FCP I/O Rate	ESS, DS6000, DS8000: port	Average number of send and receive operations per second using the FCP protocol for a port over a specified time interval.
Port FICON Send I/O Rate	ESS, DS6000, DS8000: port	Average number of send operations per second using the FICON protocol for a port over a specified time interval.
Port FICON Receive I/O Rate	ESS, DS6000, DS8000: port	Average number of receive operations per second using the FICON protocol for a port over a specified time interval.
Total Port FICON I/O Rate	ESS, DS6000, DS8000: port	Average number of send and receive operations per second using the FICON protocol for a port over a specified time interval.
Port PPRC Send I/O Rate	ESS, DS6000, DS8000:: port, storage system	Average number of send operations per second for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Port PPRC Receive I/O Rate	ESS, DS6000, DS8000: port, storage system	Average number of receive operations per second for Peer-to-Peer Remote Copy usage for a port over a specified time interval.

Column	Resources: components[1]	Description
Total Port PPRC I/O Rate	ESS, DS6000, DS8000: port, storage system	Average number of send and receive operations per second for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Data rates		
Port Send Data Rate	ESS, DS6000, DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system SMI-S BSP: port switch port, switch XIV[3]: port	Average number of megabytes (2^20 bytes) per second that were transferred for send (read) operations for a port over a specified time interval.
Port Receive Data Rate	ESS, DS6000, DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system SMI-S BSP: port switch port, switch XIV[3]: port	Average number of megabytes (2^20 bytes) per second that were transferred for receive (write) operations for a port over a specified time interval.
Total Port Data Rate	ESS, DS6000, DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system SMI-S BSP: port switch port, switch XIV[3]: port	Average number of megabytes (2^20 bytes) per second that were transferred for send and receive operations for a port over a specified time interval.
Port Peak Send Data Rate	switch port	Peak number of megabytes (2^20 bytes) per second that were sent by a port over a specified time interval
Port Peak Receive Data Rate	switch port	Peak number of megabytes (2^20 bytes) per second that were received by a port over a specified time interval.
Port to Host Send Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2^20 bytes) per second sent to host computers by a component over a specified time interval.

Column	Resources: components[1]	Description
Port to Host Receive Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2^20 bytes) per second received from host computers by a component over a specified time interval.
Total Port to Host Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2^20 bytes) per second transmitted between host computers and a component over a specified time interval.
Port to Disk Send Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2^20 bytes) per second sent to storage systems by a component over a specified time interval.
Port to Disk Receive Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2^20 bytes) per second received from storage systems by a component over a specified time interval.
Total Port to Disk Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2^20 bytes) per second transmitted between storage systems and a component over a specified time interval.
Port to Local Node Send Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2^20 bytes) per second sent to other nodes in the local SAN Volume Controller cluster by a component over a specified time interval.
Port to Local nNode Receive Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2^20 bytes) per second received from other nodes in the local SAN Volume Controller cluster by a component over a specified time interval.
Total Port to Local Node Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2^20 bytes) per second transmitted between other nodes in the local SAN Volume Controller cluster and a component over a specified time interval.
Port to Remote Node Send Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2^20 bytes) per second sent to nodes in the remote SAN Volume Controller cluster by a component over a specified time interval.
Port to Remote Node Receive Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2^20 bytes) per second received from nodes in the remote SAN Volume Controller cluster by a component over a specified time interval.
Total Port to Remote Node Data Rate	SAN Volume Controller: port, node, I/O group, storage system	Average number of megabytes (2^20 bytes) per second transmitted between nodes in the remote SAN Volume Controller cluster and a component over a specified time interval.
Port FCP Send Data Rate	ESS, DS6000, DS8000: port	Average number of megabytes (2^20 bytes) per second sent over the FCP protocol for a port over a specified time interval.
Port FCP Receive Data Rate	ESS, DS6000, DS8000: port	Average number of megabytes (2^20 bytes) per second received over the FCP protocol for a port over a specified time interval.

Column	Resources: components[1]	Description
Total Port FCP Data Rate	ESS, DS6000, DS8000: port	Average number of megabytes (2^20 bytes) per second sent or received over the FCP protocol for a port over a specified time interval.
Port FICON Send Data Rate	ESS, DS6000, DS8000: port	Average number of megabytes (2^20 bytes) per second sent over the FICON protocol for a port over a specified time interval.
Port FICON Receive Data Rate	ESS, DS6000, DS8000: port	Average number of megabytes (2^20 bytes) per second received over the FICON protocol, for a port over a specified time interval.
Total Port FICON Data Rate	ESS, DS6000, DS8000: port	Average number of megabytes (2^20 bytes) per second sent or received over the FICON protocol for a port over a specified time interval.
Port PPRC Send Data Rate	ESS, DS6000, DS8000: port, storage system	Average number of megabytes (2^20 bytes) per second sent for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Port PPRC Receive Data Rate	ESS, DS6000, DS8000: port, storage system	Average number of megabytes (2^20 bytes) per second received for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Total Port PPRC Data Rate	ESS, DS6000, DS8000: port, storage system	Average number of megabytes (2^20 bytes) per second transferred for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Response times		
Port Send Response Time	ESS, DS6000, DS8000: port, storage system XIV[3]: port	Average number of milliseconds that it took to service each send (read) operation for a port over a specified time interval.
Port Receive Response Time	ESS, DS6000, DS8000: port, storage system XIV[3]: port	Average number of milliseconds that it took to service each receive (write) operation for a port over a specified time interval.
Overall Port Response Time	ESS, DS6000, DS8000: port, storage system XIV[3]: port	Average number of milliseconds that it took to service each operation (send and receive) for a port over a specified time interval.
Port to Local Node Send Response Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds it took to service each send operation to another node in the local SAN Volume Controller cluster for a component over a specified time interval. This value is the external response time of the transfers.
Port to Local Node Receive Response Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds it took to service each receive operation from another node in the local SAN Volume Controller cluster for a component over a specified time interval. This value is the external response time of the transfers.
Total Port to Local Node Response Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds it took to service each send or receive operation between another node in the local SAN Volume Controller cluster and a component over a specified time interval. This value is the external response time of the transfers.

Column	Resources: components[1]	Description
Port to Local Node Send Queued Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds that each send operation issued to another node in the local SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.
Port to Local Node Receive Queued Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds that each receive operation from another node in the local SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.
Total Port to Local Node Queued Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds that each operation issued to another node in the local SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.
Port to Remote Node Send Response Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds it took to service each send operation to a node in the remote SAN Volume Controller cluster for a component over a specified time interval. This is the external response time of the transfers.
Port to Remote Node Receive Response Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds it took to service each receive operation from a node in the remote SAN Volume Controller cluster for a component over a specified time interval. This value is the external response time of the transfers.
Total Port to Remote Node Response Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds it took to service each send or receive operation between a node in the remote SAN Volume Controller cluster and a component over a specified time interval. This value is the external response time of the transfers.
Port to Remote Node Send Queued Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds that each send operation issued to a node in the remote SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.
Port to Remote Node Receive Queued Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds that each receive operation from a node in the remote SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.
Total Port to Remote Node Queued Time	SAN Volume Controller: node, I/O group, storage system	Average number of milliseconds that each operation issued to a node in the remote SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.
Port FCP Send Response Time	ESS, DS6000, DS8000: port	Average number of milliseconds it took to service all send operations over the FCP protocol for a port over a specified time interval.

Column	Resources: components[1]	Description
Port FCP Receive Response Time	ESS, DS6000, DS8000: port	Average number of milliseconds it took to service all receive operations over the FCP protocol for a port over a specified time interval.
Overall Port FCP Response Time	ESS, DS6000, DS8000: port	Average number of milliseconds it took to service all I/O operations over the FCP protocol for a port over a specified time interval.
Port FICON Send Response Time	ESS, DS6000, DS8000: port	Average number of milliseconds it took to service all send operations over the FICON protocol for a port over a specified time interval.
Port FICON Receive Response Time	ESS, DS6000, DS8000: port	Average number of milliseconds it took to service all receive operations over the FICON protocol for a port over a specified time interval.
Overall Port FICON Response Time	ESS, DS6000, DS8000: port	Average number of milliseconds it took to service all I/O operations over the FICON protocol for a port over a specified time interval.
Port PPRC Send Response Time	ESS, DS6000, DS8000: port, storage system	Average number of milliseconds it took to service all send operations for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Port PPRC Receive Response Time	ESS, DS6000, DS8000: port, storage system	Average number of milliseconds it took to service all receive operations for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Overall Port PPRC Response Time	ESS, DS6000, DS8000: port, storage system	Average number of milliseconds it took to service all I/O operations for Peer-to-Peer Remote Copy usage for a port over a specified time interval.
Transfer sizes		
Port Send Transfer Size	ESS, DS6000, DS8000: port, storage system SMI-S BSP: port	Average number of KB sent per I/O by a port over a specified time interval.
Port Receive Transfer Size	ESS, DS6000, DS8000: port, storage system SMI-S BSP: port	Average number of KB received per I/O by a port over a specified time interval.
Overall Port Transfer Size	ESS, DS6000, DS8000: port, storage system SMI-S BSP: port	Average number of KB transferred per I/O by a port over a specified time interval.
Port Send Frame Size	switch port, switch	The average size of a frame that is sent out through a port.
Port Receive Frame Size	switch port, switch	The average size of a frame that is received in through a port.
Overall Port Frame Size	switch port, switch	The average frame transfer size. This value includes frames that are sent and frames that are received by a port.

Column	Resources: components[1]	Description
Special computed values		
Port Send Utilization Percentage	ESS, DS6000, DS8000: port	Average amount of time that the port was busy sending data over a specified time interval.
Port Receive Utilization Percentage	ESS, DS6000, DS8000: port	Average amount of time that the port was busy receiving data over a specified time interval.
Overall Port Utilization Percentage	ESS, DS6000, DS8000: port	Average amount of time that the port was busy sending or receiving data over a specified time interval.
Port Send Bandwidth Percentage	ESS, DS8000: port SAN Volume Controller: port switch, port XIV[3]: port	The approximate bandwidth utilization percentage for send operations by a port, based on its current negotiated speed.
Port Receive Bandwidth Percentage	ESS, DS8000: port SAN Volume Controller: port switch, port XIV[3]: port	The approximate bandwidth utilization percentage for receive operations by this port, based on its current negotiated speed.
Overall Port Bandwidth Percentage	ESS, DS8000: port SAN Volume Controller: port switch, port XIV[3]: port	The approximate bandwidth utilization percentage for send and receive operations by this port.
Error rates		
Error Frame Rate	switch port, switch DS8000: port, storage system	The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.
Discarded Frame Rate	switch port, switch	The average number of frames per second that are discarded because host buffers are unavailable for the port.
Link Failure Rate	switch port, switch DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system	The number of link errors per second.

Column	Resources: components[1]	Description
Loss of Sync Rate	switch port, switch DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system	The average number of times per second that synchronization is lost. Synchronization is assumed lost after a timeout interval expires.
Loss of Signal Rate	switch port, switch DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system	The average number of times per second that the signal is lost.
CRC Error Rate	switch port, switch DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system	The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver.
Short Frame Rate	switch port, switch	The average number of frames that are received per second that are shorter than 28 octets. This number excludes start-of-frame bytes and end-of-frame bytes. The 28 octet limit is calculated based on the assumption that a frame has 24 header bytes, and 4 CRC bytes.
Long Frame Rate	switch port, switch	The average number of frames that are received per second that are longer than 2140 octets. This number excludes start-of-frame bytes and end-of-frame bytes. The 2140 octet limit is calculated based on the assumption that a frame has 24 header bytes, 4 CRC bytes, and 2112 data bytes.
Encoding Disparity Error Rate	switch port, switch	The average number of disparity errors per second that are received.
Discarded Class3 Frame Rate	switch port, switch	The average number of class 3 frames per second that are discarded.
F-BSY Frame Rate	switch port, switch	The average number of F-BSY frames per second that are generated. An F-BSY frame is issued by the fabric to indicate that a frame cannot be delivered because the fabric or destination N_port is busy.
F-RJT Frame Rate	switch port, switch	The average number of F-RJT frames per second that are generated. An F-RJT frame is issued by the fabric to indicate that delivery of a frame was denied.

Column	Resources: components[1]	Description
Primitive Sequence Protocol Error Rate	switch port, switch DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system	The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.
Invalid Transmission Word Rate	switch port, switch DS8000: port, storage system SAN Volume Controller: port, node, I/O group, storage system	The average number of bit errors per second that are detected. Errors include an 8b10 code violation in one or more characters, a K28.5 in the second, third, or fourth character positions, or an ordered set that had an incorrect Beginning Running Disparity.
Zero Buffer-Buffer Credit Timer	SAN Volume Controller: port, node, I/O group, storage system	The number of microseconds that the port was unable to send frames because of a lack of buffer credit since the last node reset.
Zero Buffer-Buffer Credit Percentage	SAN Volume Controller: port, node, I/O group, storage system	The percentage of sampling time that the port was unable to send frames because of a lack of buffer credit since the last node reset.
Link Reset Transmitted Rate	switch port, switch DS8000: port, storage system	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.
Link Reset Received Rate	switch port, switch DS8000: port, storage system	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.
Out of Order Data Rate	DS8000: port, storage system	The average number of times per second that an out of order frame was detected for a port over a specified time interval.
Out of Order ACK Rate	DS8000: port, storage system	The average number of times per second that an out of order ACK frame was detected for a port over a specified time interval.
Duplicate Frame Rate	DS8000: port, storage system	The average number of times per second that a frame was received and also was detected as previously processed for a port over a specified time interval.
Invalid Relative Offset Rate	DS8000: port, storage system	The average number of times per second that a frame was received with invalid relative offset in the frame header for a port over a specific time interval.
Sequence Timeout Rate	DS8000: port, storage system	The average number of times per second the port detected a timeout condition on receiving sequence initiative for a Fibre Channel exchange for a port over a specified time interval.

Column	Resources: components[1]	Description
Important: • [1] The "Resource: components" column shows the resources and components that are measured by a metric. Metrics that apply to Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified. • [2] This metric is available in the XIV system version 10.2.2 or later. • [3] This metric is available in the XIV system version 10.2.4 or later.		

Related concepts:

"Monitoring the performance of resources" on page 383
Tivoli Storage Productivity Center can collect information about the performance of storage systems and switches. This information includes key performance metrics and notifications of threshold violations. Metrics can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage. Threshold violations alert you when the performance of a monitored resource falls outside of a specified range.

Arguments and window locations available with the inbound launch in context feature

This section describes the arguments and window locations available with the inbound launch in context feature.

Launch in context parameters

You can use the launch in context parameters, definitions, and examples.

The command format for Java Web Start is unique in that a parameter and its value are separated by an equal sign (=). Parameter=value pairs are separated by an ampersand (&).

Not all parameters can be used with all launch types. Where this is true, it is noted in the parameter definition.

Parameters

-user The user ID. The value must be a valid user ID to log on to the IBM Tivoli Storage Productivity Center server, for example, `-user admin`.

Note: This parameter is not supported in a Web Browser's URL when using Java Web Start to launch the Tivoli Storage Productivity Center user interface.

-passwd
The password. The value must be a valid password to log on to the Tivoli Storage Productivity Center server, for example, `-passwd MyPassword`.

Note: This parameter is not supported in a Web Browser's URL when using Java Web Start to launch the Tivoli Storage Productivity Center user interface.

-encpasswd
The password encrypted. This is used if you choose to encrypt the password instead of entering the text password, for example, `-encpasswd 48735`.

Tip: To encrypt a password using Tivoli Storage Productivity Center (the preferred method), enter the following command in the Command Prompt window:

TPC-HOME/data/server/tools/tpctool encrypt *text password*

The encrypted password is returned.

Note: This parameter is not supported in a Web Browser's URL when using Java Web Start to launch the Tivoli Storage Productivity Center user interface.

-encmethod

The encryption method used to encrypt the password. The password can be encrypted using Tivoli Storage Productivity Center or Tivoli Enterprise Portal. The preferred method is to use Tivoli Storage Productivity Center to encrypt the password. If **encpasswd** is selected, you must also select **encmethod** with a value. The default is TPC. For example, `-encpasswd 48735 -encmethod TPC`

Note: This parameter is not supported in a Web Browser's URL when using Java Web Start to launch the Tivoli Storage Productivity Center user interface.

Tip: The Tivoli Enterprise Portal graphical user interface (GUI) does not start the Tivoli Storage Productivity Center GUI properly with a Tivoli Enterprise Portal-encrypted password if that encryption creates a string with spaces in it. Such a string is handled as multiple values unless it is enclosed in double quotation marks (").

Tip: Encryption using Tivoli Enterprise Portal creates a string pattern. There are multiple string patterns that can result in the same decrypted value, so it is possible that different string patterns used as an encrypted password might result in same decrypted password.

-function

Specify one of the following values:

- The task to perform on the remote data server. See "Launch-in-context task parameters" on page 1731 for tasks and their specific parameters. For example, `-function disk.create_volume -subsystem_id <ID of the subsystem>`.
- The window name to open on the target system. The window name is mapped to the name of the node on the navigation tree. See "Available windows" on page 1731 for a list of available windows. For example, `-function topology.computer`.

-server

The Tivoli Storage Productivity Center server to log on to. This argument specifies a single server. For example, `-server server1:9549`.

-servers

Not used with the Web Browser. The Tivoli Storage Productivity Center server or servers to log on to, in the format server:port. You can specify multiple server ports separated by commas. For example, `-servers "server1:9549,server2:9549,server3:9549"`. If you specify multiple server ports, Tivoli Storage Productivity Center will connect to the first server. If you later initiate a new connection from the main window (**Connection** > **New Connection**), the server list box will include the available servers in the **Connection** dialog.

-help **Not used with the Web Browser.** The help argument is used alone in the Arguments text field. It is entered as `-help functions`. This argument returns a list of the windows representing nodes in the navigation tree that can be launched. Typing `-help` with no arguments will return an error message.

You can also enter **help** as a command in the Command Prompt window as `tpc.bat -help functions` for Windows, and `TPCD.sh -help functions` for AIX. Linux, or UNIX.

Launch-in-context task parameters

Use the launch-in-context feature task parameters to complete storage tasks.

Task parameters that are associated with actions taken on the Data server also use the **-function** parameter. For example: `-function data.filesystem_capacity_by_computer -computer_id "<computer_id>"`.

Capacity report (data.filesystem_capacity_by_computer)

-computer_id
The ID of the computer for which the report is to be created. If this value is not specified, the Tivoli Storage Productivity Center GUI opens the default report panel for **Data Manager** > **Reporting** > **Capacity** > **FileSystem Capacity** > **By Computer**.

You must enter a value for **-computer_id** that matches the ID for that computer that is stored in the Tivoli Storage Productivity Center database repository. You receive a warning message if the computer ID you enter in the command line does not match the ID stored in the repository. To ensure that this report is displayed properly when you use launch-in-context, make sure to enter an ID for the computer that matches the ID that is displayed for it through the Topology Viewer or appropriate report in the product's user interface.

Wasted space report (TPC.reports_data.wasted_space)

This report has no parameters. The Tivoli Storage Productivity Center GUI opens the default report panel for **Data Manager** > **Reporting** > **System Reports** > **Data** > **Wasted Space**.

Available windows

You can use windows with the inbound launch-in-context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the **-function** window name argument.

Available window names for function argument

Table 151. Available windows for inbound launch-in-context

Tree Node	Window Name
Administrative Services -> Data Sources ->IBM Tivoli Storage Productivity Center Servers[1]	datasource.tpc_servers
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->All	alert.all
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Alerts Directed to user	alert.alerts_directed_to_user

Table 151. Available windows for inbound launch-in-context (continued)

Tree Node	Window Name
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Storage Subsystem	alert.storage_subsystem
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Computer	alert.computer
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Disk	alert.disk
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->File system	alert.filesystem
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Directory[1]	alert.directory
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->User[1]	alert.user
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->OS User Group[1]	alert.os_user_group
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Fabric	alert.fabric
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Switch	alert.switch
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Endpoint Device	alert.endpoint_device
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->External	alert.external
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Configuration Analysis[1]	alert.configuration_analysis
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Hypervisor[1]	alert.hypervisor
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Disk Space Summary[1]	data_report.disk_capacity
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Storage Access Times[1]	data_report.storage_access_times
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Most Obsolete Files[1]	data_report.most_obsolete_files
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Storage Availability	data_report.storage_availibility
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Storage Capacity[1]	data_report.storage_capacity
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Total Freespace[1]	data_report.total_freespace
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data-> User Space Usage[1]	data_report.user_space_usage
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Disk->Port Performance[1]	disk_report.port_performance
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Disk->Subsystem Performance[1]	disk_report.subsystem_performance
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Disk->Top Volumes Data Rate Performance[1]	disk_report.top_volumes_data_rate_performance

Table 151. Available windows for inbound launch-in-context (continued)

Tree Node	Window Name
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Disk-> Top Volumes I/O Rate Performance[1]	disk_report.top_volumes_io_rate_performance
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Fabric-> SAN Assets (ALL)	fabric_report.san_assets
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Fabric ->Port Connections	fabric_report.port_connections
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Fabric -> Switch Performance[1]	fabric_report.switch_performance
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Fabric ->Top Switch Ports Data Rate Performance[1]	fabric_report.top_switch_ports_data_rate_performance
Disk Manager->Storage Subsystems	disk.storage_subsystems
Fabric Manager->Fabrics	fabric.fabrics
Windows that correspond to launch-in-context tasks	
Data Manager-> Reporting -> Capacity -> File system Capacity -> By Computer[1]	data.filesystem_capacity_by_computer
IBM Tivoli Storage Productivity Center -> My Reports -> System Reports -> Data -> Wasted Space[1]	TPC.reports.data.wasted_space
[1]This window is not available in IBM Tivoli Storage Productivity Center Basic Edition.	

Opening Tivoli Storage Productivity Center on Windows operating systems

You can open Tivoli Storage Productivity Center CLIs and GUIs and administer Tivoli Storage Productivity Center on Windows operating systems.

Opening Tivoli Storage Productivity Center CLIs and GUIs

To manage and monitor storage resources, open Tivoli Storage Productivity Center CLIs and GUIs.

You can open the following GUIs and CLIs:

- "Opening Tivoli Storage Productivity Center Web-based GUI" on page 1734
- "Opening Tivoli Storage Productivity Center Stand-alone GUI" on page 1734
- "Opening Tivoli Storage Productivity Center Replication Manager GUI" on page 1734
- "Opening DB2 Command Window" on page 1735
- "Opening IBM Data Studio Administration client on Windows operating systems" on page 1735
- "Opening IBM Tivoli Monitoring Services" on page 1735

Opening Tivoli Storage Productivity Center Web-based GUI

1. Choose one of these options:

Option	Description
Windows Server 2012	1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**. 2. On the Start page, right-click, and then click the **All apps** taskbar button.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click **Start** > **All Programs**.

2. Click **Tivoli Storage Productivity Center** > **TPC Web-based GUI**.

Opening Tivoli Storage Productivity Center Stand-alone GUI

1. Choose one of these options:

Option	Description
Windows Server 2012	1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**. 2. On the Start page, right-click, and then click the **All apps** taskbar button.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click **Start** > **All Programs**.

2. Click **Tivoli Storage Productivity Center** > **TPC Stand-alone GUI**.

Opening Tivoli Storage Productivity Center Replication Manager GUI

1. Choose one of these options:

Option	Description
Windows Server 2012	1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**. 2. On the Start page, right-click, and then click the **All apps** taskbar button.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click **Start** > **All Programs**.

2. Click **Tivoli Storage Productivity Center** > **TPC Replication Manager GUI**.

Opening DB2 Command Window

1. Choose one of these options:

Option	Description
Windows Server 2012	1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**. 2. On the Start page, right-click, and then click the **All apps** taskbar button.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click **Start** > **All Programs**.

2. Click **IBM DB2** > **Command Line Tools** > **Command Window**.

Opening IBM Data Studio Administration client on Windows operating systems

1. Choose one of these options:

Option	Description
Windows Server 2012	1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**. 2. On the Start page, right-click, and then click the **All apps** taskbar button.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click **Start** > **All Programs**.

2. Click **IBM Data Studio** > **Data Studio Administration Client**.

Opening IBM Tivoli Monitoring Services

1. Choose one of these options:

Option	Description
Windows Server 2012	1. On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**. 2. On the Start page, right-click, and then click the **All apps** taskbar button.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click **Start** > **All Programs**.

2. Click **IBM Tivoli Monitoring** > **IBM Tivoli Monitoring Services**.

Accessing administration tools

To manage and maintain Tivoli Storage Productivity Center, access the Windows system administration tools.

To complete tasks in Tivoli Storage Productivity Center, you must open the following administration and maintenance facilities:

- "Accessing the Control Panel" on page 1736

- "Accessing Administrative Tools"
- "Accessing Windows Services"
- "Accessing Computer Management"
- "Accessing Programs and Program Features" on page 1737

Accessing the Control Panel

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click **Start**.

2. Click **Control Panel**

Accessing Administrative Tools

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**.
Windows 7, Windows 2008	Click **Start**.
Windows XP, Windows Vista	Click **Start** > **Control Panel**.

2. Click **Administrative Tools**

Accessing Windows Services

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**.
Windows 7, Windows 2008	Click **Start**.
Windows XP, Windows Vista	Click **Start** > **Control Panel**.

2. Click **Administrative Tools** > **Services**

Accessing Computer Management

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**.

Option	Description
Windows 7, Windows 2008, Windows Vista, Windows XP	Click **Start**.

2. Click **Control Panel** > **Computer Management**

Accessing Programs and Program Features

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click **Start**.

2. Click **Control Panel** > **Programs and Program Features**.

Accessing Window Run

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the **Server Manager** taskbar button, and then click **Start**.
Windows 7, Windows 2008, Windows Vista, Windows XP	Click **Start**.

2. Click **Run**.

Frequently Asked Questions

Use this section to view answers to common questions about IBM Tivoli Storage Productivity Center. These questions are organized according to the manager to which they apply.

General and Migration information

Learn about the issues when you work with IBM Tivoli Storage Productivity Center, upgrade to a new version of the product, or migrate from another product.

What Java version is officially supported by Tivoli Storage Productivity Center version 5.2 graphical user interface?

Java 7 is the level that is supported for the Tivoli Storage Productivity Center GUI. For information about how to start the Tivoli Storage Productivity Center graphical user interface, see the Tivoli Storage Productivity Center Information Center. Search for **Starting and stopping the console**.

You are entering namespaces manually for CIM agents and switches in the Tivoli Storage Productivity Center user interface. What are the current namespaces for the supported brands?

Table 152. Interop Namespaces for CIM Agents for switches and storage systems

Switch or system	Namespace
IBM (other than IBM System Storage TS3310 Tape Library)	/root/ibm
IBM System Storage TS3310 Tape Library	/root/cimv2
Brocade	/interop
Cisco	/root/cimv2 For version 3.2.1 or later: /root/pg_interop
Engenio	/interop
EMC	/root/emc
HDS	For HiCommand 5.6 or later, use: /root/smis/current For a HiCommand version that supports SMI-S 1.2, use following namespace to traverse the model as SMI-S 1.1: /root/smis/smis11 For HiCommand versions earlier than HiCommand 5.6, use: /root/hitachi/dm*xx*, where *xx* is represents the level of HiCommand.
HP	/root
LSI	/root/PG_InterOp
SUN	/root/sun3510 or /interop Note: This namespace is for a system and not a switch.
XYRATEX	/root/PG_interop

Check with your switch vendor to ensure you use the most current namespaces.

What if you specify an incorrect namespace?

If you specify an incorrect namespace, the following issues might occur:

- The connection test fails when the CIMOM is added.
- The discovery does not discover all information of the system that is managed by the CIMOM.
- The probe fails.
- The function that you want to start on the system might fail (for example, collecting performance data).

What are the required steps when you upgrade Tivoli Storage Productivity Center to the latest version?

For detailed information about how to upgrade the product, go to the Tivoli Storage Productivity Center Information Center and expand the **Tivoli Storage Productivity Center** > **Upgrading** node in the navigation tree.

You are using an older version of IBM Tivoli Storage Productivity Center for Data. How can I upgrade to this latest version?

For detailed information about how to upgrade the product, go to the Tivoli Storage Productivity Center Information Center and expand the **Tivoli Storage Productivity Center** > **Upgrading** node in the navigation tree.

Data Manager

Use the questions in this section to answer questions about working with the Data Manager component of IBM Tivoli Storage Productivity Center.

How do you know if your storage system is supported by Data Manager and which SMI-S agents are supported?

To confirm whether your storage system is supported and which SMI-S agent is supported for that system, review the supported products list for the current release of Tivoli Storage Productivity Center. Go to http://www-01.ibm.com/support/docview.wss?rs=40&context=SWJ50&q1=matrix&uid=swg21386446&loc=en_US&cs=utf-8&lang=en and click the applicable Tivoli Storage Productivity Center release in the Storage column.

Fabric Manager

Learn about the frequently asked questions for monitoring fabrics and switches in Tivoli Storage Productivity Center.

What type of fabric agents can be used in Tivoli Storage Productivity Center?

The following table shows the current agent types that are supported for the monitoring fabrics and switches.

Table 153. Agent types for switch and fabric functions

Function > Switch	Brocade	Cisco [1]	QLogic
Switch Performance Monitoring	CIM agent	CIM agent	CIM agent
Switch and switch port information collected	Recommended: CIM agent Also supported: SNMP agent or Storage Resource agent	Recommended: CIM agent Also supported: SNMP agent or Storage Resource agent	SNMP agent or Storage Resource agent
Topology connectivity information collected	Recommended: CIM agent Also supported: SNMP agent or Storage Resource agent	SNMP agent or Storage Resource agent	SNMP agent or Storage Resource agent
Zoning information collected [2]	Recommended: CIM agent	A Storage Resource agent is required in each VSAN [3] [4] [5]	Storage Resource agent
Tivoli Storage Productivity Center Alerts	Recommended: CIM agent Also supported: SNMP agent or Storage Resource agent	SMP agent or Storage Resource agent	SNMP agent or Storage Resource agent
Hosts, endpoint devices, device-centric and host-centric information collected	Storage Resource agent	Storage Resource agent	Storage Resource agent
Switch “Sensors and Events”	SNMP agent	SNMP agent	SNMP agent

Table 153. Agent types for switch and fabric functions (continued)

Function > Switch	Brocade	Cisco [1]	QLogic
Tips: [1] The Storage Resource agent collects the information about a VSAN basis. Each VSAN is viewed as an individual SAN. The SNMP agent and CIM agent get the physical fabric information and can correlate the VSAN information to a physical infrastructure. The SNMP agent also collects some VSAN information. [2] "Zone information collected" indicates that information about the active zone set and the inactive zoning library is collected during a fabric probe. [3] For many Cisco proprietary zone members, the values are displayed as zeros. [4] If a zoning configuration has Cisco proprietary zone members, then you cannot modify the zoning for that configuration through Tivoli Storage Productivity Center. You must use the element manager to remove proprietary zone members or do zoning. [5] Orphan Cisco zones are not reported. [6] Tivoli Storage Productivity Center does not get a transaction lock on the switch while zone changes are made on the fabric.			

What are the prerequisites for monitoring the performance of switches?
To monitor the performance of a switch, add the switch to Tivoli Storage Productivity Center. When you add the switch for monitoring, ensure that you schedule a probe and performance monitor to collect information about that switch.

Why does an Out of Band Fabric (SNMP) agent show a state of "not contacted"?
Complete the following tasks in the stand-alone GUI to help determine why an Out of Band Fabric (SNMP) agent shows a state of "not contacted":

- Check that the IP address or name is entered correctly in the Out of Band Fabric agent details panel.
- Ensure that the FA MIB is enabled on the switch.
- Ensure that the Tivoli Storage Productivity Center server IP address is included in any SNMP Access Control List that exists on the switches.
- Ensure that you are using SNMP v1.

Can a storage system be added as an Out of Band Fabric (SNMP) agent?
No. Out of Band Fabric agents within Tivoli Storage Productivity Center are meant for switches only. Disk Manager uses CIMOMs to gather subsystem information.

How do you manually change a fabric or switch label that is displayed in Tivoli Storage Productivity Center?
In the web-based GUI, go to **Network Resources** > **Fabrics** or **Network Resources** > **Switches**. Right-click a fabric or switch and select **View Properties**. On the properties notebook for a fabric or switch, click **Edit** and change the name in the **Name** field. Click **Save**.

Disk Manager

Use the questions in this section to answer questions about working with the Disk Manager component of IBM Tivoli Storage Productivity Center.

Keep in mind the following information and considerations:

You encounter errors while collecting performance data on SAN Volume Controller. You fail to associate SAN Volume Controller performance data from non-configuration node with SAN Volume Controller performance data from configuration node. You encounter incomplete SAN Volume Controller

performance data sample.

This issue is caused by a configuration issue with SAN Volume Controller (time zone). Reset the time zone on SAN Volume Controller by logging into the SAN Volume Controller through putty. Run this command first:

```
svctask settimezone -timezone 509
```

This forces the cluster into the Universal time zone. To get the time zone you want the cluster to be in, run this command:

```
svctask settimezone -timezone
```

Click on Launch Element Manager. The URL that appears is http://0.0.0.0/ica

This error is caused by the wrong configuration of the SAN Volume Controller ICAT. Use the SAN Volume Controller GUI to remove and add the cluster again to the ICAT. This populates the correct value for the access point which is stored in the database repository during the next discovery. Use this cli command to set the IP address and port of the user interface.

For example: `svcluster -x modify -p 9.155.62.91:9080`

This is the same command used by the user interface and CIMOM.

Protocols and standards

This section provides an overview of the protocols and standards that are used within IBM Tivoli Storage Productivity Center.

Web Based Enterprise Management

Web Based Enterprise Management (WBEM) is an initiative of the Distributed Management Task Force (DTMF) with the objective to enable the management of complex IT environments. It defines a set of management and internet standard technologies in order to unify the management of complex IT environments.

The WBEM initiative is composed of three main conceptual elements:

Common Interface Model (CIM)

CIM is a formal object-oriented modeling language that is used to describe the management aspects of systems.

xmlCIM

This is the grammar to describe CIM declarations and messages used by the CIM protocol.

Hypertext Transfer Protocol (HTTP)

HTTP is used as a way to enable communication between a management application and a device that both use CIM.

The WBEM architecture defines the following elements:

CIM Client

The CIM Client is a management application like IBM Tivoli Storage Productivity Center that uses CIM to manage devices. A CIM Client can reside anywhere in the network, because it uses HTTP to talk to CIM Object Managers and Agents.

CIM Managed Object

A Managed Object is a hardware or software component that can be managed by a management application by using CIM.

CIM Agent
A CIM Object Manager that includes the provider service for a limited set of resources. An agent may be embedded or hosted and can be an aggregator for multiple devices.

CIM Provider
A CIM Provider is the element that translates CIM calls to the device-specific commands. A provider is always closely linked to a CIM.

CIM Object Manager (CIMOM)
The central component of the CIM Server responsible for the communication between the CIM server components.

CIM Server
A server that receives and processes CIM Operation Message Requests and issues CIM Operation Message Responses.

Storage Management Initiative Specification

The Storage Networking Industry Association (SNIA) defines a standard that is used within IBM Tivoli Storage Productivity Center to create and develop a universal open interface for managing storage devices including storage networks.

Information on SMI-S can be found at this Web site:

`http://www.snia.org`

SNIA has fully adopted and enhanced the Common Information Model (CIM) standard for storage management in its Storage Management Initiative - Specification (SMI-S). SMI-S was launched to create and develop a universal open interface for managing storage devices including storage networks. SMI-S provides:

- A comprehensive specification for the management of heterogeneous storage and storage area networks (SANs).
- The information available to a WBEM client from an SMI-S compliant CIM server (provider).
- Profiles organized by:
  - Storage
  - Fabric
  - Host
  - Common profiles and subprofiles
- An object-oriented CIM and XML-based interface for managing SAN devices, services, and fabrics.
- An initial discovery, which is SLP based.

The idea behind SMI-S is to standardize the management interfaces so that management applications can utilize these and provide cross-device management. This means that a newly introduced device can be immediately managed as it will conform to the standards.

The models and protocols in the SMI-S implementation are platform-independent, enabling application development for any platform, and enabling them to run on different platforms. The SNIA will also provide interoperability tests which will help vendors test their applications and devices if they conform to the standard.

Service Location Protocol

The Service Location Protocol (SLP) is an Internet Engineering Task Force (IETF) standard. SLP provides a scalable framework for the discovery and selection of network services.

The Internet Engineering Task Force (IETF) is a large open international community of network designers, operators, vendors, and researchers that are concerned with the evolution of the Internet architecture and the smooth operation of the Internet. The IETF includes formal standards for SNMP and MIBs. More information on IETF can be found at this website:

`http://www.ietf.org`

SLP enables the discovery and selection of generic services, which can range in function from hardware services such as those for printers or fax machines, to software services such as those for file servers, email servers, web servers, databases, or any other possible services that are accessible through an IP network.

Traditionally, to use a particular service, a user, or client application provided the host name or network IP address for the service. With SLP, however, it is not necessary for the user or client application to know individual host names or IP addresses. Instead, the user or client application can search the network for the required service type and an optional set of qualifying attributes.

For example, a user can search for all available printers that support Adobe PostScript. Based on the service type such as printers and the attributes such as PostScript, SLP searches the user's network for matching services, and returns the discovered list to the user.

Simple Network Management Protocol

The Simple Network Management Protocol (SNMP) is an Internet Engineering Task Force (IETF) protocol for monitoring and managing systems and devices in a network. Functions supported by the SNMP protocol are the request and retrieval of data, the setting or writing of data, and traps that signal the occurrence of events.

SNMP is a method that enables a management application to query information from a managed device. The managed device has software running that sends and receives the SNMP information. This software module is usually called the SNMP agent.

An SNMP manager can read information from an SNMP agent in order to monitor a device. Therefore, the device needs to be polled on an interval bases. The SNMP manager can also change the configuration of a device, by setting certain values to corresponding variables. A device can also be set up to send a notification to the SNMP manager (this is called a trap) to asynchronously inform this SNMP manager of a status change.

Depending on the existing environment and organization it is very likely that your environment already has an SNMP management application in place. IBM Tivoli Storage Productivity Center can be set up to send traps.

SNMP uses a hierarchical structured Management Information Base (MIB) to define the meaning and the type of a particular value. A MIB defines managed objects that describe the behavior of the SNMP entity, which can be anything from an IP router to a storage subsystem. The information is organized in a tree structure.

For users planning to make use of the Tivoli Storage Productivity Center SNMP trap alert notification capabilities, an SNMP MIB is included in the server installation.

The MIB is provided for use by your SNMP management console software. Most SNMP management products provide a program called a MIB compiler that can be used to import MIBs. This will allow you to view Tivoli Storage Productivity Center generated SNMP traps from within your management console software. Refer to your management console software documentation for instructions on how to compile or import a third party MIB.

Tivoli Storage Productivity Center uses port 162 to listen for SNMP traps. This is the default port. For switches, you must configure the switch to send SNMP traps to the Device server IP address. If you need to change the default port number, use the **setdscfg** command. The attribute to set is **SNMPTrapPort**.

System administrators must set up their SNMP trap ringer with the provided MIB files in order to receive SNMP traps from Tivoli Storage Productivity Center. These files are located in the following directories on the product installation DVD:

For Data server:

```
data\snmp\tivoliSRM.mib
```

For Device server:

```
device\snmp\fabric.mib
```

Fibre Channel Methodologies of Interconnects

IBM Tivoli Storage Productivity Center supports the ANSI T11 Fibre Channel FC-MI (Fibre Channel Methodologies of Interconnects) for the automated discovery of FC SAN assets and topology.

ANSI T11 Fibre Channel FC-MI includes the following for the automated discovery of FC SAN assets and topology:

- Hosts (HBAs)
- FC interconnects
- FC storage devices

The T11 FC-MI also includes the following:

- FC-GS-3/4 (discovery, zoning, and so forth)
- RNID (advanced device recognition)
- Platform registration (device recognition and launch)
- Common HBA API (fabric and storage views)
- Name server (connectivity)
- Management server (SAN connectivity and topology)
- RSCN (advanced event detection)
- SCSI queries (storage views, volume information, and so forth)
- SNMP Fabric Element (FE) MIB
- SNMP FC Management MIB (discovery, performance statistics, and so forth)
- SNMP alerts

Appendix B. Accessibility features for Tivoli Storage Productivity Center

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

The following list includes the major accessibility features in IBM Tivoli Storage Productivity Center:

- Keyboard-only operation in the stand-alone GUI and the web-based GUI. Restriction: In the stand-alone GUI, you must use the mouse to navigate the Topology Viewer and report graphs.
- Interfaces that are commonly used by screen readers.
- An information center that includes the following accessibility features:
  - The information center is provided in XHTML 1.0 format, which is viewable in most web browsers. With XHTML, you can view documentation according to the display preferences that are set in your browser. XHTML supports screen readers and other assistive technologies.
  - All documentation for Tivoli Storage Productivity Center is available in Adobe Portable Document Format (PDF) by using the Adobe Acrobat Reader. You can access the PDFs from the Printable PDFs topic in the information center at http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp.
  - All images in the information center are provided with alternative text, so that visually impaired users can understand the contents of the images.

Keyboard navigation in the stand-alone GUI

The stand-alone GUI uses standard Microsoft Windows navigation keys. However, you must use the mouse to navigate the Topology Viewer and report graphs. Standard operating system keystrokes are used for standard operating system operations.

Keyboard navigation in the web-based GUI

Most of the features of the web-based GUI are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI), except as noted in the product release notes.

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. The following sections describe the keys or key combinations for different parts of the web-based GUI:

For navigating the web-based GUI and the context-sensitive help system:

- To navigate to the next link, button, or topic within a panel, press Tab.
- To move to the previous link, button, or topic within a panel, press Shift+Tab.
- To select an object, when the object is in focus, press Enter.

For actions menus:

- To navigate to the grid header, press Tab.

© Copyright IBM Corp. 2001, 2013

- To reach the drop-down field, press the Left Arrow or Right Arrow key.
- To open the drop-down menu, press Enter.
- To select the menu items, press the Up Arrow or Down Arrow key.
- To start the action, press Enter.

For filters:

To specify a filter option and text:

1. Press Tab to navigate to the magnifying glass icon.
2. Press the Up Arrow or Down Arrow key to navigate the filtering list.
3. Press Enter to select a filtering option.
4. When a filtering option is selected, the cursor moves to the filter text box. Type the filter text and press Enter. To reset a filter, press Enter.

For text fields:

- To navigate to text fields, press Tab.
- To navigate to the fields that are available for editing, press Tab.
- To navigate to the next field or to the **Submit** button, press Tab.

For tables or lists:

- To navigate between column headers, focus on a column header and use the Left Arrow and Right Arrow keys to move to other column headers.
- To navigate between data cells, focus on a data cell and use the Left, Right, Up, Down, Pageup, and Pagedown Arrow keys.
- To sort a column, focus on a column header and press Enter. The focus remains on the column header after the sort occurs.
- To change the size of a column, focus on the column header, hold Shift+Control, and press the Left or Right Arrow keys.
- To follow a link in a data cell, focus on a data cell and press Shift+F9.
- To open a menu for a table row, focus on the row and press Shift+F10.
- To select consecutive rows, select the first row and hold Shift, press the Up or Down Arrow keys to go to the last row in the range, and press the Space bar to add the new rows to the selection.
- To select non-consecutive rows, select a row and hold Control, press the Up or Down Arrow keys, and press the Space bar to add the new row to the selection.

Restriction: For Chinese languages, the keyboard combination Control+Space bar is not enabled for selecting multiple rows at the same time.

IBM and accessibility

For more information about IBM's commitment to accessibility, see the IBM Human Ability and Accessibility Center website at http://www.ibm.com/able/.

Appendix C. Accessibility features for Tivoli Storage Productivity Center for Replication

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

The following list includes the major accessibility features in Tivoli Storage Productivity Center for Replication:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

For more information about the commitment that IBM has for accessibility, see the IBM Human Ability and Accessibility Center website at www.ibm.com/able.

Accessibility and keyboard shortcuts in the information center

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. Using the major accessibility features in this product, users can perform these tasks:

- Use assistive technologies, such as screen-reader software and digital speech synthesizer, to hear what is displayed on the screen. Consult the product documentation of the assistive technology for details on using those technologies with this product.
- Operate specific or equivalent features by using only the keyboard.
- Magnify what is displayed on the screen.

In addition, the documentation was modified to include the following features to aid accessibility:

- All documentation is available in HTML formats to give the maximum opportunity for users to apply screen-reader software technology.
- All images in the documentation are provided with alternative text so that users with vision impairments can understand the contents of the images.

Use the following key combinations to navigate the interface by keyboard:

- To go directly to the Topic pane, press Alt+K, and then press Tab.
- In the Topic pane, to go to the next link, press Tab.
- To go directly to the Search Results view, press Alt+R, and then press the Enter or Up-Arrow key to enter the view.
- To go directly to the Navigation (Table of Contents) view, press Alt+C, and then press the Enter or Up-Arrow key to enter the view.
- To expand and collapse a node in the navigation tree, press the Right and Left-Arrow keys.
- To move to the next topic node, press the Down-Arrow or Tab key.
- To move to the previous topic node, press the Up-Arrow key or Shift+Tab.
- To go to the next link, button, or topic node from inside on of the views, press Tab.

© Copyright IBM Corp. 2001, 2013

- To scroll all the way up or down in a pane, press Home or End.
- To go back, press Alt+Left Arrow; to go forward, press Alt+Right Arrow.
- To go to the next pane, press F6.
- To move to the previous pane, press Shift+F6.
- To print the active pane, press Ctrl+P.

Related accessibility information for sight-impaired users

The following list contains hints and tips that can help you more fully use the graphical user interface:

Drop-down lists are positioned directly over or before the radio button that activates it.

If you use a screen reader, you should be aware that there are radio buttons to activate drop-down lists for several GUI pages. The way to activate the drop-down list is by selecting the associated radio button. The drop-down list is positioned directly over or before the radio button that activates it. When you use a screen reader that processes the fields and controls of a page sequentially, you might select the radio button, but not know that the associated drop-down list has been activated. The screen reader processes inactive drop-down lists first, and then processes the next radio button. The drop-down list is activated if you select the radio button.

On the following pages, keep in mind that radio buttons activate a drop-down list:
- Administration
- ESS/DS Paths
- Sessions
- Session Details
- Storage Systems

Tables are best understood by reviewing the surrounding text and the table row and column number of the table.

On some graphical user pages, tables use the header or row ID attributes when reading a single cell. The screen reader reads the table row and column number, along with cell data. Therefore, you can infer the column header and row ID.

Experiment with and fine-tune the way your screen reader pronounces some of the product abbreviations.

Your screen reader might pronounce abbreviations as if they were words. For example, the common abbreviation for Enterprise Storage Server is ESS. Your screen reader might read ESS as the word "ess". With some screen readers you can hear alternate pronunciations. If you frequently use the software you might prefer to fine-tune such associations in your settings. When an association is created, the screen reader can recognize the abbreviation as a word. If you can add dictionary words with your screen reader, replace the capitalized character sequence with the sequence E space S space S.

Typically, this abbreviation is used in the combination form of ESS/DS. This term refers to the Enterprise Storage Server 800, the DS6000, or the DS8000.

Some decorative artifacts might persist if the cascading style sheet is disabled.

Enable cascading style sheets when possible; otherwise, some decorative elements might persist in the web browser GUI. These artifacts do not affect performance. If they become too distracting, consider using the command-line interface instead.

For efficiency, confirmation dialogs place initial focus on the Yes button.

When a confirmation dialog box is displayed, focus is given to the **Yes** button. Therefore, the screen reader reads "Yes" but does not read the confirmation text. The software processes the information in this way when you do the following types of tasks:

- Perform an action on a session
- Remove a connection to a storage system
- Click the **About link**
- Create a high-availability connection

To read the confirmation text before clicking the **Yes**, **No**, or **OK** button, view the previous heading before the button.

Dojo components are not read by all screen readers.

The Job Access for Windows and Speech (JAWS) screen reader does not read some Dojo components on Windows Internet Explorer. Use the command-line interface instead of the GUI with JAWS on Windows Internet Explorer.

Firefox is the preferred browser for use with a screen reader.

Use Firefox as the screen reader because other browsers might not fully expose assistive technology content to the screen reader.

Glossary

This glossary includes terms and definitions for IBM Tivoli Storage Productivity Center.

The following cross-references are used in this glossary:

- See refers you from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- See also refers you to a related or contrasting term.

To view glossaries for other IBM products, go to www.ibm.com/software/globalization/terminology (opens in new window).

A

active management server
A management server from which the storage environment can be monitored and managed. The active management server replicates its database to the standby server.

B

business continuity
The capability of a business to withstand outages and to operate mission-critical services normally and without interruption in accordance with predefined service-level agreements.

C

CIM See Common Information Model.

CIM agent
The code that consists of common building blocks that can be used instead of proprietary software or device-specific programming interfaces to manage devices that are compliant with the Common Information Model (CIM).

CKD See count key data.

cluster

1. In Storwize® V7000, a pair of nodes that provides a single configuration and service interface.
2. A loosely coupled collection of independent systems (or nodes) organized into a network for the purpose of sharing resources and communicating with each other.
3. A collection of complete systems that work together to provide a single, unified computing capability.
4. In IBM System Storage DS8000, a partition capable of performing all DS8000 series functions. With two clusters in the DS8000 storage unit, any operational cluster can take over the processing of a failing cluster.

Common Information Model (CIM)
An implementation-neutral, object-oriented schema for describing network management or systems management information. The Distributed Management Task Force (DMTF) develops and maintains CIM specifications.

community name
The part of an SNMP message that represents a password-like name and that is used to authenticate the SNMP message.

consistency group
A group of copy relationships between virtual volumes or data sets that are maintained with the same time reference so that all copies are consistent in time.

copy set
The set of source volumes or target volumes involved in a FlashCopy operation.

count key data (CKD)
In mainframe computing, a data-record format employing self-defining record formats in which each record is represented by up to three fields: a count field identifying the record and specifying its format, an optional key field that can be used to identify the data area contents, and an optional data field that typically contains the user data.

© Copyright IBM Corp. 2001, 2013

D

data collection
The process of obtaining performance and availability monitoring data and providing that data to a metric evaluator. Examples of data collectors include Domain Name System (DNS) probes, web page analyzers, or database analyzers. See also discovery.

data source
A storage resource or agent that provides data about a storage environment.

discovery
The process of finding resources within an enterprise, including finding the new location of monitored resources that were moved. See also data collection.

E

ECKD See extended count key data.

enterprise repository
A component of the data server that records and stores all information about the monitored computers' storage assets and their usage over time. The repository is organized into relational database tables and is accessed by the data server using Java Database Connectivity (JDBC).

event An occurrence of significance to a task or system. Events can include completion or failure of an operation, a user action, or the change in state of a process.

extended count key data (ECKD)
An extension of the count-key-data (CKD) architecture. It includes additional commands that can be used to improve performance.

F

FlashCopy
An optional feature of the Storage System DS family that can make an instant copy of data, that is, a point-in-time copy of a volume.

freeze An operation in which a storage system blocks I/O from the host system to the affected volumes on the primary site. A freeze operation stops mirroring between the primary and secondary volumes to ensure data consistency at the secondary site. See also thaw.

G

GB See gigabyte.

GiB See gibibyte.

gibibyte (GiB)
In digital information storage, a base-2 unit of measurement equal to 1,073,741,824 bytes (2 to the 30th power). See also gigabyte.

gigabyte (GB)
For processor storage, real and virtual storage, and channel volume, 10 to the power of nine or 1,073,741,824 bytes. For disk storage capacity and communications volume, 1,000,000,000 bytes. See also gibibyte.

Global Copy
A non-synchronous long-distance copy option for data migration and backup. See also Global Mirror, Metro Mirror, remote mirror and copy.

Globally Unique Identifier (GUID)
An algorithmically determined number that uniquely identifies an entity within a system.

Global Mirror
An optional capability of the remote mirror and copy feature that provides a two-site extended-distance remote copy. Data that is written by the host to the storage unit at the local site is automatically maintained at the remote site. See also Global Copy, Metro Mirror, remote mirror and copy.

grain size
The unit size for allocating space on thin-provisioned volumes, such as 32, 64, 128, and 256 kibibyte. The grain size is defined when a volume is created.

GUID See Globally Unique Identifier.

H

Hardware Management Console (HMC)
In a system storage environment, a system that acts as the focal point for configuration, management of Copy Services functions, and maintenance.

heat map
A color-coded data chart in which colors are used to differentiate values in a data set.

HMC See Hardware Management Console.

host volume
A volume that represents the volume functional role from an application point of view. The host volume can be connected to a host or server. It receives read, write, and update application I/O, depending on the site to which the application is writing.

hypervisor
Software or a physical device that enables multiple instances of operating systems to run simultaneously on the same hardware.

I

in-band discovery
The process of discovering information about the storage area network (SAN), including topology and attribute data, through the Fibre Channel data paths. See also out-of-band discovery.

J

journal volume
A volume that holds a consistent copy of data until a new consistent copy is formed. The journal volume restores the last consistent point during a recovery.

K

KB See kilobyte.

KiB See kibibyte.

kibibyte (KiB)
A base-2 unit of measurement that is equal to 1,024 bytes (2 to the 10th power). See also kilobyte.

kilobyte (KB)
For processor storage, real and virtual storage, and channel volume, 2 to the power of 10 or 1,024 bytes. For disk storage capacity and communications volume, 1,000 bytes. See also kibibyte.

L

launch-in-context
An operation in which a user starts a secondary application from a primary application to perform a specific task. Using the parameters, navigation instructions, and user credentials that are supplied by the primary application, the secondary application opens to the specific place in which to complete the task.

logical unit number (LUN)
In the Small Computer System Interface (SCSI) standard, a unique identifier used to differentiate devices, each of which is a logical unit (LU).

LUN See logical unit number.

M

Management Information Base (MIB)
In the Simple Network Management Protocol (SNMP), a database of objects that can be queried or set by a network management system.

management server
A system that provides a central point of control for managing data replication.

management server relationship
A connection between two replication servers, where one server acts as the active server and replicates the data that is necessary for the standby server to take control of the replication environment.

MB See megabyte.

mebibyte (MiB)
A base-2 unit of measurement that is equal to 1,048,576 bytes (2 to the 20th power). See also megabyte.

megabyte (MB)
For processor storage, real and virtual storage, and channel volume, 2 to the 20th power or 1,048,576 bytes. For disk storage capacity and communications volume, 1,000,000 bytes. See also mebibyte.

Metro Global Mirror
A three-site, high availability, disaster recovery solution. Metro Global Mirror uses synchronous replication to mirror data between a local site and an

intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site.

Metro Mirror
A function of the remote mirror and copy feature that constantly updates a secondary copy of a volume to match changes made to a source volume. See also Global Copy, Global Mirror, remote mirror and copy.

MiB See mebibyte.

MIB See Management Information Base.

N

namespace
The scope within which a Common Information Model (CIM) schema applies.

native interface
An interface that is specific to a system or subsystem.

O

out-of-band discovery
The process of discovering storage area network (SAN) information, including topology and device data, without using the Fibre Channel data paths. A common mechanism for out-of-band discovery is the use of SNMP Management Information Base (MIB) queries, which are invoked over an IP network. See also in-band discovery.

P

ping

1. The command that sends an Internet Control Message Protocol (ICMP) echo-request packet to a gateway, router, or host with the expectation of receiving a reply.
2. A job that tracks the availability of assets and that is performed by an agent. Several ping jobs can be used to monitor the availability of any computer or subset of computers in the network.

pool A grouping of storage space that consists of volumes, logical unit numbers (LUNs), or addresses that share a common set of administrative characteristics.

practice volume
A volume that can be used to test disaster-recovery actions while maintaining disaster-recovery capability.

primordial pool
Storage capacity that is unallocated on a storage device. Storage pools are created by allocating storage capacity from primordial pools.

probe A data collection job that itemizes and creates an inventory of assets, such as computers, controllers, disk drives, file systems, and logical units.

R

recovery point objective
The maximum amount of data loss that can be tolerated during a service interruption.

remote mirror and copy
A feature of a storage server that constantly updates a secondary copy of a logical volume to match changes made to a primary logical volume. The primary and secondary volumes can be on the same storage server or on separate storage servers. See also Global Copy, Global Mirror, Metro Mirror.

resource
In a storage environment, an entity that is monitored. Resources can include fabrics, switches, computers, and storage systems.

role A job function that identifies the tasks that a user can perform and the resources to which a user has access. A user can be assigned one or more roles.

role pair
The association of two volume roles in a session that take part in a copy relationship. For example, in a Metro Mirror session, the role pair can be the association between host volumes at the primary site and host volumes at the secondary site (H1-H2).

S

scan A data collection job that monitors storage usage and file statistics on the resources in an environment.

session
A collection of source and target volumes that are managed to create consistent copies of data. The type of data replication that is associated with the session determines the actions that can be conducted for the volumes.

site awareness
The association of a location with each storage system in a session. Site awareness ensures that a volume can only be selected for a session if it matches the location of the site. Site awareness helps to prevent both reversing a hardware relationship and selecting volumes at the wrong location.

SMI-S See Storage Management Initiative Specification.

standby management server
A management server that is a backup for the active server. The replication environment cannot be monitored or managed from the standby server.

Storage Management Initiative Specification (SMI-S)
A design specification developed by the Storage Networking Industry Association (SNIA) that specifies a secure and reliable interface with which storage management systems (SMSs) can identify, classify, monitor, and control physical and logical resources in a storage area network (SAN). The interface integrates the various devices to be managed in a SAN and the tools used to manage them.

storage pod
A logical entity of components of a system consisting of two storage nodes, and one or two storage subsystems directly connected with these storage nodes.

storage resource group
A named collection of logically related resources that are monitored by Tivoli Storage Productivity Center. Monitored resources can include fabrics, switches, computers, storage systems, and other storage resource groups.

T

target volume
A volume that receives data from a host volume or another intermediate volume.

thaw An operation in which a storage system releases the block of the I/O from the host system to the affected volumes on the primary site. A thaw operation can occur after a freeze operation ends and consistency is formed on the secondary site. See also freeze.

track space-efficient volume (TSE volume)
A volume in which storage space is allocated on an as-needed basis by using space on the target volume only when tracks are copied from the source volume to the target volume.

TSE volume
See track space-efficient volume.

V

virtualization
The substitution of virtual resources for actual resources, where the virtual resources have the same functions and external interfaces as their counterparts, but differ in attributes, such as size, performance, and cost. Virtualization is commonly applied to physical hardware resources by combining multiple physical resources into shared pools from which users receive virtual resources.

virtual storage area network (VSAN)
A fabric within the storage area network (SAN).

VSAN See virtual storage area network.

Z

zone A logical grouping of switches, switch ports, and their attached devices in a fabric.

zone alias
A name that is given to a collection of one or more zone members to be managed together.

zone set

A group of zones that function together on a fabric.

Index

A

about this document 1813
Access File Summary report 1513
access load reports
  access time reports 1420
  creation time reports 1421
  modification time reports 1421
  overview 1420
access path host replacement 409
access privileges
  granting 785
  modifying 786
  removing 787
  viewing 786
Access Time Summary report 1511
accessibility features
  for Tivoli Storage Productivity Center 1745
accessibility features for users with disabilities 1747
active management servers
  about 634
  reconnecting 641
  reinstalling during active session 640
  taking over 641
adding
  CIM agents 205
  copy sets to a FlashCopy session 728
  copy sets to a Global Mirror session 733
  copy sets to a Metro Global Mirror session 735
  copy sets to a Metro Mirror session 731
  copy sets to a Snapshot session 730
  ESS/DS series logical paths 663
    using a CSV file 663
  host system connections 659
  SNMP managers 642
  storage connections
    DS6000 651, 653
    DS8000 650, 651
    SAN Volume Controller 652
    Storwize 652
    TotalStorage Enterprise Server 651, 653
    z/OS 653
adding resources 257
  fabrics 241, 293
  hypervisors 229, 271, 272
  IBM SONAS 288
  servers without agents 264
  storage systems 232, 272
    DS4000 273
    DS5000 273
    DS6000 275
    DS8000 277
    ESS 279
    other 291
    SAN Volume Controller 281
    Storwize V3500 283
adding resources *(continued)*
  storage systems *(continued)*
    Storwize V3700 283
    Storwize V7000 283
    Storwize V7000 Unified 286
    XIV 278
  switches 241, 293
administering
  users 782
advanced options, Preferences menu 480
advanced tiering 97, 429
  file usage 97
  I/O density 97
agents
  adding Storage Resource agents 210, 267
  changing credential for a Storage Resource agent 219
  checking the health 205
  CIM 82
  CIM agents 205
  deployment considerations for Storage Resource agents 83, 224
  enabling a Storage Resource agent 218
  enabling or disabling scripts for aStorage Resource agent 221
  enabling or disabling the monitoring of fabrics by aStorage Resource agent 222
  preparing Storage Resource agents for installation 83, 224
  registering with a different server 223
  scheduling a Storage Resource agent upgrade 216
  starting a Storage Resource agent upgrade 216
  testing the connection with the server where a Storage Resource agent is deployed 219
  uninstalling a Storage Resource agent 223
  upgrading Storage Resource agents 216
  viewing information aboutStorage Resource agents 214
  viewing Storage Resource agent log files 215
alert
  performance threshold 100
alert conditions and violations 341
alert log 356, 495
alert overlay 495, 582
alert tab 495, 567, 583
alerting 341
  alert log 341, 356
  alerting
    conditions and violations 341
  computer alerts 347
  database alerts 352
alerting *(continued)*
  directory alerts 349
  fabrics 345, 346, 354
  file system alerts 348
  hypervisor alerts 349
  instance alerts 351
  notifications 341
  other NAS alerts 350
  overview 75, 1663
  running scripts 555
  table alerts 353
  tablespace alerts 352
alerts
  acknowledging in the web-based GUI 359
  configuring for performance monitors 302
  configuring for probes 298
  database connection 354
  endpoint device 354
  fabric 346
  performance thresholds 345
  removing from the web-based GUI 362
  sharing for VASA provider 1641
  switch 345
  triggering conditions 1665
alias 80
analysis period
  performance data, change 435
analyze tiering
  analysis period, change 435
  conditions, setting 97
  hypervisors 434
  servers 433
  storage pools 434
  storage virtualizers 434
  volume placement 97
  volumes 435
analyze tiering wizard 96
architecture
  description 19
architecture, product 47
archive job 103
archive/backup actions 539
  creating job definitions 1426
  defining within constraints 539
  overview 539
  Reports 1426
arrays
  performance reports 1538
asset reporting
  agents 1356
  by cluster 1335
  by computer 1338, 1467
  by DBMS type 1467
  by IBM SONAS 1341
  by IBM Storwize V7000 Unified 1341
  by OS type 1351
  by storage subsystems 1351
  information in asset reports 1466

© Copyright IBM Corp. 2001, 2013

asset reporting *(continued)*
overview 1334, 1466
system-wide 1355, 1471
audience xv
automatic tiering 96, 429
availability reporting
overview 1358
awareness, site 649, 726

B

backing up
considerations for 631
backup
Tivoli Storage Productivity Center 251
backup job 103
backup reporting
backup storage requirements reports 1430
modified files not backed up reports 1430
Most at Risk Files report 1430
overview 1429
backup storage requirements reports
overview 1430
balance pools 94
Balance Pools
wizard 436
Basic HyperSwap
about 680
session commands 707
batch reports
creating 1562
creating for performance 1569
exporting and importing 1574
specifying output formats 1564
block storage provisioning
data store
creating in vSphere Web Client extension 459
datastore
See data store
on vSphere Web Client extension 458
vSphere Web Client extension 459
block storage reports
viewing in vSphere Web Client extension 1654
vSphere Web Client extension 1655
Brocade 1662
browsers supported 482

C

capacity
adding to a capacity pool 415
removing from a capacity pool 416
volumes
Storage Virtualizer - Pools Space Report 826
Storage Virtualizer - Volumes Space Report 830
Storage Virtualizers - Space Report 820
capacity pool 65, 408
adding resources to 415
capacity pool *(continued)*
creating 415
deleting 417
modifying 416
removing resources from 416
viewing capacity of 416
viewing details of 416
viewing resources of 416
capacity reporting 1360
overview 1477
change overlay 490
chargeback
overview 620
purpose 620
setting up invoices 622
charts
types 1314
checking configuration 495
CIM agents 82, 205
adding 206
collecting logs 210
removing 209
show managed devices 209
testing a connection 209
updating user credentials 208
verify they are running 210
viewing information 207
Cisco 1662
Cisco MDS 9000 81
client ports 50
cloud configuration 65, 408
collection
overview 61
columns
hiding in vSphere Web Client extension 1653
sorting in vSphere Web Client extension 1653
commands
repocopy 194
comments, sending xviii
common user, about 125
communication-failure trap descriptions 638
Compare Performance of Multiple Controllers and Modules report 991
Compare Performance of Multiple Disks report 1001
Compare Performance of Multiple IO Groups report 1012
Compare Performance of Multiple Managed Disks report 1020
Compare Performance of Multiple Nodes report 992
Compare Performance of Multiple Pools report 1032
Compare Performance of Multiple Ports report 1039
Compare Performance of Multiple RAID Arrays report 1045
Compare Performance of Multiple Storage Systems report 979
Compare Performance of Multiple Switch Ports report 946
Compare Performance of Multiple Switches report 939
Compare Performance of Multiple Volumes report 1057
Compare Performance of One Pool over Time Ranges report 1033
Compare Performance of One Storage System over Time Ranges report 980
Compare Performance of One Switch over Time Ranges report 940
Compare Performance of One Switch Port over Time Ranges report 947
components 19
computer alerts 347
overview 347
triggering conditions 1665
Computer Disk Space report 1521
computer quotas 513
concept diagram
VASA provider 74
conditions
associating with groups 325
configuration
analysis, using configuration 495
analyzing changes to 490, 493, 495
checking compliance to policies 495
history, using 490
configuration change, SNMP trap descriptions 637
configuring 66, 409
alerts for performance monitors 302
alerts for probes 298
logical paths 662
storage systems 655
Connection menu 478
connections
about 645
direct 646
Hardware Management Console (HMC) 647
modifying
host systems 660
modifying properties 654
z/OS 647
connections, displaying in topology viewer 594
considerations, back up and recovery 631
consistency group 671
time interval 671
constraint 103
constraint violation report 103
constraint violations
affected hosts 1502
affected volumes 1502
creating reports 1501
report 1428, 1509
viewing 1502
constraints
archive/backup actions
overview 539
creating 533
defining archive/backup actions 539
editing 536
overview 533
content pane
using 477
controller cache performance
performance reports 1539

controllers
performance reports 1541
controls for performance view
export performance information 390
hide and show resources 390
hide and show the legend 390
remove resources 390
select metrics 390
specify a time range 390
specify granularity 390
view chart 390
view chart in a new window 390
view table 390
copy sets
about 665
adding 728
exporting 633, 739
importing 633, 739
removing 741
supported number of role pairs and volumes 665
viewing 781
creating 411, 413
constraint violation reports 1501
FlashCopy session 728
Global Mirror session 733
Metro Global Mirror session 735
Metro Mirror session 731
profile 317
Snapshot session 730
TPC server probes 332
CSV file
exporting vSphere reports 1647
custom reports
See also reports, custom
Hitachi Universal Storage Platform V systems 886
nodes 900
ports on Hitachi Universal Storage Platform V systems 922
storage virtualizer I/O groups 908, 909
storage virtualizer nodes 900
storage virtualizer pools 922, 927
storage virtualizer ports 918
storage virtualizer volumes 890, 894
storage virtualizers 873, 877
storage-virtualizer managed disks 887, 889
volumes on Hitachi Universal Storage Platform V systems 898

D

dashboard 365, 468
Data
system reports 1511
data collection 61
creating performance monitors in the web-based GUI 300
modifying performance monitors in the web-based GUI 301
performance monitors 300
starting and stopping performance monitors in the web-based GUI 303
data collection *(continued)*
verifying that a performance monitor is running in the web-based GUI 301
verifying that a probe is running in the web-based GUI 297
viewing all performance monitors in theweb-based GUI 304
viewing performance monitor logs in the web-based GUI 305
viewing performance monitors for a resource in theweb-based GUI 304
viewing performance monitors in theweb-based GUI 304
data copying symbols 749
data exposure 671
Data Manager
overview 484
reports 1331
Data Manager for Chargeback
generating invoices 623
overview 485
planning 621
purpose 485
reports 1490
viewing invoices 624
Data Manager for Databases
reports 1462
running scripts with alerts 555
system reports 1525
what is it? 485
data migration utility 742
data path
monitoring and troubleshooting 584
Data Path Explorer
overview 584
prerequisites 584
shortcut to launching 564
Data Path Explorer view 592
data paths
changing the layout 407
exporting a data path to a file 408
fabrics 402
hypervisors 402
navigating a data path 406
overview 402
printing a view 407
resources that you can view 403
saving the view settings 407
servers 402
storage systems 402
topology view 403
topology view for data paths
changing the layout 407
exporting a data path to a file 408
navigating a data path 406
printing a view 407
saving the view settings 407
viewing information about internal resources 406
viewing information about top-level resources 406
zooming in 406
zooming out 407
viewing information about internal resources 406
data paths *(continued)*
viewing information about top-level resources 406
zooming in 406
zooming out 407
Data server
broadcasting location 143
changing DB2 password 187
configuring audit logging 143
displaying services 141
increasing memory 174
increasing memory on AIX 175
increasing memory on Linux 175
increasing memory on UNIX 175
increasing memory on Windows 176
shutting down services 143
starting a server service 144
Data Server
checking status 174
Data server service 140
data sources
IBM Tivoli Storage Productivity Center servers 331
data store
creating in vSphere Web Client extension 458
Data Studio
data studio 254
database alerts 352
description 352
database connection alerts 354
database repository
accessing views 193
views overview 191
database storage by computer report 1525, 1527, 1531, 1533
database-table space alerts
triggering conditions 1669
database-tablespace quotas 521
datastore
See data store
Datastores report
vSphere Web Client extension 1655, 1656
DB2
command line 252
monitoring 254
DB2 password
changing for Data server 187
changing for Device server 187
DB2, stop 253
default
groups 337
zone 80
deleting
storage resource group sessions 555
Device server
changing DB2 password 187
checking status 174
setting timeout values 176
device triggering conditions 345
directories reports
largest directories reports 1427
monitored directory space reports 1427
overview 1427
directory alerts 349

directory alerts *(continued)*
description 349
triggering conditions 1668
directory groups
creating 339
overview 339
disabilities, accessibility features for users with 1747
disabling
Metro Mirror heartbeat 739
Open HyperSwap 705
OpenSwap 705
disaster recovery, practicing 745
discover 306
discovery 489
fabrics 489
Disk Defect report 1514
Disk Manager
overview 486, 487
reports 1490
Disk Space Summary report 1512
Disks Capacity report 994
Distributed Management Task Force (DTMF) 1741
domain account
how to grant local administrative privileges 190
domain port zoning 80
Domain_ID 80
drop-down lists, limitations of to the sight-impaired 1747
DS6000
port thresholds 1697
DS8000
port thresholds 1697

E

Easy Tier heat map transfer
about 644
configuring for 656
results, viewing 658
editing
storage usage quotas 523
storage usage quotas for databases 525
enabling, Metro Mirror heartbeat 739
endpoint device alerts
creating 354
triggering conditions 1677
entities in topology viewer
arranging in a group 574
displaying information about 594
labels, specifying 595
locating 595
pinning 596
removing from database 597
selecting 598
ESS/DS storage systems
adding logical paths 663
using a CSV file 663
removing logical paths 665
exported file systems 306
exporting
copy sets 633, 739
historical data for Global Mirror role pairs 740
exporting and importing
batch reports 1574
exporting information in the web-based GUI 441

F

fabric 79
fabric alerts
creating 346
triggering conditions 1675
fabric connections
viewing in vSphere Web Client extension 1654
Fabric Manager
alerting 345, 346, 354
create fabric groups 338
overview 488
reports 1502
system reports 1554
zone 80
zone set 80
Fabric Performance Report - By Port 1507
Fabric zone configuration 1662
fabric.mib file 1743
fabrics
adding 241, 293, 294
agent types for monitoring 294
alerting 345, 346, 354
create performance monitors 1503
discovery 489
groups 338
modifying properties 437
removing 244
testing connection 244
updating credentials 242
viewing data paths 402
viewing information 242, 243
viewing information in the web-based GUI 363
viewing status in the web-based GUI 374
failback 665, 671
failover 665, 671
FCID zoning 80
federated repositories
adding an LDAP repository 148
removing an LDAP repository 164
fibre channel
methodologies of interconnects 1744
file filters
wildcards 326
File menu 477
file share
access path 409
file storage provisioning 462
on vSphere Web Client extension 461
file system alerts
description 348
triggering conditions 1666
file system quotas 516
file system reports
viewing in vSphere Web Client extension 1656
File system reports
vSphere Web Client extension 1656
File System to Volume Relationships report 935
File Systems Capacity report 967
file usage
volumes, tiering 97
files reports
duplicate files reports 1422
file size distribution reports 1424
file summary reports 1425
file types reports 1425
largest files reports 1422
most obsolete files reports 1423
orphan files reports 1424
overview 1422
filtering lists in the web-based GUI 443
filtering resources in topology viewer 595
FlashCopy
about 681
creating a session 728
session commands 707
full disk encryption
icons, storage 569
support for 1663
fully qualified host name
checking for on AIX 188
checking for on Linux 189
checking for on Solaris 190
Windows, verify 188

G

Global Mirror
about 689
creating a session 733
session commands 712
session details, viewing 755
Global Mirror with Practice, session commands 715
glossary 1751
granting
access privileges 785
groups 1328
creating a storage subsystem 599
creating invoices for 334
Data Manager for Chargeback 334
default objects 337
fabrics 338
monitoring 334
storage subsystem 599
groups in topology viewer
creating custom 583
default 565
groups, using 1328
GUI
managing resources in the stand-alone GUI 464
menu bar 477
tool bar 483

H

hardware zoning 80
health, viewing summary of 748
heartbeat, enabling and disabling 739
Help menu 481

historical reports
Storage Virtualizer - Pools Historical Space Report 860
Storage Virtualizer - Volumes Historical Space Report 863
Storage Virtualizers - Historical Space Report 858
historical trending 1317
configuration history 490
exporting history data to .csv 1320
history charts 1317
host systems
adding a connection to 659
managing 659
modifying a connection to 660
removing 661, 662
viewing connection information for 782
hypervisor alerts 349
triggering conditions 1678
Hypervisor Data Stores Capacity report 956
Hypervisor Disks Capacity report 957
hypervisors 434
adding 229, 271, 272
modifying properties 437
tiering, analyze 434
updating credentials 231
viewing data paths 402
viewing information in the web-based GUI 363
viewing status in the web-based GUI 374
Hypervisors Capacity report 949

I

I/O density
volumes, tiering 97
I/O groups
performance reports 1542
performance statistics 846
Storage Virtualizers - Most Active IO Groups Report 834
IBM Tivoli Storage Productivity Center
architecture 19
component overview 19
menu bar 477
overview 19
tool bar 483
IBM Tivoli Storage Productivity Center for Replication
starting 629
IBM Tivoli Storage Productivity Center servers
overview 331
icons
session status 749
storage systems status 643
identifying the version of IBM Tivoli Storage Productivity Center for Replication 631
importing copy sets 633, 739
installing
Storage Resource agent considerations 83, 224
installing *(continued)*
Storage Resource agents remotely 83, 224
instance alerts 351
overview 351
triggering conditions 1669
instance quotas 518
instance storage by computer report 1536
interface
content pane 477
navigation tree 476
navigation tree nodes 466
interfaces
changes in V5.2 38
display differences 373
overview 33
when to use 34
Internet Engineering Task Force (IETF) 1743
invoice 620, 621
samples 624
invoices
generating 623
setting up 622
using the Online Execution window 623
viewing 624

J

journal volume, restoring data from 656

L

launch in context
available windows 1731
context parameters 1729
context task parameters 1731
starting from the command line 615
starting with Java Web Start 613
using outbound 92, 619
LDAP repository
adding to federated repositories 148
removing 164
levels of detail in topology viewer 563
license
check 138
license restrictions 23, 128
licenses 23
Linux
IBM Data Studio, start 254
locations
modifying for a site 741
modifying for a storage system 654
logical paths
adding 663
using a CSV file 663
configuring 662
removing 665
viewing 663, 781
logical SANs (LSANs), viewing topology 563
logical storage subsystems
viewing volumes 781
LSANs (logical SANs), viewing topology 563
LTPA token expiration 173

M

managed disk groups
performance reports 1543
managed disks
performance statistics 848
Storage Virtualizer - MDisk Details Report 847
Storage Virtualizers - Most Active MDisks Report 835
Storage Virtualizers - Most Active Nodes Report 837
Storage Virtualizers - Most Active Pools Report 838
Storage Virtualizers - Most Active Ports Report 842
Storage Virtualizers - Most Active Volumes Report 840
Managed Disks Capacity report 825, 1013
Management Information Base (MIB) files 1743
management servers
about 634
reconnecting 641
setting as the standby server 639
setting up a remote server as the standby server 640
state change SNMP trap descriptions 638
taking over 641
viewing health summary 748
managing
host systems 659
storage subsystems 599
storage systems 643
MDisk Group Details report 1084
MDisk Group - VDisk Workload Activity Details report 1080
MDisk groups
performance reports 1543
MDisk Groups - VDisk Workload Activity report 1067
MDisk Groups - Workload Activity report 1087
MDisks
performance statistics 848
Storage Virtualizer - MDisk Details Report 847
Storage Virtualizers - Most Active MDisks Report 835
memory
increasing allocation for Data server on AIX 175
increasing allocation for Data server on Linux 175
increasing allocation for Data server on Windows 176
memory allocation
Data server 174
menu bar 477
Connection menu 478
File menu 477

menu bar *(continued)*
Help menu 481
Preferences menu 478
View menu 478
Windows menu 480
messages, viewing 782
metricsperformance reports 1701
Metro Global Mirror
about 694
creating a session 735
session commands 718
Metro Global Mirror scenarios
disaster recovery (unplanned outages) 113
planned outages 110
practice 116
Metro Global Mirror with Practice, session commands 722
Metro Mirror
about 683
creating a session 731
session commands 709
Metro Mirror heartbeat, enabling and disabling 739
Metro Mirror with Practice, session commands 711
minimap in Configuration History view 495
minimap in topology viewer 564
modified files not backed up reports 1430
modifying
access privileges 786
host system connections 660
site location 741
storage connection properties 654
storage system location 654
monitor
creating a storage-subsystem group 599
monitoring 1690
groups 334
performance with Data Path Explorer 584
Monitoring Agent
displaying reports 89, 1636
most active
I/O groups
storage virtualizers 834
managed disks
storage virtualizers 835
MDisks
storage virtualizers 835
nodes
storage virtualizers 837
pools
storage virtualizers 838
ports
storage virtualizers 842
Storage Virtualizers - Most Active MDisks Report 836
volumes
storage virtualizers 840
Most Active Controllers or Modules report 983
Most Active Disks report 996
Most Active Host Connections report 1002
Most Active Hypervisors report 951
Most Active IO Groups report 1008
Most Active Managed Disks report 1015
Most Active Nodes report 987
Most Active Pools report 1028
Most Active Ports report 1035
Most Active RAID Arrays report 1040
Most Active Servers report 961
Most Active Storage Systems report 975
Most Active Switch Ports report 942
Most Active Volumes report 1053
Most at Risk Files report 1430, 1518
Most Obsolete Files report 1518

N

N_Port address identifier 80
N_Port_Name 80
NAS alerts
triggering conditions 1678
NAS mounts report
vSphere Web Client extension 1656
navigation tree
using 476
Network Appliance OS user quotas
viewing 531
Network Appliance quotas
creating 528
viewing directory (QTree) quotas and violations 532
Network Appliance user quotas
viewing 531
network quotas 508, 511
Node_Name 80
nodes
most active 837
performance statistics 850
Storage Virtualizer - Node Details Report 850
Storage Virtualizers - Most Active Nodes Report 837
non standard zones
fabric assist 82
protocol 82
quick loop 82

O

objects
default 337
Oldest Orphaned Files report 1519
Online Execution window 623
Open HyperSwap
disabling 705
Opening the management GUI for a resource 440
OpenSwap
disabling 705
OS Disk Capacity By Computer
sample invoice 625
OS Storage Usage By User
sample invoice 624, 625
OS user groups reports 1427
other NAS alerts 350
other NAS alerts *(continued)*
overview 350
overview 66

P

passwords
changing 177
changing DB2 password for Data server 187
changing DB2 password for Device server 187
changing using password tool 178, 180, 182, 184
paths
adding 663
using a CSV file 663
removing 665
PDF file, saving as 1312
performance 1690
controls for the performance view 390
defining thresholds 397
hot spots 93
metrics
disk 1701
port 1701
volume 1701
monitoring 383
monitoring with Data Path Explorer 584
performance view 388
storage virtualizer reports
managed disks 847
pools 856
Storage Virtualizers - Most Active IO Groups Report 834
viewing alerts for threshold violations 401
viewing thresholds 400
performance chart
viewing in vSphere Web Client extension 1647
vSphere Web Client extension 1648
performance data
analysis period, change 435
Performance Data Export report 982
performance monitor, creating 1503
Performance of One Controller or Module report 985
Performance of One Disk report 998
Performance of One IO Group report 1010
Performance of One Managed Disk report 1017
Performance of One Node report 988
Performance of One Pool report 1030
Performance of One Port report 1037
Performance of One RAID Array report 1043
Performance of One Storage System report 977
Performance of One Switch Port report 943
Performance of One Switch report 937
Performance of One Volume report 1055

Performance of Volumes by Host Connection report 1006
Performance of Volumes by Hypervisor report 954
Performance of Volumes by Server report 964
performance overlay 495
performance reporting
  Hitachi Universal Storage Platform V systems 886
  nodes 900
  ports on Hitachi Universal Storage Platform V systems 922
  storage virtualizer I/O groups 908, 909
  storage virtualizer nodes 900
  storage virtualizer pools 922, 927
  storage virtualizer ports 918
  storage virtualizer volumes 890, 894
  storage virtualizers 873, 877
  storage-virtualizer managed disks 887, 889
  volumes on Hitachi Universal Storage Platform V systems 898
performance reports
  array performance 1538
  controller cache performance 1539
  controller performance 1541
  creating 1497
  I/O groups 1542
  managed disk groups 1543
  port performance 1546
  predefined reports 1537
  SAN Volume Controller node cache 1544
  storage subsystem performance 1548
  top n active volumes cache hit performance 1549
  top volumes data rate 1550
  top volumes disk performance 1550
  top volumes I/O rate 1552
  top volumes response performance 1553
performance statistics
  volumes 855
performance thresholds 345, 1690
performance triggering conditions 345
performance view
  accessing 388
  controls 390
  viewing the performance of internal resources 394
  viewing the performance of related resources 394
performing
  takeover 641
pings
  creating 327
    pings 327
pinning entities in topology viewer 596
planned outage of remote (H3) and return 112
planned outage of the local site (H1) while in normal configuration, with a production move to the remote site (H3) 111
planned outage of the local site (H1), with a production move to the intermediate site (H2) and an optional return 110
planned outage on site 2 (H2) and return three sites to the original configuration 111
planning
  authorization for users 125
  for Data Path Explorer view 592
policies
  tiering 428
policy management
  constraints 533
  creating computer quotas 513
  creating constraints 533
  creating database-tablespace quotas 521
  creating file system quotas 516
  creating instance quotas 518
  creating Network Appliance quotas 528
  creating network quotas 508
  creating network quotas for databases 511
  editing constraints 536
  editing storage a usage quota 523
  editing storage a usage quota for databases 525
  enforcing storage policies 505
  quotas 507, 527
  scheduled actions 542
  viewing imported Network Appliance OS user quotas 531
  viewing imported Network Appliance user quotas 531
  viewing Network Appliance directory (QTree) quotas and violations 532
pool utilization 429
pools
  balance 94, 96
  balance, analyze 436
  hot spots, identify 94
  modifying properties 437
  most active 839
  performance statistics 852
  space usage 861
  Storage Virtualizer - Pool Details Report 852
  Storage Virtualizer - Pools Space Report 826, 827
  Storage Virtualizers - Most Active Pools Report 838
  tier 96
  tiering, analyze 434
Pools Capacity report 1021
Pools Historical Capacity report 868, 1026
Port Connections report 1554
port name zoning 80
port performance
  performance reports 1546
port zoning 80
PortId zoning 80
ports 50
  most active 842
  performance statistics 857
ports *(continued)*
  Storage Virtualizers - Most Active Ports Report 842
practice at H3 for Metro Global Mirror 118
practice recover to H1 with production continuing on H2 117
practice recover to H2 with production continuing on H1 in a normal configuration 116
practice volumes 671, 745
practicing disaster recovery 745
predefined performance reports 1537
predefined reports
  historical information 858
  Storage Virtualizer - IO Group Details Report 845, 846
  Storage Virtualizer - MDisk Details Report 847, 848
  Storage Virtualizer - Node Details Report 850
  Storage Virtualizer - Performance Details Report 844
  Storage Virtualizer - Pool Details Report 852
  Storage Virtualizer - Pools Historical Space Report 860, 861
  Storage Virtualizer - Pools Space Report 826, 827
  Storage Virtualizer - Port Details Report 856, 857
  Storage Virtualizer - Volume Details Report 854
  Storage Virtualizer - Volumes Historical Space Report 863, 864
  Storage Virtualizer - Volumes Space Report 830, 831
  Storage Virtualizer Performance Details Report 843
  Storage Virtualizers - Historical Space Report 858, 859
  Storage Virtualizers - Most Active IO Groups Report 834
  Storage Virtualizers - Most Active MDisks Report 835, 836
  Storage Virtualizers - Most Active Nodes Report 837
  Storage Virtualizers - Most Active Pools Report 838, 839
  Storage Virtualizers - Most Active Ports Report 842
  Storage Virtualizers - Most Active Volumes Report 840, 841
  Storage Virtualizers - Space Report 820, 821
Preferences menu 478
Preserve Mirror option, recommendations for 727
primordial 567
printing 1312
probe 103
probes
  configuring alerts 298
problems
  report xix
product
  license, check 138

product *(continued)*
version, check 138
product architecture 47
profiles
associating conditions with groups 325
creating 317
using wildcards 326
properties
session
SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified sessions 776
TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 sessions 756
XIV system sessions 779
viewing session 756
protected volumes, about 649
protecting volumes 655
Provision Storage wizard 418
provisioning 306, 418
custom requirements for 410
overview 71, 417
scenarios
constrained set of storage 106
using vSphere Web Client extension 106
shares 420
volumes 419

Q

QLogic 1662
quota violations reporting
all quotas reports 1429
computer quotas reports 1429
file system quotas reports 1429
network-wide quotas reports 1429
quotas
creating computer quotas 513
creating database-tablespace quotas 521
creating file system quotas 516
creating instance quotas 518
creating Network Appliance quotas 528
creating network quotas 508
creating network quotas for databases 511
editing a storage usage quota 523
editing a storage usage quota for databases 525
overview 507, 527
viewing imported Network Appliance OS user quotas 531
viewing imported Network Appliance user quotas 531
viewing Network Appliance directory (QTree) quotas and violations 532

R

read I/O capability
calculating for MDisk groups 1092
calculating for pools 1092
formula 1092
reader feedback, sending xviii
recommendations for using the Preserve Mirror option 727
reconnecting the active and standby management servers 641
recovering, considerations for 631
refreshing, storage system configuration 655
removing
copy sets 741
fabrics 244
host systems 661, 662
logical paths 665
sessions 741
storage system 241
storage systems 653
switches 244
user access 787
replication alerts
triggering conditions 1680
repocopy command 194
exporting data 195
importing data 196
repocopy tool
exporting data 195
report
constraint violation 1501
storage subsystem 1495
storage subsystem constraint violation 1501
storage subsystem performance 1497
report displays
customize in vSphere Web Client extension 1653
reporting 1328
"No data available" message 1292
archive location, specifying 807
archive/backup 1426
asset reports 1334, 1466
availability 1358
backup 1429
backup copies of reports, creating 809
build ID, showing 1294
by clauses 1333, 1465
capacity 1360
capacity reports 1477
categories 1295, 1297
CM-REQ-4158 error message 1293
Cognos
See Tivoli Common Reporting
configuring, overview 804
connections to TPCDB, multiple 808
considerations 1307
copying a predefined report 805
creating a folder 809
creating custom reports 798
creating reports from within Tivoli Common Reporting 799
custom reports
See reports, custom
dashboard 365, 468
reporting *(continued)*
DPR-ERR-2109 error message 1292
drilling down 1316
email
configuring to send reports by 806
sending report by 806
file system
configuring report to save to 808
saving reports to 807
specifying archive location 807
specifying locations for output 807
specifying root directory 807
filters, limit of 5000 items 1293
format and delivery options, setting 804
generating reports 1320
generating rollup reports 1576
getting started 789
historical trending 1317
history charts 1317
IBM Cognos Configuration GUI, starting 806
limitations 1307
logo, customizing 809, 810
long time to complete performance report 1294
long time to display report 1293
overview 1295, 1297
package version number, showing 1294
PDF output, maximum number of columns 1290
Portuguese (Brazilian) text in reports 1294
predefined reports
See reports, predefined
printing 1312
Query Studio, performance report 1294
remote connection, logging off 1293
report contains no data 1293
report titles
customizing 809
modifying 810
report views 804
reporting
categories 1295
scenarios
analyzing hypervisor volumes 801
identifying relationships 803
investigating pool performance 803
monitoring volume performance 802
scheduling a report with predefined options 804
search index, creating 811
security
access to functions, restricting 817
access to reports, specifying 816
administrative privileges, assigning 814
administrative privileges, restricting 815

reporting *(continued)*
security *(continued)*
configuring for Tivoli Common Reporting 812
roles 814, 815
user IDs, creating 812
WebSphere Application Server, user IDs 812
selecting an object 1324
sorting columns 1312
Storage Volume properties 1292
tabs, navigating 1316
Tivoli Common Reporting
access to functions, restricting 817
access to reports, specifying 816
administrative privileges, assigning 814
administrative privileges, restricting 815
configuring security 812
creating and viewing reports within 799
opening 813
password change 1291
relationship to Cognos 789
timeout 1289
TPCDB
multiple connections to 808
password change 1291
password prompt 1291
troubleshooting 1289
types 1295
usage 1420
usage reports 1483
usage violation 1427, 1488
viewing predefined reports 798
viewing reports from within Tivoli Common Reporting 799
what can I find out? 1303
reporting problems xix
reports
Data Manager 1331
Data Manager for Chargeback 1490
Data Manager for Databases 1462
Disk Manager 1490
Fabric Manager 1502
storage tier 1058
Storage Virtualizer - IO Group Details Report 845, 846
Storage Virtualizer - MDisk Details Report 847
Storage Virtualizer - Node Details Report 850
Storage Virtualizer - Performance Details Report 844
Storage Virtualizer - Pool Details Report 852
Storage Virtualizer - Pools Historical Space Report 861
Storage Virtualizer - Pools Space Report 826, 827
Storage Virtualizer - Port Details Report 856, 857
Storage Virtualizer - Volume Details Report 854
Storage Virtualizer - Volumes Historical Space Report 864
reports *(continued)*
Storage Virtualizer - Volumes Space Report 830, 831
Storage Virtualizer – Volume Details Report 855
Storage Virtualizer Performance Details Report 843
Storage Virtualizers - Historical Space Report 859
Storage Virtualizers - Most Active IO Groups Report 834
Storage Virtualizers - Most Active MDisks Report 835, 836
Storage Virtualizers - Most Active Nodes Report 837
Storage Virtualizers - Most Active Pools Report 838, 839
Storage Virtualizers - Most Active Ports Report 842
Storage Virtualizers - Most Active Volumes Report 840, 841
Storage Virtualizers - Space Report 820, 821
system 1511, 1525
threshold values 1094
upgrade from V5.1, V5.1.1, or V5.1.1.1 to V5.2 819
VMware capacity 1460
vSphere 1643
vSphere Web Client extension 1643
reports, custom
"No data available" message 1292
block server storage systems, performance 1119
capacity and relationships reports
creating 1219
overview 1218
clusters
hypervisors, capacity 1234
servers, capacity 1224
controllers, storage systems
block server systems, performance 1158
DS8000 and DS6000, performance 1153
creating 798
data stores on hypervisors, capacity 1236
DS8000 and DS6000 storage systems, performance 1109
file system groups on servers, capacity 1231
file systems
hypervisors, capacity 1234
NAS systems, capacity 1241
servers, capacity 1227
filters
adding to capacity and relationships reports 1220
hypervisors 1222, 1223
limit of 5000 items 1293
servers 1221
storage resource groups 1224
storage systems 1223
help 1291
reports, custom *(continued)*
host connections, back-end storage
SAN Volume Controller, performance 1203
Storwize V7000 Unified, performance 1203
Storwize V7000, performance 1203
host connections, storage systems
DS8000 and DS6000, performance 1206
XIV systems, performance 1209
hypervisors, capacity 1232
I/O groups
SAN Volume Controller, performance 1191
Storwize V7000 Unified, performance 1191
Storwize V7000, performance 1191
IBM SONAS 1246
disks, capacity 1243
pools, capacity 1245
local disks, storage systems
capacity 1266
performance 1200, 1201
logical volumes
NAS systems, capacity 1243
servers, capacity 1228
managed disks
SAN Volume Controller, performance 1179
Storwize V7000 Unified, performance 1179
Storwize V7000, performance 1179
managed disks, storage systems
capacity 1288
performance 1178
managed disks, storage virtualizers
capacity 1264
modules, storage systems
XIV systems, performance 1160
NAS disks, capacity 1243
NAS exports, capacity 1242
NAS systems, capacity 1239
NetApp disks, capacity 1243
nodes
SAN Volume Controller, performance 1182
Storwize V7000 Unified, performance 1182
Storwize V7000, performance 1182
overview 797
performance reports
creating 1095
overview 1095
pools
SAN Volume Controller, performance 1168
Storwize V7000 Unified, performance 1168
Storwize V7000, performance 1168

reports, custom *(continued)*
pools, storage systems
DS8000 and DS6000, performance 1171
XIV systems, performance 1175
ports
SAN Volume Controller, performance 1139
Storwize V7000 Unified, performance 1139
Storwize V7000, performance 1139
ports, storage systems
block server systems, performance 1147
DS8000 and DS6000, performance 1142
XIV systems, performance 1146
SAN Volume Controller, performance 1101
server disks on hypervisors, capacity 1237
server disks, capacity 1230
server groups, capacity 1231
servers, capacity 1225
storage host connections, performance 1202
storage resource groups, capacity 1289
storage system controllers, performance 1153
storage system disks, capacity 1285
storage system groups, capacity 1287
storage system I/O groups, performance 1190
storage system modules, performance 1159
storage system nodes, performance 1181
storage system pools, capacity 1272
storage system pools, performance 1163
storage system ports, performance 1137
storage system RAID arrays, performance 1147, 1148
storage system volumes, capacity 1278
storage system volumes, performance 1120
storage systems, capacity 1268
storage systems, performance 1097
storage virtualizer groups, capacity 1267
storage virtualizer pools, capacity 1251
storage virtualizer volumes, capacity 1257
storage virtualizers, capacity 1247
Storage Volume properties 1292
Storwize V7000 Unified
disks, capacity 1243
file sets, capacity 1246
pools, capacity 1245
Storwize V7000 Unified, performance 1101
Storwize V7000, performance 1101
reports, custom *(continued)*
switch ports, performance 1215, 1216
switches, performance 1212, 1213
virtual machines managed by hypervisors, capacity 1238
volume groups on servers, capacity 1229
volumes
SAN Volume Controller, performance 1126
Storwize V7000 Unified, performance 1126
Storwize V7000, performance 1126
volumes, storage systems
block server systems, performance 1136
DS8000 and DS6000, performance 1130
XIV systems, performance 1133
XIV systems, performance 1116
reports, predefined
Compare Performance of Multiple Controllers and Modules 991
Compare Performance of Multiple Disks 1001
Compare Performance of Multiple IO Groups 1012
Compare Performance of Multiple Managed Disks 1020
Compare Performance of Multiple Nodes 992
Compare Performance of Multiple Pools 1032
Compare Performance of Multiple Ports 1039
Compare Performance of Multiple RAID Arrays 1045
Compare Performance of Multiple Storage Systems 979
Compare Performance of Multiple Switch Ports 946
Compare Performance of Multiple Switches 939
Compare Performance of Multiple Volumes 1057
Compare Performance of Multiple Volumes, scenario 802
Compare Performance of One Pool over Time Ranges 1033
Compare Performance of One Pool over Time Ranges, scenario 803
Compare Performance of One Storage System over Time Ranges 980
Compare Performance of One Switch over Time Ranges 940
Compare Performance of One Switch Port over Time Ranges 947
copying 805
Disks Capacity 994
File System to Volume Relationships 935
File Systems Capacity 967
Hypervisor Data Stores Capacity 956
Hypervisor Disks Capacity 957
Hypervisors Capacity 949
Managed Disks Capacity 825, 1013
reports, predefined *(continued)*
MDisk Group Details 1084
MDisk Group - VDisk Workload Activity Details 1080
MDisk Groups - VDisk Workload Activity 1067
MDisk Groups - Workload Activity 1087
Most Active Controllers or Modules 983
Most Active Disks 996
Most Active Host Connections 1002
Most Active Hypervisors 951
Most Active Hypervisors, scenario 801
Most Active IO Groups 1008
Most Active Managed Disks 1015
Most Active Nodes 987
Most Active Pools 1028
Most Active Ports 1035
Most Active RAID Arrays 1040
Most Active Servers 961
Most Active Storage Systems 975
Most Active Switch Ports 942
Most Active Volumes 1053
overview 789
Performance Data Export report 982
performance metrics, searching for 800
Performance of One Controller or Module 985
Performance of One Disk 998
Performance of One IO Group 1010
Performance of One Managed Disk 1017
Performance of One Node 988
Performance of One Pool 1030
Performance of One Port 1037
Performance of One RAID Array 1043
Performance of One Storage System 977
Performance of One Switch 937
Performance of One Switch Port 943
Performance of One Volume 1055
Performance of Volumes by Host Connection 1006
Performance of Volumes by Hypervisor 954
Performance of Volumes by Hypervisor, scenario 801
Performance of Volumes by Server 964
Pools Capacity 1021
Pools Historical Capacity 868, 1026
properties, searching for 800
report views 804
scenarios
analyzing hypervisor volumes 801
identifying relationships 803
investigating pool performance 803
monitoring volume performance 802
scheduling with predefined options 804

reports, predefined *(continued)*
security, specifying access to reports 816
Server Disks Capacity 966
Servers Capacity 959
Storage Resource Group - VDisk Workload Activity 1075
Storage Resource Group - VDisk Workload Activity Details 1059
Storage Resource Relationships Summary 930
Storage Resource Relationships Summary (Configurable) 932
Storage Resource Relationships Summary, scenario 803
Storage Systems Capacity 969
Storage Systems Historical Capacity 866, 973
Storage Virtualizer - VDisk Workload Activity 1071
Summarized Performance of Volumes by Host Connection 1004
Summarized Performance of Volumes by Hypervisor 952
Summarized Performance of Volumes by Hypervisor, scenario 801
Summarized Performance of Volumes by Server 962
VDisk Details 1064
viewing descriptions of 798
viewing predefined reports 798
Volumes Capacity 1047
Volumes Historical Capacity 872, 1051
Repository Copy tool 194
importing data 196
resources
adding 257
adding agentless servers for monitoring 264
adding fabrics for monitoring 241, 293
adding hypervisors for monitoring 229, 271, 272
adding storage systems for monitoring 232, 272
DS4000 273
DS5000 273
DS6000 275
DS8000 277
ESS 279
other 291
SAN Volume Controller 281
Storwize V3500 283
Storwize V3700 283
Storwize V7000 283
Storwize V7000 Unified 286
XIV 278
adding switches for monitoring 241, 293
adding to a capacity pool 415
discovery 489
modifying properties 437
removing from a capacity pool 416
Storage Virtualizer - Pools Space Report
capacity 826
resources *(continued)*
Storage Virtualizer - Pools Space Report *(continued)*
run 826
used and available space 826
Storage Virtualizer - Volumes Space Report
capacity 830
run 830
used and available space 830
Storage Virtualizers - Space Report
capacity 820
run 820
used and available space 820
updating credentials for hypervisors 231
updating credentials for storage systems 234
restoring data from a journal volume 656
restrictions
based on license 23, 128
based on role 23, 128
return codes 1659
role pairs 670
role pairs, status messages 754
role restrictions 23, 128
roles 78, 126
run
Storage Virtualizer - Pools Historical Space Report 860
Storage Virtualizer - Volumes Historical Space Report 863
Storage Virtualizers - Historical Space Report 858

S

SAN (storage area network), viewing topology 563
SAN Assets (ALL) report 1555
SAN Assets (Connected Devices) report 1555
SAN Assets (Switches) report 1555
SAN Volume Controller
publications xv
SAN Volume Controller nodes
cache performance reports 1544
scans
creating 307
scans 307
for archive 103
monitoring
scans 306
overview 306
scenarios
overview of scenarios for practicing at H3 for Metro Global Mirror 119
practice scenario 1 119
practice scenario 2 119
practice scenario 3 120
planned outage of remote (H3) and return 112
planned outage of the local site (H1) with a production move to the remote site (H3) 111
scenarios *(continued)*
planned outage of the local site (H1), with a production move to the intermediate site (H2) and an optional return 110
planned outage on site 2 (H2) and return three sites to the original configuration 111
practice at H3 for Metro Global Mirror 118
practice recover to H1 with production continuing on H2 117
practice recover to H2 with production continuing on H1 in a normal configuration 116
unplanned outage of local (H1) with production move to remote (H3) 115
Unplanned outage of local (H1), production move to intermediate (H2) 113
unplanned outage of remote (H3) 114
unplanned outage of Site 2 (H2) 114
scenarios for Metro Global Mirror 110, 113, 116
scenarios overview for practicing at H3 for Metro Global Mirror 119
practice scenario 1 119
practice scenario 2 119
practice scenario 3 120
scheduled actions
create script window 542, 556
creating 542, 556
edit script window 544
editing 544
overview 542
scripts
resolving script names 560
running 558
running scripts with alerts 555
using with jobs 555
SCSI Volumes (LUNs) report
vSphere Web Client extension 1655
Secure Socket Layer (SSL)
disabling 163
enabling 161
security 782
Tivoli Common Reporting, configuring 812
segments with wasted space report 1529
sending comments xviii
server alerts
triggering conditions 1665
Server Disks Capacity report 966
servers
adding
importing servers from a list 266
methods 265
adding without agents 264
modifying properties 437
removing 223
tiering, analyze 433
viewing data paths 402
viewing information in the web-based GUI 363

servers *(continued)*
viewing status in the web-based GUI 374
Servers Capacity report 959
service xviii
Data server
starting 144
Data Server 140
service class 65, 66, 408, 409, 411, 413
deleting 414
satisfying custom requirements of a 410
viewing properties of 413
viewing storage constraints of 413
viewing users who can provision by using 413
viewing volumes or shares associated with 413
Service Location Protocol (SLP)
overview 1743
service management connect xviii
Service tool
for agents 200
for servers 199
services
Data server
broadcasting location 143
configuring audit logging 143
configuring tracing 142
configuring tracing services 142
displaying 141
displaying service log 142
shutting down 143
Device server 145
changing the password 146
displaying information 146
viewing log 146
starting 122
stopping 121
viewing and managing 139
session commands
Basic HyperSwap 707
FlashCopy 707
Global Mirror 712
Global Mirror with Practice 715
Metro Global Mirror 718
Metro Global Mirror with Practice 722
Metro Mirror 709
Metro Mirror with Practice 711
Snapshot 708
session state change trap descriptions 636
session types
Basic HyperSwap 680
FlashCopy 681
Global Mirror 689
Metro Global Mirror 694
Metro Mirror 683
Snapshot 682
sessions
about 665
FlashCopy 728
Global Mirror 733
images 749
Metro Global Mirror 735
Metro Mirror 731
sessions *(continued)*
properties
SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified sessions 776
TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 sessions 756
XIV system sessions 779
removing 741
Snapshot 730
states 751
status icons 749
viewing 748
viewing details 754
Global Mirror 755
viewing health summary 748
viewing properties 756
setting
remote standby management server 640
SNMP 642
standby management server 639
volume protection 655
setting units of measurement for reports 480
setting up invoices 622
shares 306
shortcuts in minimap 564
sight-impaired 1747
single sign-on
overview 89
user authentication method used with 89
site awareness 649, 726
site, modifying location of 741
SMC xviii
Snapshot
about 682
creating a session 730
session commands 708
snapshot selection panel 490
SNMP
adding managers 642
MIBs 1743
overview 1743
setting up 642
traps 1743
viewing alerts 748
SNMP trap descriptions
communication-failure 638
configuration change 637
management servers state change 638
session state change 636
suspending-event notification 638
software zoning 80
solid-state drives
icons, storage 569
support for 1663
sorting rows in the web-based GUI 444
space usage
storage virtualizers 859
stand-alone GUI
managing resources 464
stand-alone GUI *(continued)*
user interface elements 464
standby management servers
about 634
ports 50
setting 639
setting a remote server 640
starting
IBM Tivoli Storage Productivity Center for Replication 629
Tivoli Storage Productivity Center 41
Tivoli Storage Productivity Center for Replication 44
Tivoli Storage Productivity Center stand-alone GUI 41
Tivoli Storage Productivity Center web-based GUI 43
starting services 122
states, session 751
status messages, role pairs 754
status of resources in the web-based GUI
acknowledging the status of a resource 380
acknowledging the status of a top-level resource 380
acknowledging the status of an internal resource 382
how the status of resources is determined 375
viewing the overall status of resources 378
viewing the status of internal resources 379
viewing the status of top-level resources 379
status, checking in topology viewer 574
status, session icons 749
stopping
Tivoli Storage Productivity Center for Replication 630
stopping services 121
Storage Access Times report 1513
storage area networks (SANs), viewing topology 563
Storage Availability report 1521
storage capacity 1490
storage connections
viewing details 781
Storage Management Initiative - Specification (SMI-S) 1742
storage mapping
viewing in vSphere Web Client extension 1644
vSphere Web Client extension 1644
Storage mapping
vSphere Web Client extension 1647
Storage Modification Times report 1516
storage network, viewing 563
Storage Networking Industry Association (SNIA) 1742
storage pools
tier level, setting 433
tiering, analyze 434
storage provisioning
vSphere Web Client extension 75
storage reporting
vSphere Web Client extension 75

Storage Resource agent 1659
changing credentials 219
changing the Windows service logon 223
Storage Resource agents
adding 210, 267
deployment considerations 83, 224
enabling 218
enabling or disabling scripts 221
enabling or disabling the monitoring of fabrics 222
registering with a different server 223
scheduling an upgrade 216
starting an upgrade 216
testing the connection with Tivoli Storage Productivity Center 219
uninstalling using GUI 223
upgrading 216
upgrading manually 217
viewing information 214
viewing log files 215
Storage Resource Group - VDisk Workload Activity Details report 1059
Storage Resource Group - VDisk Workload Activity report 1075
storage resource group sessions
deleting 555
groups
storage resource groups 555
viewing 555
working with 555
storage resource groups
groups
storage resource groups 552
overview 552
volumes space usage 864
Storage Resource Relationships Summary (Configurable) report 932
Storage Resource Relationships Summary report 930
storage subsystem
creating a group 599
group 599
performance reports 1548
storage subsystem constraint violation reporting
generate report 1501
storage subsystem reporting 1490
computer views 1490
generate report 1495
storage subsystem views 1490
Volume to HBA assignment 1490
storage subsystems
managing 599
overview 232
viewing information 599
storage system
removing 241
testing connection 241
storage system alerts
triggering conditions 1671
storage system metrics
customizing report displays in vSphere Web Client extension 1653
hiding columns 1653
sorting columns 1653
storage system metrics *(continued)*
viewing in vSphere Web Client extension 1647
vSphere Web Client extension 1648
Storage system metricsvSphere Web Client extension vSphere Web Client extension
Exporting information to a CSV file 1647
storage systems
about 643
adding 232, 272
DS4000 273
DS5000 273
DS6000 275
DS8000 277
ESS 279
IBM SONAS 288
other 291
SAN Volume Controller 281
Storwize V3500 283
Storwize V3700 283
Storwize V7000 283
Storwize V7000 Unified 286
XIV 278
adding a connection to
DS6000 651, 653
DS8000 650, 651
SAN Volume Controller 652
Storwize 652
TotalStorage Enterprise Server 651, 653
z/OS 653
icons 643
managing 643
modifying connection properties 654
modifying location of 654
modifying properties 437
ports 50
refreshing the configuration 655
removing 653
supported 643
updating credentials 234
viewing data paths 402
viewing details 781
viewing health summary 748
viewing information in the web-based GUI 363
viewing status in the web-based GUI 374
viewing volumes 781
Storage Systems Capacity report 969
Storage Systems Historical Capacity report 866, 973
storage tier reports
threshold values 1094
threshold values, change 1094
storage tiering
advanced 429
analysis 428
automatic 429
benefits 429
policies 428
storage tiers
workload activity reports 1058
Storage Usage By Database User
sample invoice 626
Storage Virtualizer - IO Group Details Report
performance statistics 846
Storage Virtualizer - MDisk Details Report
performance statistics 848
run 847
Storage Virtualizer - Node Details Report
performance statistics 850
run 850
Storage Virtualizer - Performance Details Report
performance statistics 844
Storage Virtualizer - Pool Details Report
performance metrics 852
run 852
Storage Virtualizer - Pools Historical Space Report
space usage 861
Storage Virtualizer - Pools Space Report
capacity 826
capacity and space report 827
run 826
used and available space 826
Storage Virtualizer - Port Details Report
performance statistics 857
run 856
Storage Virtualizer - Volume Details Report
run 854
Storage Virtualizer - Volumes Historical Space Report
space usage 864
Storage Virtualizer - Volumes Space Report
capacity 830
capacity and space report 831
run 830
used and available space 830
Storage Virtualizer – IO Group Details Report
run 845
Storage Virtualizer – Volume Details Report
performance statistics 855
predefined reports 855
Storage Virtualizer Performance Details Report
run 843
Storage Virtualizer - VDisk Workload Activity report 1071
storage virtualizers 433, 873
custom reports
I/O groups 908, 909
managed disks 887, 889
nodes 900
pools 922, 927
ports 918
storage virtualizers 873, 877
volumes 890, 894
disks 821
most active reports 834
performance details 843
pools 821
pools, balance 96
space usage 859

storage virtualizers *(continued)*
Storage Virtualizer - IO Group Details Report 845
Storage Virtualizer - Node Details Report
nodes 850
Storage Virtualizer - Performance Details Report 844
Storage Virtualizer - Pool Details Report 852
Storage Virtualizer Performance Details Report 843
Storage Virtualizers - Most Active IO Groups Report 834
Storage Virtualizers - Most Active MDisks Report 835
Storage Virtualizers - Most Active Nodes Report 837
Storage Virtualizers - Most Active Pools Report 838
Storage Virtualizers - Most Active Ports Report 842
Storage Virtualizers - Most Active Volumes Report 840
Storage Virtualizers - Space Report 820
tiering, analyze 434
virtual volumes 821
volumes, analyze placement 93
volumes, optimize 93
volumes, tier 96
Storage Virtualizers - Historical Space Report
space usage 859
Storage Virtualizers - Most Active IO Groups Report
most active 834
running 834
Storage Virtualizers - Most Active MDisks Report
most active 836
running 835
Storage Virtualizers - Most Active Nodes Report
performance statistics 837
running 837
Storage Virtualizers - Most Active Pools Report
performance statistics 839
running 838
Storage Virtualizers - Most Active Ports Report
performance statistics 842
running 842
Storage Virtualizers - Most Active Volumes Report 841
running 840
Storage Virtualizers - Space Report
capacity 820
capacity and space 821
run 820
used and available space 820
storage virtualizers reports
view 820
storage-subsystem controllers
performance reports 1541
storage-subsystem performance reports 1497
Summarized Performance of Volumes by Host Connection report 1004
Summarized Performance of Volumes by Hypervisor report 952
Summarized Performance of Volumes by Server report 962
support xviii
Support
contact xix
suspending-event notification trap descriptions 638
switch alerts
creating 345
triggering conditions 1676
Switch Performance Management 1662
Switch Performance report 1506, 1556
switch performance report, define and generate 1503
Switch Port Errors report 1557
switches
adding 241, 293, 294
agent types for monitoring 294
modifying properties 437
removing 244
testing connection 244
updating credentials 242
viewing information 242, 243
viewing information in the web-based GUI 363
viewing status in the web-based GUI 374
symbols
data copying 749
session status 749
volume roles 749
SYSMAINT authority
checking and setting 254
system maintenance authority
checking and setting 254
system reports 1511, 1525
Access File Summary 1513
Access Time Summary 1511
Computer Disk Space 1521
database storage by computer 1525, 1527, 1531, 1533
Disk Defects 1514
Disk Space Summary 1512
Fabric Manager 1554
instance storage by computer 1536
Most at Risk Files 1518
Most Obsolete Files 1518
node cache performance 1544
Oldest Orphaned Files 1519
Port Connections 1554
SAN Assets (ALL) 1555
SAN Assets (Connected Devices) 1555
SAN Assets (Switches) 1555
segments with wasted space 1529
Storage Access Times 1513
Storage Availability 1521
Storage Modification Times 1516
Switch Performance 1506, 1556
Switch Port Errors 1557
system reports *(continued)*
Top Switch Ports Data Rate Performance 1559
Top Switch Ports Packet Rate Performance 1560
total database free space 1528
total database freespace 1531, 1534
total device free space 1534
total DMS container free space 1536
Total Free Space 1523
user database quota violations 1527
user database space usage 1526, 1528, 1532, 1535, 1536
User Quota Violations 1520
User Space Usage 1524
using 1510
Wasted Space 1525

T

T11 FC-MI 1744
table alerts 353
description 353
triggering conditions 1670
tablespace alerts 352
description 352
tabs
alert 567
taking over management from the active server 641
tasks
analysis execution 454, 455
configuring the display of tasks in the web-based GUI 446
creating provisioning tasks in the web-based GUI 450
creating tiering storage, balancing pools, and transforming-storage tasks in the web-based GUI 452
deleting in the web-based GUI 457
exporting provisioning tasks in the web-based GUI 451
implementing optimization recommendations in theweb-based GUI 454
managing in the web-based GUI 445
managing provisioning tasks in the web-based GUI 449
managing tiering storage, balancing pools, and transforming-storage tasks in the web-based GUI 452
pausing, resuming, and canceling analysis-execution tasks in the web-based GUI 456
running provisioning tasks in the web-based GUI 450
running tiering analysis, transform-plan, and balance-analysis tasks in the web-based GUI 453
scheduling the implementation of optimization recommendations in theweb-based GUI 455
setting the number of task runs displayed 446
setting the retention of removed resources 446

tasks *(continued)*
viewing all tasks in the web-based GUI 447
viewing details in the web-based GUI 448
viewing in the web-based GUI 446
viewing task logs in the web-based GUI 448
viewing tasks for a resource in the web-based GUI 447
viewing tasks for a resource type in the web-based GUI 447
temporary zone sets 80
testing connection
fabrics 244
storage system 241
switches 244
threshold
alert 100
array 1692
array threshold 1692
controller 1696
controller thresholds 1696
threshold values
storage tier reports, change 1094
thresholds 1690
defining 397
port 1697
port thresholds 1697
viewing 400
viewing alerts 401
tier levels
storage pools, setting 433
tiering
advanced 429
advanced tiering 428
analysis period 435
automatic 429
automatic tiering 428
automatic tiering policy 96
performance data, collect 428
resources, select 428
target pools, select 428
wizard 428
tiering policies
advanced 428
automatic 428
tiering policy 429
timeout values
setting for Device server 176
Tivoli Storage Manager archive/backup actions, see archive/backup actions 539
Tivoli Storage Productivity Center
VASA provider 74
Tivoli Storage Productivity Center alerts
sharing with VMware 1641
Tivoli Storage Productivity Center for Replication
stopping 630
Tivoli Storage Productivity Center servers
overview 246
tivoliSRM.mib file 1743
tool bar 483
tools
repository copy 194
top n active volumes cache hit performance
performance reports 1549
Top Switch Ports Data Rate Performance report 1559
Top Switch Ports Packet Rate Performance report 1560
top volumes data rate
performance reports 1550
top volumes disk performance
performance reports 1550
top volumes I/O rate
performance reports 1552
top volumes response performance
performance reports 1553
topology settings, shortcut to 564
topology view for data paths 403
topology viewer
arranging entities within a group 574
connections, displaying 594
Data Path Explorer 584
entities, removing from database 597
entities, selecting 598
entity information, displaying 594
filtering resources 595
groups, creating custom 583
groups, default 565
icons, storage 569
labels, specifying 595
levels of detail 563
locating entities 595
manage storage subsystems 596
minimap 564
missing entities 597
monitoring and troubleshooting a data path 584
pinning entities to monitor 596
reports 594
show alert logs 598
show monitoring status 598
status, checking 574
table view 567
user-defined properties (UDPs) 584
viewing a single data path 584
zones and zone sets 593
total database free space report 1528
total database freespace report 1531, 1534
total device free space report 1534
total DMS container free space 1536
Total Free Space report 1523
Total Size By Database -Tablespace
sample invoice 627
TotalStorage Enterprise Storage Server
port thresholds 1697
TPC server probes
creating 332
TPCD.cfg
saveNonRoot parameter, add 306
trademarks 1761
translations
browser locale requirement xvii
traps for SNMP 1743
trending 1317
history charts 1317
triggering conditions
alerts 1665
triggering conditions *(continued)*
computer alerts 1665
database-table space alerts 1669
device 345
directory alerts 1668
endpoint device alerts 1677
fabric alerts 1675
file system alerts 1666
hypervisor alerts 1678
instance alerts 1669
NAS alerts 1678
performance 345
replication alerts 1680
server alerts 1665
storage system alerts 1671
switch alerts 1676
table alerts 1670
troubleshooting
performance 383
reports 1289

U

UDPs (user-defined properties) in topology viewer 584
units of measurement, setting for reports 480
unplanned outage of local (H1) with production move to remote (H3) 115
Unplanned outage of local (H1), production move to intermediate (H2) 113
unplanned outage of remote (H3) 114
unplanned outage of Site 2 (H2) 114
updating credentials
fabrics 242
hypervisors 231
storage systems 234
switches 242
upgrading
scheduling a Storage Resource agent upgrade 216
starting a Storage Resource agent upgrade 216
Storage Resource agents 216
Storage Resource agents manually 217
usage reporting
access load reports 1420
directories reports 1427
files reports 1422
OS user groups reports 1427
overview 1420, 1483
users reports 1427
usage violation reporting
constraint violations report 1428
overview 1427, 1428, 1488
quota violations reports 1428
Usage Violation Reporting
information in Usage Violation reports 1489
Quota Violations 1489
user administration 782
granting access privileges 785
modifying access privileges 786
removing user access 787
viewing access privileges 786

user authentication
Secure Socket Layer (SSL)
disabling 163
enabling 161
user authentication configuration
changing 146
user database quota violations report 1527
user database space usage report 1526, 1528, 1532, 1535, 1536
user interface
content pane 477
menu bar 477
navigation tree 476
navigation tree nodes 466
tool bar 483
user interfaces
changes in V5.2 38
when to use 34
User Quota Violations report 1520
User Space Usage report 1524
user-defined properties (UDPs) in topology viewer 584
users
about 783
assigning roles 78, 126
groups 783
roles 784
users reports 1427
using practice volumes 745

V

VASA provider
architecture 74
Tivoli Storage Productivity Center 74
VASA provider for Tivoli Storage Productivity Center
architecture 74
VDisk Details report 1064
ver command 631
version
check 138
version of IBM Tivoli Storage Productivity Center for Replication 631
View menu 478
viewing
access privileges 786
configuration changes in graphical view 493
configuration changes in table view 495
constraint violations 1502
copy sets 781
error messages 782
health summary 748
host system connection details 782
logical paths 663, 781
session details 754
Global Mirror 755
session properties 756
sessions 748
SNMP alerts 748
storage capacity 1490
storage connection details 781
storage resource group sessions 555
viewing *(continued)*
storage system details 781
storage-subsystem information 599
Tivoli Storage Productivity Center VASA provider reports 1642, 1657
VASA provider reports 1642, 1657
volume details 781
volumes affected by constraint violations 1502
Viewing
the overall status of a storage environment 365
viewing invoices 624
viewing your topology
arranging entities within a group 574
connections, displaying 594
entities, removing from database 597
entities, selecting 598
entity information, displaying 594
filtering resources 595
groups, creating custom 583
groups, default 565
icons, storage 569
labels, specifying 595
levels of detail 563
locating entities 595
minimap 564
pinning entities to monitor 596
reports 594
status, checking 574
table view 567
user-defined properties (UDPs) 584
zones and zone sets 593
views
overview 191
retrieving information 191, 193
virtual SANs (VSANs), viewing topology 563
VMware
capacity reports 1460
VMware vSphere
provisioning 106
volume protection
about 649
setting 655
volume roles
about 669
symbols 749
volumes
most active 841
optimize 429
practice 671, 745
role pairs 670
roles 669
Storage Virtualizer - Volume Details Report 854
Storage Virtualizer - Volumes Space Report 830, 831
Storage Virtualizers - Most Active Volumes Report 840
tiering 97
tiering, analyze 435
viewing details 781
workload 436
Volumes Capacity report 1047
Volumes Historical Capacity report 872, 1051
volumesstorage resource groups
storage virtualizers 864
VSANs 81
VSANs (virtual SANs), viewing topology 563
vSphere
reports 1643, 1644, 1647, 1648, 1653, 1654, 1655, 1656
vSphere reports
exporting to a CSV file 1647
vSphere resources
managing 458, 459, 461, 462
vSphere storage
provisioning 75
vSphere storage reports
viewing 75
vSphere Web Client extension 462
block storage reports 1655
datastores report 1656
Datastores report 1655
file system reports 1656
managing resources 458
NAS mounts report 1656
performance chart 1647, 1648
provisioning block storage 458, 459
provisioning file storage 461, 462
provisioning storage 75
SCSI Volumes (LUNs) report 1655
storage mapping to Tivoli Storage Productivity Center 1644
storage system metrics 1648
viewing block storage reports 1654
viewing fabric connections 1654
viewing file system reports 1656
viewing storage mapping 1644
viewing storage reports 75
viewing storage system metrics 1647
vSphere Web Client extension reports 1643

W

Wasted Space report 1525
Web Based Enterprise Management (WBEM) 1741
web-based GUI
acknowledging alerts 359
acknowledging the status of a resource 380
acknowledging the status of a top-level resource 380
acknowledging the status of an internal resource 382
alert severities 343
alerts
acknowledging 359
removing 362
configuring the display of tasks 446
customizing columns in lists 445
customizing lists 443
data collection
creating performance monitors 300
modifying performance monitors 301
performance monitors 300

web-based GUI *(continued)*
data collection *(continued)*
starting performance monitors 303
stopping performance monitors 303
verifying that a performance monitor is running 301
verifying that a probe is running 297
display differences with the stand-alone GUI 373
exporting information about resources, jobs, and alerts 441
fabrics 79
filtering 443
how information is organized 363
how the status of resources is determined 375
internal and related resources, list 259
managing tasks 445
Monitoring the status of resources 374
Opening the management GUI for a resource 440
removing alerts 362
resources that can be monitored 259
saving information about resources, jobs, and alerts 441
sorting rows in lists 444
status of resources 374
switches 79
tasks
analysis execution 454, 455
creating provisioning tasks 450
creating tiering storage, balancing pools, and transforming-storage tasks 452
deleting 457
exporting provisioning tasks 451
implementing optimization recommendations 454
managing provisioning tasks 449
managing tiering storage, balancing pools, and transforming-storage tasks 452
pausing, resuming, and canceling analysis-execution tasks 456
running provisioning tasks 450
running tiering analysis, transform plan, and balance-analysis tasks 453
scheduling the implementation of optimization recommendations 455
viewing all performance monitors 304
viewing all tasks 447
viewing information about internal and related resources 372
viewing information about resources 363
viewing information about top-level resources 371
viewing performance monitor logs 305
web-based GUI *(continued)*
viewing performance monitors 304
viewing performance monitors for a resource 304
viewing task details 448
viewing task for a resource 447
viewing task for a resource type 447
viewing task logs 448
viewing tasks 446
viewing the overall status of resources 378
viewing the status of internal resources 379
viewing the status of top-level resources 379
zones 79
windows
create script 542, 556
edit script 544
instance alerts 351
Online Execution 623
table alerts 353
Windows 1733, 1735
host names, verify 188
IBM Data Studio, start 254
launch-in-context 616
Windows menu 480
Windows service logon
changing for a Storage Resource agent 223
wizards
Analyze Tiering 96
working with
storage resource group sessions 555
workload
volumes 436
WWN zoning 80

X

XIV Storage System
port thresholds 1697

Z

zone 80
concepts 80
default 80
hardware 80
membership 80
non standard 82
port 80
software 80
supported switches 82
zone set 80
temporary 80
zones 79
zones and zone sets in topology viewer 593

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

© Copyright IBM Corp. 2001, 2013

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
U.S.A

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE: This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at http://www.ibm.com/software/info/product-privacy.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml.

Adobe and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Product Number: 5725-F92, 5725-F93, 5725-G33

Printed in USA

SC27-4060-01