IBM Tivoli Storage Productivity Center
IBM Tivoli Storage Productivity Center for Replication
Version 4.2.2.1

Installation and Configuration Guide



SC27-2337-06

IBM Tivoli Storage Productivity Center
IBM Tivoli Storage Productivity Center for Replication
Version 4.2.2.1

Installation and Configuration Guide



SC27-2337-06

Note:
Before using this information and the product it supports, read the information in "Notices" on page 443.

This edition applies to version 4, release 2, modification 2, fix pack 1 of IBM Tivoli Storage Productivity Center (product numbers 5608-WB1, 5608-WC0, 5608-WC3, 5608-WC4, 5725-A01, 5725-E18, 5608-WB2, and 5608-WB3) and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces SC27-2337-05.

© Copyright IBM Corporation 2001, 2012.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures vii

Tables ix

Preface xi

Who should read this guide xi
Publications xi
Accessing publications online xi
IBM Redbooks xiii
Translation xiii
Ordering publications. xiii
Providing feedback about publications xiv
Contacting the IBM Support Center xiv
Reporting a problem xiv
Conventions used in this guide. xv

New for Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2.1 . xvii

New for Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2 . . xix

New for Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.1 xxiii

Chapter 1. Planning for the Tivoli Storage Productivity Center family . . . 1

Planning for installation 2
Hardware requirements. 2
Software requirements 5
Running IBM Tivoli Storage Productivity Center in a 64-bit environment 20
Planning for IBM Tivoli Storage Productivity Center for Replication 21
Planning for Open HyperSwap replication . . . 22
General planning 23
Licenses for Tivoli Storage Productivity Center 23
General configuration guidelines 31
Planning for IBM Tivoli Storage Productivity Center authorization 33
TCP/IP ports used by the IBM Tivoli Storage Productivity Center family 33
User IDs and passwords for installing and using Tivoli Storage Productivity Center 37
Supported storage systems, devices, file systems, databases, and switches 42
Planning for multipath subsystem device drivers 47
Planning for Storage Resource agents. 49
Storage Resource agent protocol support . . . 50
Planning for Internet Protocol Version 6 51
Planning for native interface support 57
Planning for LDAP support 58
Planning for the single sign-on feature 59
Planning for enterprise-rollup reports. 61
Planning for the SAN configuration tools . . . 63
SAN Planner overview 65
Storage Optimizer overview 71
Planning for the Data Path Explorer 72
Planning for the Data Manager 74
Planning for the **Storage Resource agents** on Oracle Solaris. 75
Planning for the Disk Manager 75
Planning for the IBM CIM agents 77
Planning for the TagmaStore CIM agents . . . 78
Planning for Disk Manager configuration . . . 79
Planning for TotalStorage Enterprise Storage Server FlashCopy reports 82
Storage capacity of volumes 82
Planning for Tivoli Storage Productivity Center for Disk Select 82
Planning for the DS8000 86
Planning for DS8000 performance 87
Planning for the SAN Volume Controller 87
Planning for Storwize V7000. 90
Planning for Storwize V7000 Unified 92
Planning for the XIV system. 96
Planning for the Fabric Manager 98
Planning for Brocade or McData fabric management agents 99
Planning for QLogic or Cisco fabric management agents 103
Collecting data with the Fabric Manager . . . 105
Storage Resource agent or in-band Fabric agent 109
Out-of-band Fabric agent 110
Configuration guidelines for agent placement 111
Planning for Fabric Manager configuration . . 111
Planning for Fabric Manager switches and directors 114
Planning for host bus adapters (HBAs) 115
Planning for private switch networks 116
Planning for performance management. 116
Performance metrics 116
Performance management configuration . . . 119
Tivoli Provisioning Manager 5.1 for storage workflows 121
Microsoft Cluster Server. 121
Microsoft Cluster Server environment 121
Microsoft Cluster Server support 122
Planning for PowerHA 124
PowerHA environment 126
NAS and NetWare support 127
Network Attached Storage system requirements 127
Planning for NetApp device support 129

© Copyright IBM Corp. 2001, 2012

Planning for System Storage N Series Gateway servers 131
Planning for VMware 132
VMware configuration 135
VMware capacity reports 135
Planning for the Virtual I/O Server 136
Planning for SQL access to IBM Tivoli Storage Productivity Center's views. 139
Setting up a view-only user on Windows . . . 140
Monitoring your DB2 environment 141
Planning for the Discovery Library Adapter . . . 141
Tivoli Storage Productivity Center Monitoring Agent 141
Tivoli Agent Management Services 143
Tivoli Storage Productivity Center installation . . 144
Upgrading and migrating the Tivoli Storage Productivity Center family 147
Upgrading IBM Tivoli Storage Productivity Center from version 3 148
Upgrading Tivoli Storage Productivity Center scenarios 148

Chapter 2. Installing the Tivoli Storage Productivity Center family 153

Installing Tivoli Storage Productivity Center overview 153
Installing DB2 155
Preparing to install DB2 155
Installing DB2 on Windows 157
Installing DB2 on UNIX or Linux - GUI install 159
Installing DB2 on AIX using the command line 162
Installing DB2 on AIX 6.1 POWER7 164
Verifying that DB2 is installed correctly. . . . 166
Installing the IBM Tivoli Storage Productivity Center family 167
Preparing for installation 167
Starting the installation program 169
Using typical installation 172
Installing the Two Site or Three Site Business Continuity license 184
Using custom installation 185
Installing the agents 228
Installing the Language Pack 244
Installing and configuring Tivoli Common Reporting 248

Chapter 3. Configuring IBM Tivoli Storage Productivity Center 255

Configuration overview 255
Configuration Utility 261
Configuration utility tasks 261
Configure Devices 266
Configure storage subsystems 267
Configure fabrics and switches 276
Configure computers 278
Configure tape libraries 281
Managing schedules, jobs, and runs 283
Job Management - Recommendations 289
Configuration nodes 291
Role-to-Group Mappings 291
License Keys 293
Alert disposition 299
Log File and Cached Batch Report Retention 300
Quota and Constraint e-mail Address Rules . . 301
Scan/Probe Agent Administration 302
Manual NAS/NetWare Server Entry. 302
Agent Manager Registration 307
History Aggregator 308
Storage Resource agent deployments 308
Storage Resource Agent Upgrades 326
Configuration - Migrate Data/Fabric Agent . . 327
Update Storage Subsystem Credentials 328
NetWare Tree Logins 332
Resource History Retention. 332
Removed Resource Retention 333
Resource History Retention for Databases . . . 334
Removed Resource Retention for Databases . . 334
Configuration History Settings 335
Configuration: Device server specific 335
Configuring switches 335
Configuring the inactive zone sets for a fabric data source 337
Excluding devices from Fabric agent scans . . 338
Managing a SAN without agents 338
Setting timeout values for the Device server . . 338
Service Location Protocol (SLP) configuration considerations 340
Configuring IBM Tivoli Storage Productivity Center for a Windows domain. 342
Configuring the Data server for a Windows domain 342
Granting local administrative privileges to a domain account 343
Adding Tivoli Storage Productivity Center to the Windows domain. 344
Configuring IP addressing 345
Configuring Tivoli Storage Productivity Center with multiple IP addresses 345
Changing the HOSTS file 346
Starting the Tivoli Storage Productivity Center GUI using a web browser 347
Configuring tracing 351
Configuring the DB2 JDBC driver 352
Configuration guidelines for 500 or more agents 353
Using a master image to set up Storage Resource agent servers 353
Configuring LUN provisioning for Oracle Solaris 354
Assigning TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUNs to Oracle Solaris HBAs 355
Modifying the HBA configuration file 355
Setting Persistent Name Binding for QLA2200 and QLA2300 HBAs by using the **scfx** command 355
Modifying the SCSI disk configuration file . . 356
Checking for TotalStorage Enterprise Storage Server, DS6000, or DS8000 multipaths in VxDMP 357
How to import authentication information for a Storage Resource agent 358

Configuring Tivoli Storage Productivity Center for DS8000 LDAP authentication 359
Configuring multiple Tivoli Storage Productivity Center servers with one DS8000 R4.2 363
Setting up dual Tivoli Storage Productivity Center servers for high availability. 364
Configuring Tivoli Storage Productivity Center and SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 for single sign-on 366
Disabling unlicensed or unused Tivoli Storage Productivity Center components 368
Configuring and controlling the Tivoli Storage Productivity Center Monitoring Agent 370
Configuring the Tivoli Storage Productivity Center Monitoring Agent on Windows 370
Configuring the Tivoli Storage Productivity Center Monitoring Agent on AIX or Linux . . 371
Starting and stopping the Tivoli Storage Productivity Center Monitoring Agent 371
Installing and configuring the Tivoli Storage Productivity Center server with multiple NIC cards 372
Creating an SSH certificate for the root user ID . . 376
Configuring DB2, AIX, and Linux for IPv6-only environment. 377

Chapter 4. Upgrading and migrating the Tivoli Storage Productivity Center family. 379
General information for upgrading and migrating Tivoli Storage Productivity Center components . . 379
Upgrading Tivoli Storage Productivity Center for Replication 380
Fix packs and patches 380
Performance database migration tool (Tivoli Storage Productivity Center only) 380
Storage Subsystem Credential Migration tool (Tivoli Storage Productivity Center only) . . . 381
Upgrading TotalStorage Productivity Center 3.3.2 or later to Tivoli Storage Productivity Center 4.2.2.1 or later 382
Upgrading Tivoli Storage Productivity Center for Replication from V3.4 or later to V4.2.2.1 or later . 383
Upgrading DB2 from DB2 9.1 or DB2 9.5 to DB2 9.7 383
Upgrading DB2 8.2 to DB2 9.7. 386
Migrating the database repository for Windows 387
Migrating the database repository for AIX. . . 391
Migrating the database repository for Linux . . 394
Preparing for an upgrade 397
Backing up IBM Tivoli Storage Productivity Center. 397
Upgrading Tivoli Storage Productivity Center . . 398
Upgrading the remote database 405
Upgrading Storage Resource agents 406
Upgrading Storage Resource agents using a command 406
Migrating Data agents and Fabric agents to Storage Resource agents 407
Migrating the Data agent or Fabric agent using the user interface 407
Migrating Data agents and Fabric agents using a command 408
Migrating Data agents and Fabric agents using the installation wizard 409
Upgrading CIM agents for storage systems . . . 409
Upgrading CIM agents 410
Upgrading System Storage N Series Gateway servers 410

Chapter 5. Uninstalling the Tivoli Storage Productivity Center family . . 417
Uninstalling Tivoli Storage Productivity Center in a single-server environment by using the wizard . . 418
Uninstalling IBM Tivoli Storage Productivity Center for Replication with InstallShield 419
Uninstalling IBM Tivoli Storage Productivity Center for Replication manually 419
Uninstalling IBM Tivoli Storage Productivity Center for Replication manually on Windows . 419
Uninstalling IBM Tivoli Storage Productivity Center for Replication manually on AIX and Linux 420
Uninstalling the agents 421
Deleting or uninstalling an agent using the GUI 421
Uninstalling the Storage Resource agent manually 421
Uninstalling the Language Pack 422
Uninstalling the Tivoli Storage Productivity Center Monitoring Agent 423
Uninstalling the IBM Tivoli Storage Productivity Center Monitoring Agent on Windows 423
Uninstalling the IBM Tivoli Storage Productivity Center Monitoring Agent on UNIX or Linux . . 424
Uninstalling DB2 424
Uninstalling DB2 on UNIX or Linux. 424

Chapter 6. Reference 427
agent.sh command 427
Estimating tablespace allocation 427
Valid characters for user IDs and passwords . . . 428
TCP/IP ports used by the IBM Tivoli Storage Productivity Center family 429

Appendix A. Work sheet for user IDs and passwords 435

Appendix B. Accessibility features for Tivoli Storage Productivity Center . . 437

Appendix C. Accessibility features for Tivoli Storage Productivity Center for Replication 439

Notices 443
Trademarks 444

Glossary 447

Index 453

Figures

1. Select the type of installation you want to run panel 174
2. User ID and password panel 175
3. Tivoli Integrated Portal panel 177
4. Authentication selection panel 178
5. LDAP server information panel 179
6. LDAP user and group information panel 180
7. LDAP TPC Administrator user and group information panel 181
8. System prerequisites check panel 182
9. Directory Name panel 183
10. TPC-R Administrator ID and Password panel 183
11. Default ports panel 184
12. Select the type of installation you want to run panel 186
13. Select one or more components to install panel 187
14. Storage Resource Agent Information panel 188
15. Storage Resource Agent service information panel 189
16. Database administrator Information panel 191
17. New database schema information panel 191
18. Database schema creation information panel 192
19. Data server, Device server, and Storage Resource Agent information panel 193
20. Advanced security roles mapping panel 195
21. NAS discovery information panel. 196
22. Storage Resource Agent options panel 196
23. Tivoli Integrated Portal installation panel 197
24. Authentication selection panel 198
25. LDAP Server information panel 199
26. LDAP user and group information panel 200
27. LDAP TPC Administrator user and group information panel 201
28. New database schema information panel 204
29. Database schema creation information panel 205
30. New database schema information panel 206
31. Authentication selection panel 209
32. LDAP Server information panel 210
33. LDAP user and group information panel 211
34. LDAP TPC Administrator user and group information panel 212
35. Select the type of installation you want to run panel 215
36. Select one or more components to install panel 216
37. Database administrator Information panel 217
38. New database schema information panel 218
39. Database schema creation information panel 219
40. Summary information panel 220
41. Existing database schema information panel 221
42. Update Storage Subsystem Credentials panel 329
43. User ID and password panel 373
44. Data server, Device server, and Storage Resource Agent information panel 374
45. Delete Manual NAS/Netware Server Entry prompt. 411
46. Save Confirmation prompt 411
47. Add NAS Server panel 412
48. Verify that the N Series Gateway server is added 412
49. File system Agent Editor panel. 413
50. Scan/Probe Agent Administration panel 413
51. View information for N Series Gateway server 414
52. View information for a file system 414

© Copyright IBM Corp. 2001, 2012

Tables

1. Locations of publications for Tivoli Storage Productivity Center and related products . . xii
2. Hardware requirements for the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication servers . . 2
3. Storage system requirements for Tivoli Storage Productivity Center for Replication 3
4. Hardware requirements for the Tivoli Storage Productivity Center agents 4
5. Hardware requirements for the Tivoli Storage Productivity Center GUI and command-line interface 5
6. Operating system support for Tivoli Storage Productivity Center server, Tivoli Storage Productivity Center for Replication server, and GUI 6
7. Operating system support for the Tivoli Storage Productivity Center agents. 10
8. Tivoli Storage Productivity Center for Replication licenses 22
9. Available functions in each Tivoli Storage Productivity Center license 25
10. Available reports in each IBM Tivoli Storage Productivity Center license 28
11. License upgrade scenarios for IBM Tivoli Storage Productivity Center 30
12. TCP/IP ports used by IBM Tivoli Storage Productivity Center 34
13. Other TCP/IP ports used by IBM Tivoli Storage Productivity Center 35
14. TCP/IP ports used by IBM Tivoli Storage Productivity Center for Replication for incoming communication 36
15. TCP/IP ports used by IBM Tivoli Storage Productivity Center for Replication for outgoing communication 37
16. Installing the Device server, Data server, or Tivoli Storage Productivity Center server. . . 38
17. Installing the Storage Resource agent, Data agent, or Fabric agent 39
18. Agent service logon for Windows 39
19. DB2 administration server 40
20. User IDs and passwords for Tivoli Storage Productivity Center installation 40
21. Multipath subsystem device drivers supported by Tivoli Storage Productivity Center 48
22. IPv6 communication for internal Tivoli Storage Productivity Center components 52
23. Supported and preferred interfaces for fabric management 100
24. Supported and preferred interfaces for fabric management 104
25. Information collected by agent 106
26. Supported platforms and storage systems for VMware 133
27. Upgrade scenarios for Tivoli Storage Productivity Center V4.1.1 (or earlier) . . . 149
28. Upgrade scenario for Tivoli Storage Productivity Center V4.2 or later 149
29. Upgrade scenario for Tivoli Storage Productivity Center agents using V4.2 GUI . 149
30. Upgrade scenario for Tivoli Storage Productivity Center agents using Tivoli Storage Productivity Center V4.2 command line 150
31. Platform support for Tivoli Storage Productivity Center agents 150
32. Upgrade scenario for Tivoli Storage Productivity Center V4.1.1. agents with V4.1.1 installation program 151
33. DB2 Setup wizard default user and group names 157
34. Operating system for Storage Resource agents 169
35. Parameters required for each protocol 230
36. Storage Resource agent installation images 230
37. Language Pack files 246
38. Storage Resource agent files copied to the Data server after the Language Pack installation 247
39. Run-time jar files and installed locations 247
40. Location of Windows, UNIX, or Linux agent after an upgrade 248
41. Roles and authorization levels 291
42. Database permissions 293
43. Location of MIB and Baroc files 299
44. Operating system for Storage Resource agents 399
45. Tablespace allocation for the Tivoli Storage Productivity Center database schema . . . 428
46. Valid characters for user IDs and passwords 429
47. TCP/IP ports used by IBM Tivoli Storage Productivity Center 430
48. Other TCP/IP ports used by IBM Tivoli Storage Productivity Center. 431
49. TCP/IP ports used by IBM Tivoli Storage Productivity Center for Replication for incoming communication 432
50. TCP/IP ports used by IBM Tivoli Storage Productivity Center for Replication for outgoing communication. 433
51. User IDs and passwords for Tivoli Storage Productivity Center installation 435

© Copyright IBM Corp. 2001, 2012

Preface

IBM® Tivoli® Storage Productivity Center is a storage infrastructure management software product that can centralize, automate, and simplify the management of complex and heterogeneous storage environments.

Who should read this guide

This publication is intended for administrators or users who are installing and using IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication. It also describes the hardware and software requirements for installing the products and provides an overview of the installation procedures.

Administrators and installers should be familiar with the following topics:

- General procedures for installing software on Microsoft Windows, AIX®, Linux, HP-UX, and Oracle Solaris.
- SAN concepts
- IBM Tivoli Storage Productivity Center concepts
- IBM Tivoli Storage Productivity Center for Replication concepts
- DB2® Database for Linux, UNIX, and Windows
- Simple Network Management Protocol (SNMP) concepts
- IBM Tivoli Enterprise Console

This installation and configuration guide includes information for the IBM Tivoli Storage Productivity Center family of products which includes the following programs:

- IBM Tivoli Storage Productivity Center
- IBM Tivoli Storage Productivity Center for Replication

Publications

This section lists publications in the IBM Tivoli Storage Productivity Center library and other related publications. It also describes how to access publications online, how to order publications, and how to submit comments on 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 IBM Tivoli Storage Productivity Center and related products is available online. Use Table 1 on page xii to view and download these publications. Translated documents are available for some products.

© Copyright IBM Corp. 2001, 2012

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

Product	Online location (see Note)
IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication	http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp. In the left 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, in the left navigation pane click **Previous Versions** and click the version that you want to view (for example, **4.1.1**, **4.1**, or **3.4.1**).
IBM Tivoli Storage Productivity Center for Replication for System z®	http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp. In the left 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, in the left navigation pane click **Previous Versions** and click the version that you want to view (for example, **4.1.1**, **4.1**, or **3.4.1**).
IBM WebSphere® Application Server	http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp
IBM System Storage® Productivity Center	http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp. In the left navigation pane, click **System Storage Productivity Center**. To obtain PDF documents, click **System Storage Productivity Center** > **Printable documentation**. To view previous versions of the IBM System Storage Productivity Center Information Center, in the left navigation pane, click **Previous Versions** and click the version that you want to view (for example, **1.3.2**, **1.3.1**, or **1.3**).
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.

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

Product	Online location (see Note)
IBM System Storage SAN Volume Controller	http://publib.boulder.ibm.com/infocenter/svc/ic/index.jsp
IBM Storwize® V7000	http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp
IBM Storwize V7000 Unified	http://publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp
IBM XIV® Storage System	http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp
IBM DB2 Database for Linux, UNIX, and Windows	http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/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/.
Note: If you print PDF documents on other than letter-sized paper, in the **Print** window select **Fit to Printable Area** in the Page Scaling field. This option is available when you click **File** > **Print**. **Fit to Printable Area** ensures that the full dimensions of a letter-sized page print on the paper that you are using.	

IBM Redbooks

The IBM Redbooks® are books on 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/redbooks.

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.

Ordering publications

Information is provided for the ordering of IBM publications on the Internet or by telephone.

You can order many IBM publications online at http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss.

You can also order by telephone. In the United States and Canada, call 800-879-2755. In other countries, contact your IBM service representative.

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/v4r1/index.jsp and click the **Feedback** link at the bottom of the information center Welcome page or topic pages.

Contacting the IBM Support Center

You can find support information for IBM products from various sources.

For support for IBM Tivoli Storage Productivity Center, you can contact the IBM Support Center in one of the following ways:

- 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_Standard_Edition.

  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 for customer support telephone numbers.

You can also review the *IBM Software Support Handbook*, which is available on our website 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 the IBM Support Center, 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

This topic provides a list of what information you should have ready when you encounter a problem.

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.

Conventions used in this guide

This section provides information about the conventions used in this publication.

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

The following typeface conventions are used in this publication:

Bold

- Lower-case and mixed-case commands that appear with text
- Command options that appear with text
- Flags that appear with text
- Graphical user interface (GUI) elements (except for titles of windows and dialogs)
- Names of keys

Italic

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

monospace

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

For syntax notation, these conventions are used:

- # is the prompt for the root user on UNIX platforms.
- Uppercase and lowercase characters do matter for UNIX and Linux. Type in commands exactly as shown.

New for Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2.1

Use this information to learn about new features and enhancements in Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2.1. This topic highlights the changes since Tivoli Storage Productivity Center 4.2.2.

Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication 4.2.2.1 add the following new features, functions, and enhancements:

IBM Storwize V7000 Unified

Storwize V7000 Unified is a virtualized storage system that includes Storwize V7000 and Storwize V7000 File Module. Storwize V7000 Unified is designed to consolidate block and file workloads into a single storage system for simplicity of management, reduced cost, highly scalable capacity, performance, and high availability. Storwize V7000 Unified also offers improved efficiency and flexibility through built-in solid state drive (SSD) optimization, thin provisioning and non-disruptive migration of data from existing storage. The system can virtualize and reuse existing disk systems offering a greater potential return on investment.

Data for Storwize V7000 Unified storage systems is collected, monitored, displayed, and reported in Tivoli Storage Productivity Center and data replication is supported by Tivoli Storage Productivity Center for Replication.

System Storage DS8000 R6.2

Starting with System Storage DS8000 R6.2, you are no longer required to access the DS Storage Manager through the Tivoli Storage Productivity Center GUI. The DS Storage Manager can be accessed from any location that has network access using a web browser. For a list of browsers that are supported for DS8000, see the DS8000 information center.

© Copyright IBM Corp. 2001, 2012

New for Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2

Use this information to learn about new features and enhancements in Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2. This topic highlights the changes since Tivoli Storage Productivity Center 4.2.1.

Tivoli Storage Productivity Center

Tivoli Storage Productivity Center 4.2.2 adds the following new features, functions, and enhancements:

IBM Tivoli Storage Productivity Center Select

The Select Edition includes the same functions and features available in IBM Tivoli Storage Productivity Center Standard Edition:

- Disk, fabric, tape, and data management functions
- Storage system performance monitoring, fabric performance monitoring, chargeback, and database monitoring
- Analytical functions, such as:
  - configuration analysis
  - storage area network (SAN) Planner
  - Storage Optimizer
  - configuration history

This product supports the IBM, XIV Storage System, IBM Storwize V7000, IBM System Storage DS3000, IBM System Storage DS4000, and IBM System Storage DS5000 as stand-alone devices, or when attached to the IBM System Storage SAN Volume Controller. This product also supports any storage devices that are attached to Storwize V7000.

IBM Tivoli Storage Productivity Center for Disk Select

Tivoli Storage Productivity Center for Disk Midrange Edition has been renamed to Tivoli Storage Productivity Center for Disk Select. This product provides the same function of theTivoli Storage Productivity Center for Disk Midrange Edition:

- Basic disk, fabric, tape, and data management functions
- Storage system performance monitoring

This product supports the XIV system, Storwize V7000, DS3000, DS4000, and DS5000 as stand-alone devices, or when attached to the SAN Volume Controller. This product also supports any storage devices that are attached to Storwize V7000.

Storage tier reports

Storage tier reports help you decide where to place accessed data for the best performance from your system, based on whether the data is accessed frequently or infrequently. Infrequently accessed data can be placed on lower-cost storage systems and frequently accessed data can be placed on faster performing storage systems.

To create the reports, you must have Tivoli Storage Productivity Center Standard Edition or Tivoli Storage Productivity Center Select installed. You

© Copyright IBM Corp. 2001, 2012

must also have Tivoli Common Reporting installed. These reports are supported for IBM System Storage SAN Volume Controller and Storwize V7000. To generate the reports, you must set the back-end subsystem pool properties, and then run a performance monitor job for the SAN Volume Controller or Storwize V7000. You can then display the reports using Tivoli Common Reporting.

Tivoli Storage Productivity Center reports
For SAN Volume Controller clusters and Storwize V7000 systems, you can drill down from storage subsystems and expose managed disks, managed disk groups, and virtual disks. For Storwize V7000 systems, you can also drill down from the internal disks node to view the internal storage on the system or generate a Storage Subsystem Disks report to view the internal storage on the system.

SAN Planner
The SAN Planner now supports workload profiles and replication planning for XIV systems.

Tivoli Storage Productivity Center for Replication

Tivoli Storage Productivity Center for Replication 4.2.2 supports IBM XIV Storage System. You can use the following session types for the XIV system:

Snapshot
Snapshot is a new session type that creates a point-in-time copy (snapshot) 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.

Metro Mirror Failover/Failback
Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 kilometers apart. You can use failover and failback to switch the direction of the data flow.

Global Mirror Failover/Failback
Global Mirror is a method of asynchronous, remote data replication that operates between two sites that are over 300 kilometers apart. You can use failover and failback to switch the direction of the data flow.

Tivoli Storage Productivity Center for Replication support for the XIV system includes the following new features.

Support for volume nickname
For the XIV system sessions, you can provide the volume ID or the volume nickname as a parameter value when you add or remove copy sets by using the command line interface (CLI) commands **mkcpset** and **rmcpset**.

In addition, you can include the XIV system volume ID or the volume nickname in a comma-separated value (CSV) file that is used to import copy set information. You can import the CSV file by using the **importcsv** command or the Tivoli Storage Productivity Center for Replication graphical user interface (GUI). CSV files that are exported from Tivoli Storage Productivity Center for Replication for XIV system sessions include the volume nickname rather than the volume ID. CSV files are exported by using the **exportcsv** command.

New CLI commands
The following CLI commands are new. For more information about new

and updated CLI commands for Tivoli Storage Productivity Center for Replication 4.2.2, see the IBM Tivoli Storage Productivity Center Command-line Interface Reference.

cmdsnapgrp
Use the **cmdsnapgrp** command to run a specific action against a snapshot group that is in an XIV system Snapshot session. A snapshot group is a grouping of snapshots of individual volumes in a consistency group at a specific point in time.

lssnapgrp
Use the **lssnapgrp** command to view snapshot groups that are in an XIV system Snapshot session.

lssnapgrpactions
Use the **lssnapgrpactions** command to specify the session and snapshot group name that you want to view available actions for.

lssnapshots
Use the **lssnapshots** command to view snapshots that are in a snapshot group in an XIV system session.

New for Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.1

Use this information to learn about new features and enhancements in Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.1. This topic highlights the changes since Tivoli Storage Productivity Center 4.2.

Tivoli Storage Productivity Center

Tivoli Storage Productivity Center 4.2.1 adds the following new features, functions, and enhancements:

SAN Volume Controller Version 6.1

Tivoli Storage Productivity Center supports SAN Volume Controller Version 6.1.

Tivoli Storage Productivity Center also supports launch in context and single sign-on for SAN Volume Controller.

Note: Launch in context and single sign-on was not supported for SAN Volume Controller Version 5.1 or earlier.

IBM® Storwize® V7000

Storwize V7000 is a hardware and software solution that provides unmatched performance, availability, advanced functions, and highly scalable capacity. Storwize V7000 offers IBM storage virtualization, SSD optimization and "thin provisioning" technologies built in to improve storage utilization. The storage system can be reconfigured to meet changing needs quickly and easily. This solution helps to reduce costs without performance degradation.

Tivoli Storage Productivity Center supports Storwize V7000 as a native interface device.

The Storwize V7000 hardware consists of a set of drive enclosures. Control enclosures contain disk drives and two node canisters. The two nodes within the canisters make an I/O group that is attached to the SAN fabric. A single pair of nodes is responsible for serving I/O on a given volume. Because a volume is served by two nodes, there is no loss of availability if one node fails or is taken offline.

Storwize V7000 can be used as a traditional RAID storage system where the internal drives are configured into arrays, and volumes are created from those arrays. Storwize V7000 can also be used to virtualize other storage systems.

Storwize V7000 supports both regular and solid-state drives (SSDs). A Storwize V7000 system without any internal drives can be used as a storage virtualization solution.

Storwize V7000 includes IBM System Storage Easy Tier®, a function that supports solid-state drives (SSDs) in a storage pool that also contains hard

© Copyright IBM Corp. 2001, 2012

disk drives (HDDs). The system automatically and nondisruptively moves frequently accessed data from HDD MDisks to SSD MDisks, thus placing such data in a faster tier of storage.

Easy Tier eliminates manual intervention when assigning highly active data on volumes to faster responding storage. In this dynamically tiered environment, data movement is seamless to the host application regardless of the storage tier in which the data is located. Manual controls exist so that you can change the default behavior, for example, such as turning off Easy Tier on storage pools that have both types of MDisks.

Tivoli Storage Productivity Center provides basic support for this product, including discovery, probes, performance monitoring, Storage Optimizer, and SAN Planner. Tivoli Storage Productivity Center reports and topology views also provide information about the tier, tier capacity, and tier free capacity for Storwize V7000.

Like other systems, Storwize V7000 is displayed as an entity within data sources, reports, data collection schedules, the topology viewer, and so on.

Performance Manager enhancements

New performance metrics, counters, and thresholds for DS8000 and SAN Volume Controller

The Performance Manager now collects, reports, and alerts users on various performance metrics related to fiber channel link problems in the SAN environment.

Some of the thresholds you can define are:

- Error (illegal) frame rate for DS8000 ports
- Link failure rate for SAN Volume Controller and DS8000 ports
- Invalid CRC rate for SAN Volume Controller, DS8000, and switch ports
- Invalid transmission word rate for SAN Volume Controller, DS8000, and switch ports
- Zero buffer-buffer credit timer for SAN Volume Controller ports

XIV system enhancements

In Tivoli Storage Productivity Center 4.2, users can add an XIV storage system through a single IP address. However, XIV storage systems have the capability of defining up to three administrative nodes, each with their own IP address. In Tivoli Storage Productivity Center 4.2.1, if an XIV storage system is configured with multiple administrative nodes, Tivoli Storage Productivity Center detects the IP addresses for these nodes. If Tivoli Storage Productivity Center fails to connect to one of the IP addresses, then an attempt is made to connect to the XIV storage system using one of the other IP addresses.

LUN correlation Support for Virtual SCSI disks on the Virtual I/O Client (LPAR)

The Virtual SCSI environment consists of a Virtual SCSI Server Adapter or Adapters that are created on the Virtual I/O Server and mapped to a Virtual I/O Client. If you have a Virtual SCSI environment and the virtual target device has a one to one mapping to a storage volume, Tivoli Storage Productivity Center can extract the storage system serial number and correlate the Virtual SCSI disk on the LPAR to a corresponding storage volume.

This function means that the Storage Resource agent installed on the Virtual I/O Client (LPAR), can collect this information.

Tivoli Storage Productivity Center for Replication

Tivoli Storage Productivity Center for Replication 4.2.1 adds the following new features, functions, and enhancements:

SAN Volume Controller Version 6.1 and IBM Storwize V7000

Tivoli Storage Productivity Center for Replication supports SAN Volume Controller version 6.1 and IBM Storwize V7000 as described in "Tivoli Storage Productivity Center" on page xxiii.

Chapter 1. Planning for the Tivoli Storage Productivity Center family

The following sections provide information to help plan your Tivoli Storage Productivity Center environment before you install the product. In a complex environment, good planning helps you avoid delays and problems in getting your system up and running.

The Tivoli Storage Productivity Center family includes the following programs:

- Tivoli Storage Productivity Center
- Tivoli Storage Productivity Center for Replication

The Tivoli Storage Productivity Center documentation uses the term "storage subsystem" and the Tivoli Storage Productivity Center for Replication documentation uses the term "storage system". Both terms refer to the devices used for storage management.

For the most current information about operating systems supported, storage systems, switches and directors, and HBAs, see the Tivoli Storage Productivity Center support Web site at http://www-01.ibm.com/support/docview.wss?rs=40&uid=swg21386446&loc=en_US&cs=utf-8&lang=en.

For the Tivoli Storage Productivity Center for Replication user

The installation program installs both Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication.

If you do not want the options available with Tivoli Storage Productivity Center, the quickest way to install Tivoli Storage Productivity Center for Replication is to use typical installation.

Tivoli Storage Productivity Center for Replication uses an internal database repository and does not use DB2. However, you must still install DB2 because Tivoli Storage Productivity Center requires the use of DB2 for the database repository.

The general steps to install Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication are:

1. Install DB2. See “Installing DB2” on page 155.
2. Install Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication. See “Using typical installation” on page 172. The minimum components to install would be:
   - Tivoli Integrated Portal
   - Data server
   - Device server
   - Tivoli Storage Productivity Center for Replication

© IBM Corporation 2009 © IBM 2001, 2012

Planning for installation

Before installing Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication for the first time, be familiar with your operating systems, storage devices, communication protocols, and system configurations.

The Tivoli Storage Productivity Center server must be a dedicated computer for Tivoli Storage Productivity Center operations and not shared with other applications. For more information about general configuration guidelines, see "General configuration guidelines" on page 31.

Hardware requirements

The Tivoli Storage Productivity Center server can require a large amount of memory, network bandwidth, and processor resources. In many cases, the server performs best when other applications are not installed on the same system.

See Table 2 for information about the minimum hardware requirements needed for the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication servers.

Table 2. Hardware requirements for the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication servers

Item	Hardware requirements
Processor	For Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication: • For Windows or Linux: 1 x Intel Quad-Core Xeon or greater • For AIX: IBM POWER4 or later processor, 1 GHz
Memory	8 GB of RAM. If you have at least 4 GB but less than 8 GB of RAM, you can still install Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication. However, you get a warning message. If you have less than 8 GB of RAM, run only one component: either Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication because of system load. For information about how to disable Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication after installation, see "Disabling unlicensed or unused Tivoli Storage Productivity Center components" on page 368. **Note:** • If you are running your system in an environment with less than 8 GB of RAM, do not install Tivoli Common Reporting on the Tivoli Storage Productivity Center server. You can install Tivoli Common Reporting on a separate server. • To run Tivoli Storage Productivity Center on AIX requires at least 6 GB of RAM and 12 GB of swap space. For information about paging swap space, see the man pages mkps or chps.

Table 2. Hardware requirements for the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication servers (continued)

Item	Hardware requirements
Disk space	• For installations on Windows operating systems, your system must have 6 GB of available disk space and 500 MB in the Windows temporary directory. • For installations on AIX operating systems, 10 GB of free disk space is required. • For installations on AIX or Linux operating systems, you must have the following free space: – 2.25 GB for the `/tmp` directory – 3 GB for the `/opt` directory – 250 MB in the `/home` directory – 10 KB in the `/etc` directory – 1.4 GB in the `/usr` directory (for Tivoli Integrated Portal) **Note:** After you install Tivoli Storage Productivity Center and start collecting data, a large amount of disk space is required. The amount of data collected depends on many factors: • How many devices you have • How long you keep the data • How frequently you collect data Some users have experienced disk space usage of about 40 - 80 GB.
Network interface card (NIC)	If you have multiple NIC cards on the Tivoli Storage Productivity Center server, see "Installing and configuring the Tivoli Storage Productivity Center server with multiple NIC cards" on page 372 for special considerations.
Console	A console or remote-connectivity application such as KDE, Remote Desktop, or VNC is required during installation of Tivoli Storage Productivity Center. The installation program is interactive (as opposed to silent).

See Table 3 for information about the storage system requirements for Tivoli Storage Productivity Center for Replication.

Table 3. Storage system requirements for Tivoli Storage Productivity Center for Replication

Storage system	Hardware requirements
IBM Storage System DS8000	• Release 2 minimum firmware level 6.2.410.30 or Release 3 minimum firmware level 63.1.32.0 • Advanced Copy services licenses • Optional Ethernet adapters for Tivoli Storage Productivity Center for Replication – Single image - feature code 1801 – Dual image - feature code 1802 and 1803
IBM System Storage DS6000	• Minimum firmware level 6.2.2.64 • Advanced Copy services licenses

Table 3. Storage system requirements for Tivoli Storage Productivity Center for Replication (continued)

Storage system	Hardware requirements
IBM Enterprise Storage Server® (ESS) 2105-800	• Minimum firmware level 2.4.4.72 • Feature codes 240 - PRC and 2240 - PRC
• SAN Volume Controller 4.2.1, 4.3.0, 5.1.0, and 6.1 • IBM Storwize V7000 • IBM Storwize V7000 Unified	Copy services licenses
Host adapter feature codes	One or more of the following host adapter feature codes: • 3021 - Fibre Channel/FICON Host Adapter • 3023 - Fibre Channel/FICON Host Adapter • 3024 - 2 Gb Fibre Channel/FICON Host Adapter • 3025 - 2 Gb Fibre Channel/FICON Host Adapter

See Table 4 for information about the minimum hardware requirements needed for the Tivoli Storage Productivity Center agents.

Table 4. Hardware requirements for the Tivoli Storage Productivity Center agents

Operating System	Item	Hardware requirements
• Windows and Linux on Intel • UNIX • AIX on IBM Power Systems™	Processor	Pentium 400® MHz processor, or above
	Disk space	Minimum memory and space requirements to run: • 256 MB of RAM. • For a locally installed Storage Resource agent, the temporary space required is 100 MB. For a remotely installed Storage Resource agent, the temporary space required is 250 MB.

Table 4. Hardware requirements for the Tivoli Storage Productivity Center agents (continued)

Operating System	Item	Hardware requirements
All	Network interface card (NIC)	Tivoli Storage Productivity Center supports multiple NIC cards. When you install an agent on a machine locally or remotely and the machine has more than one NIC card, the Tivoli Storage Productivity Center agent installer determines the NIC card to use. If there are no NIC cards that can be used for two-way communication, the installation program returns an error message. For more information about multiple NIC cards, see "Installing and configuring the Tivoli Storage Productivity Center server with multiple NIC cards" on page 372.

See Table 5 for information about the minimum hardware requirements needed for the Tivoli Storage Productivity Center GUI and command-line interface.

Table 5. Hardware requirements for the Tivoli Storage Productivity Center GUI and command-line interface

Operating system	Item	Hardware requirements
Windows	Processor	Minimum is PII 500 MHz processor, suggest PII 1 GHz processor or above
	Disk space	• 256 MB of RAM • Approximately 525 MB of hard disk space
	Console	1024 x 768 console or above
UNIX	Disk space	• 256 MB of RAM • Approximately 525 MB of hard disk space
	Console	1024 x 768 console or above

Software requirements

This section describes the software that is required to install and run your system. This section includes operating systems supported, browsers supported, databases supported, and other software required such as DB2.

Software requirements for operating systems

Use this information to learn what operating systems are supported by the IBM Tivoli Storage Productivity Center family.

Note: Tivoli Storage Productivity Center software is bundled with its own IBM Java™ JRE. You must use IBM Java with Tivoli Storage Productivity Center. When using the Java Web Start method of remotely accessing the Tivoli Storage Productivity Center GUI, you must also use IBM Java.

Before installing Tivoli Storage Productivity Center, check the Tivoli Storage Productivity Center support site for the latest platform support. Go to the Tivoli

Storage Productivity Center support Web site at http://www-01.ibm.com/support/docview.wss?rs=40&uid=swg21386446&loc=en_US&cs=utf-8&lang=en.

All the support in this table applies to a 32-bit operating system or setting up the operating system to run a 32-bit application using a 64-bit processor.

Table 6. Operating system support for Tivoli Storage Productivity Center server, Tivoli Storage Productivity Center for Replication server, and GUI

Operating system or platform		**Tivoli Storage Productivity Center server**	**Tivoli Storage Productivity Center for Replication server**	**Tivoli Storage Productivity Center GUI**
AIX, see "Notes® for AIX"	IBM AIX 5.3 (64-bit) • POWER4 • POWER5 • POWER6®	Yes	Yes	Yes
	IBM AIX 6.1 (64-bit) • POWER4 • POWER5 • POWER6 • POWER7®	Yes	Yes	Yes
	IBM AIX 7.1 (64-bit) • POWER5 • POWER6 • POWER7	Yes	Yes	Yes
Linux, see "Notes for Linux"	Red Hat Enterprise Linux Base Server and Advanced Platform 5 (x86-64)	Yes	Yes	Yes
	Red Hat Linux Base Server and Advanced Platform 6 (x86-64)	Yes	Yes	Yes

Table 6. Operating system support for Tivoli Storage Productivity Center server, Tivoli Storage Productivity Center for Replication server, and GUI (continued)

Operating system or platform		Tivoli Storage Productivity Center server	Tivoli Storage Productivity Center for Replication server	Tivoli Storage Productivity Center GUI
Windows, see "Notes for Windows"	Windows 2003 (x86, x86-64) • Standard Edition • Enterprise Edition	Yes	Yes	Yes
	Windows 2003 R2 (x86, x86-64) • Standard Edition • Enterprise Edition	Yes	Yes	Yes
	Windows 2008 (x86, x86-64) • Standard Edition • Enterprise Edition	Yes	Yes	Yes
	Windows 2008 R2 (x86, x86-64) • Standard Edition • Enterprise Edition	Yes	Yes	Yes
	Windows 2008 R2 Standard Edition • 64-bit	Yes	Yes	Yes
	Windows XP (x86)	No	No	Yes
	Windows Vista (x86)	No	No	Yes
	Windows Vista SP2 (x86)	No	No	Yes, see note 4 on page 8
	Windows 7 (x86) (see note 3 on page 8)	No	No	Yes
VM on VMware	VMware ESX and ESXi 3.0.x, 3.5.x, 4.0.x, and 4.1.x • Red Hat Enterprise Linux 5 • Red Hat Enterprise Linux 6 • Windows 2003 • Windows 2003 R2 • Windows 2008 • Windows 2008 R2	Yes	Yes	Yes
z/OS® V1.10, V1.11, V1.12 (31-bit)	-	No	Yes	No

Notes for AIX operating systems

Important:

1. The following AIX fixes must be applied.
   For AIX 5.3, the minimum maintenance levels required are:
   - Technology Level 5300-11 Service Pack 4

- Technology Level 5300-12 Service Pack 1

For AIX 6.1, the minimum maintenance level required are:

- Technology Level 6100-04 Service Pack 5
- Technology Level 6100-05 Service Pack 1
- Technology Level 6100-06

You can determine the Technology Level and Service Pack installed by running the following command:

```
oslevel -s
```

You can download the package from http://www.ibm.com/eserver/support/fixes/fixcentral/pfixpacks/53 or http://www.ibm.com/eserver/support/fixes/fixcentral/pfixpacks/61.

2. AIX 6.1.0.4 or later comes with IBM Systems Director. IBM Systems Director has a common agent service that can use port 9510. The Storage Resource agent also uses port 9510. If you do not need IBM Systems Director, you can uninstall IBM Systems Director before installing Tivoli Storage Productivity Center. If you install Tivoli Storage Productivity Center and require IBM Systems Director, specify a different port for the Storage Resource agent. For information about the ports IBM Systems Director uses, go to http://publib.boulder.ibm.com/infocenter/director/v6r2x/index.jsp?topic=/com.ibm.director.plan.helps.doc/fqm0_r_ports_for_the_management_server.html.
3. AIX 6.1 comes with Lightweight Infrastructure (LWI). If this package is enabled, the package can cause Tivoli Storage Productivity Center installation issues:
   - There can be a port conflict with Tivoli Storage Productivity Center at port 9510.
   - Can cause the Storage Resource agent installation to fail on hosts where only the agent is installed.
4. Use the GNU tar program version 1.14 or later rather than the native AIX tar program to extract files from Tivoli Storage Productivity Center installation images. The AIX tar program might truncate long file names, which can cause errors in Tivoli Storage Productivity Center and Tivoli Integrated Portal installation.

   The GNU tar file archiving program can be found from the following link. http://www-03.ibm.com/systems/power/software/aix/linux/toolbox/alpha.html.

Notes for Windows operating systems

Important:

1. Windows 2008 requires DB2 9.1 with fix pack 5 or later.
2. For information about IPv6 support in Windows Server 2003 Enterprise Server, see http://support.microsoft.com/kb/325449.
3. The Tivoli Storage Productivity Center GUI can be installed on Windows 7. However, the Tivoli Storage Productivity Center installation wizard does not support the installation of the GUI on Windows 7. Therefore, you must run the GUI using Java Web Start with Java 6. Java 6 provides additional features to work with the enhanced security of Windows 7. The CLI installation on Windows 7 is not supported.
4. Use the Windows setting for "Windows Classic" when installing the Tivoli Storage Productivity Center GUI on Windows Vista.

Notes for Linux operating systems

Important:

1. For Red Hat Linux Version 5, Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication require the following libraries:
   - libXp-1.0.0-8.i386.rpm or later
   - libXmu-1.0.2-5.i386.rpm or later
   - compat-libstdc++-33-3.2.3-61.i386.rpm or later
   - compat-libstdc++-296-2.96-138.i386.rpm or later
   - libXp-1.0.0-8.1.el5.i386.rpm or later
   - libXmu.i386
   - libXft

   Tivoli Storage Productivity Center for Replication also requires the following libraries on Red Hat Linux Version 5.4:
   - XFree86-libs-data-4.3.0-2.90.55.i386.rpm
   - XFree86-libs-4.3.0-2.90.55.i386.rpm

   For Red Hat Linux Version 6, Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication also require the following library:
   - libXtst-1.0.99.2-3.el6.i686.rpm or later
   - 32-bit version of pam_krb5 (required on 32-bit and 64-bit machine)

   To determine whether the i386 RPMs are installed and what version is installed, use the **--queryformat** option.

   Note: The **queryformat** option has two hyphen marks before the option.

   ```
   # rpm -q --queryformat='%{N}-%{V}-%{R}.%{arch}\n' libXp libXmu |
   grep -v x86_64
   libXp-1.0.0-8.1.el5.i386
   libXmu-1.0.2-5.i386
   #
   ```

   To install the required libraries, search the installation media for these RPMs:

   ```
   # find /mnt/ -name libXp-*.i386.rpm
   # find /mnt/ -name libXmu-*.i386.rpm
   ```

   Once you have located the directory containing the RPM files, install the RPM files (if they are not installed) or update the RPM files to the latest version (if they are already installed). Use the **-U** option:

   ```
   # rpm -U libXp-1.0.0-8.1.el5.i386.rpm libXmu-1.0.2-5.i386.rpm
   ```

   Note: The RPM versions on your installation media might differ from this example.

   To verify that you have the RPMs installed, use the **--queryformat** option to see the architecture of the installed RPMs.

   Note: The **queryformat** option has two hyphen marks before the option.

   ```
   # rpm -q --queryformat='%{N}-%{V}-%{R}.%{arch}\n' libXp libXmu |
   grep -v x86_64
   libXp-1.0.0-8.1.el5.i386
   libXmu-1.0.2-5.i386
   #
   ```

2. The Tivoli Storage Productivity Center installation program is unable to run in graphical mode on Red Hat Enterprise Linux AS Version 4 and 5 or SUSE Linux Enterprise Server version 9 and 10. You must install a Linux RPM

Package Manager (RPM) package on these operating systems. For more information see http://www-01.ibm.com/support/docview.wss?uid=swg21303823.

3. DB2 Version 9.7 does not support Red Hat Linux version 4 or SUSE Linux version 9. (Note that the Tivoli Storage Productivity Center server also is not supported on SUSE Linux.) You must upgrade your operating systems to a supported level before attempting to upgrade to DB2 Version 9.7.
4. Before installing Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication on Red Hat Linux version 5.2 or later, check to see if the shadow utility requirements are met. For information about the shadow utility, go to the information center. Search for **shadow password**.
5. Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

Table 7 shows the operating system support for the Tivoli Storage Productivity Center agents. Unless otherwise specified, all the support applies to a 32-bit operating system or setting up the operating system to run a 32-bit application using a 64-bit processor. For zSeries®, the support applies to 31-bit and 64-bit processor.

Table 7. Operating system support for the Tivoli Storage Productivity Center agents

Platform		**Storage Resource agent**	**File systems**	**Volume managers**
AIX, see "Notes for AIX"	AIX 5.3 (32-bit, 64-bit) • POWER4 • POWER5 • POWER6	Yes	• JFS • JFS2 • GPFS™ 3.2 • NFS • Veritas File System (VxFS) r4.x and 5.x	Veritas Volume Manager (VxVM) 4.x and 5.x
	AIX 6.1 • POWER5 • POWER6	Yes		
	AIX Virtual I/O server 1.5.2 or later (32-bit, 64-bit)	Yes		
HP-UX, see "Notes for HP-UX"	HP-UX 11i v3 (• Itanium • 64-bit	Yes	• HP_HFS • NFS • Veritas File System (VxFS) r4.x and 5.x	• HP-UX Logical Volume Manager • Veritas Volume Manager (VxVM) 4.x and 5.x

Table 7. Operating system support for the Tivoli Storage Productivity Center agents (continued)

Platform		Storage Resource agent	File systems	Volume managers
Red Hat Linux, see "Notes for Linux"	Red Hat Enterprise Linux AS 4.0 • x86 • POWER5 and POWER6 • zSeries (31-bit, 64-bit)	Yes	• EXT2 • EXT3 • NFS • TMPFS • Veritas File System (VxFS) r4.x and 5.x	Veritas Volume Manager (VxVM) 4.x and 5.x
	Red Hat Enterprise Linux Base Server and Advanced Platform 5: • x86, x86-64 • POWER5 and POWER6 • zSeries (31-bit, 64-bit)	Yes		
SUSE Linux, see "Notes for Linux"	SUSE Linux Enterprise Server 9 • x86, x86-64 • POWER5 and POWER5 (32-bit, 64-bit) • zSeries (31-bit, 64-bit)	Yes	• EXT2 • EXT3 • NFS • REISERFS • TMPFS • Veritas File System (VxFS) r4.x and 5.x	Veritas Volume Manager (VxVM) 4.x and 5.x
	SUSE Linux Enterprise Server 10 • x86, x86-64 • POWER5 and POWER6 (32-bit, 64-bit) • zSeries (31-bit, 64-bit)	Yes		
	SUSE Linux Enterprise Server 11 • x86, x86-64 • POWER5 and POWER6 (32-bit, 64-bit) • zSeries (31-bit, 64-bit)	Yes		
Oracle Solaris	Solaris 9 • SPARC architecture • 32-bit, 64-bit	Yes (patches are required, see 1 on page 15)	• NFS • TMPFS • UFS • Veritas File System (VxFS) r4.5 and 5.x	Veritas Volume Manager (VxVM) 4.x and 5.x
	Solaris 10 • SPARC architecture • 32-bit, 64-bit	Yes see 2 on page 15		

Table 7. Operating system support for the Tivoli Storage Productivity Center agents (continued)

Platform		Storage Resource agent	File systems	Volume managers
Windows, see "Notes for Windows"	Windows 2003 • Data Center Edition • Enterprise Edition • Standard Edition • x86, x86-64	Yes	• FAT • FAT32 • NTFS	
	Windows 2003 R2 • Data Center Edition • Enterprise Edition • Standard Edition • x86, x86-64	Yes		
	Windows 2008 • Data Center Edition • Enterprise Edition • Standard Edition • x86, x86-64	Yes		
	Windows 2008 R2 • Data Center Edition • Enterprise Edition • Standard Edition • x86, x86-64	Yes		
VM on VMware, see "Notes for VMware"	VMware ESX 3.0.x, 3.5.x, 4.0.x, and 4.1.x (guest operating system): • Red Hat Enterprise Linux AS 4.0 (x86) • Re Hat Enterprise Linux 5 (x86) • SUSE Linux Enterprise Server 9 (x86) • SUSE Linux Enterprise Server 10 (x86) • SUSE Linux Enterprise Server 11 (x86) • Windows 2003 (x86) • Windows 2008 (x86)			

Notes for Windows operating systems

Important:

1. For information about IPv6 support in Windows Server 2003 Enterprise Server, see http://support.microsoft.com/kb/325449.

Notes for AIX operating systems

Important:

1. The following AIX fixes must be applied:

   For AIX 5.3, the minimum maintenance levels required are:

   - Technology Level 5300-11 Service Pack 4
   - Technology Level 5300-12 Service Pack 1

   For AIX 6.1, the minimum maintenance level required are:

   - Technology Level 6100-04 Service Pack 5
   - Technology Level 6100-05 Service Pack 1
   - Technology Level 6100-06

   You can determine the Technology Level and Service Pack installed by running the following command:

   ```
   oslevel -s
   ```

   You can download the package from http://www.ibm.com/eserver/support/fixes/fixcentral/pfixpacks/53 or http://www.ibm.com/eserver/support/fixes/fixcentral/pfixpacks/61.
2. AIX 6.1.0.4 or later comes with IBM Systems Director. IBM Systems Director has a common agent service that can use port 9510. The Storage Resource agent also uses port 9510. If you do not need IBM Systems Director, you can uninstall IBM Systems Director before installing Tivoli Storage Productivity Center. If you install Tivoli Storage Productivity Center and require IBM Systems Director, specify a different port for the Storage Resource agent. For information about the ports IBM Systems Director uses, go to http://publib.boulder.ibm.com/infocenter/director/v6r2x/index.jsp?topic=/com.ibm.director.plan.helps.doc/fqm0_r_ports_for_the_management_server.html.
3. AIX 6.1 comes with Lightweight Infrastructure (LWI). If this package is enabled, the package can cause Tivoli Storage Productivity Center installation issues:

   - There can be a port conflict with Tivoli Storage Productivity Center at port 9510.
   - Can cause the Storage Resource agent installation to fail on hosts where only the agent is installed.
4. Tivoli Storage Productivity Center also supports PowerHA® on AIX (previously known as HACMP™) for the following versions:

   - 5.5.x
   - 5.4.0.1 (APAR IY87447) or later
   - 5.3.0.4 (APAR IY87534) or later

   Also check the PowerHA version compatibility matrix on the support site at http://www-03.ibm.com/support/techdocs/atsmastr.nsf/84279f6ed9ffde6f86256ccf00653ad3/c03127c863a313898625711100635c80?OpenDocument
5. Your fabric probe or zone control job can fail on AIX with error message: AGT0430I.

AIX has a constraint wherein the size of the response buffer passed in for GS-3 commands made to the switch cannot exceed 4 K. The Storage Resource agent passes 4 K buffer sizes for AIX.

If the zone configuration for a fabric probe or zone control job exceeds the 4 K limitation for AIX, Tivoli Storage Productivity Center flags this condition. The fabric probe or zone control operation fails under this condition with error message AGT0430I.

Consider using Storage Resource agents deployed on another platform that is connected to the same fabric. Use this Storage Resource agent for collecting zone information and for performing zone changes. In this case, the fabric functions for the Storage Resource agent that is failing can be disabled.

If the fabric probe fails with the error message, and you try to make zone changes for the fabric, then you might get a warning message that the zoning has changed. Also, the zone control job log displays the AGT0432E error message.

6. Use the GNU tar program version 1.14 or later rather than the native AIX tar program to extract files from Tivoli Storage Productivity Center installation images. The AIX tar program might truncate long file names, which can cause errors in Total Storage Productivity Center and Tivoli Integrated Portal installation.

   The GNU tar file archiving program can be found from the following link. http://www-03.ibm.com/systems/power/software/aix/linux/toolbox/alpha.html.

Notes for Linux operating systems

Important:

1. The Tivoli Storage Productivity Center installation program is unable to run in graphical mode on Red Hat Enterprise Linux AS Version 4 and 5 or SUSE Linux Enterprise Server version 9 and 10. You must install a Linux RPM package on these operating systems. For more information see http://www-01.ibm.com/support/docview.wss?uid=swg21303823.

   For Storage Resource agents installed on Red Hat Linux AS release 4, you must have Update 5. You must have the following package installed:

   ```
   compat-libstdc++-33
   ```

   For Storage Resource agents installed on Red Hat Linux, you must have the following package installed:

   ```
   compat-libstdc++-296-2.96-138.i386.rpm
   ```

2. For Storage Resource agent for fabrics connected to hosts running on Linux IBM Power Systems, or zSeries hardware:
   - Does not allow zone configuration changes
   - Does not report on HBA, fabric topology, or zoning information

   These limitations also apply to Storage Resource agents on all guest operating systems for VMware configurations.
3. Tivoli Storage Productivity Center does not support software RAID on Linux systems. A Storage Resource agent does not install and run on Linux systems with software RAID.
4. Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

Notes for HP-UX operating systems

Important:

1. Tivoli Storage Productivity Center agents running on HP-UX must have the HP `libc cumulative patch PHC_34275` installed on those machines. To download patch PHC_34275, go to this Web site: https://www1.itrc.hp.com/service/home/home.do.

Notes for Oracle Solaris operating systems

Important:

1. To install Storage Resource agents on Solaris 9, patches 111711 and 111712 are required. For information about downloading patches and obtaining Solaris 9 support, contact Oracle.
2. For Solaris 10, only Global Zone is supported for the Tivoli Storage Productivity Center agents. The Non-Global (or local) zones are not supported.

   If you upgrade from Solaris 9 to Solaris 10, all agents need to be reinstalled completely.

Notes for VMware

Important:

1. For more information about VMware platforms, see "Planning for VMware" on page 132.
2. The Storage Resource agent supports the VMware ESX environment in tolerance mode only. No fabric functions are available with the Storage Resource agent for any guest operating systems.

Software requirements for the database repository

Use this information to understand what the software requirements are for the database repository.

Tivoli Storage Productivity Center supports DB2 as the database repository.

DB2 database

Tivoli Storage Productivity Center supports the following products for the repository database:

- DB2 Enterprise Server Edition version 9.1 with fix pack 2 or later.
- DB2 Enterprise Server Edition version 9.5 with fix pack 3a or later.
- DB2 Enterprise Server Edition version 9.7

  Note:

  – For DB2 9.7, use the version of DB2 shipped with Tivoli Storage Productivity Center. Do not use DB2 9.7 with Fix Pack 1 or Fix Pack 2 (this causes issues with Tivoli Storage Productivity Center).

  – For DB2 9.5, do not use Fix Pack 6 (this causes issues with Tivoli Storage Productivity Center).

  – Tivoli Storage Productivity Center does not support the DB2 that is installed in the LDAP environment.

  – You can use the LDAP or Domain user ID for the Tivoli Storage Productivity Center Administrator; however, you must use a local OS user ID for DB2.

DB2 9.1

- Tivoli Storage Productivity Center supports DB2 V9.1 with fix pack 2 or later (32–bit). For AIX, use 64–bit. On Windows 2008, use DB2 Enterprise Server Edition version 9.1 with fix pack 5 or later.
- DB2 version 9.1 does not support Red Hat Enterprise Linux AS Version 3. Red Hat Enterprise Linux AS Version 4 (32–bit) is supported. For information about DB2 Linux support, see http://www-306.ibm.com/software/data/db2/linux/validate/.
- Tivoli Storage Productivity Center supports DB2 V9.1 with fix pack 2 on AIX and Windows in 64–bit mode.

  For AIX 5.3 64–bit kernel, the minimum C++ runtime level is `xlC.rte 8.0.0.4` and `xlC.aix50.rte 8.0.0.8`. You can check for the level by running the following command:

  ```
  lslpp -al | grep xlC
  ```

  Also, the Technology Level must be 5300.08-2 and SP 5300-04-2. You can check for this level by running the following command:

  ```
  oslevel -r
  ```

  For current information about DB2 operating system support, see http://www-01.ibm.com/support/docview.wss?&uid=swg21165448 http://www-01.ibm.com/support/docview.wss?&uid=swg21165448 .
- DB2 version 9.1 supports i386 GNU/LINUX only when the kernel is at least 2.6.9.
- For DB2 9.1 on AIX, a 64-bit AIX kernel is required.

DB2 9.5

DB2 9.5 provides increased performance and scalability and improved security. For more information, see http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/index.jsp?topic=/com.ibm.db2.luw.wn.doc/doc/c0023235.html.

Note: For DB2 9.5 on AIX, a 64-bit AIX kernel is required.

DB2 9.7

DB2 9.7 provides more manageability, monitoring, performance, security, and high-availability enhancements over DB2 9.5. For more information, see http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.wn.doc/doc/c0052035.html.

Note:

- For DB2 9.7, use the version of DB2 shipped with Tivoli Storage Productivity Center. Do not use DB2 9.7 with fix pack 1 or fix pack 2 (this causes issues with Tivoli Storage Productivity Center). This issue will be fixed in a future DB2 9.7 fix pack.
- DB2 Version 9.7 does not support Red Hat Linux version 4 or SUSE Linux version 9. (Note that the Tivoli Storage Productivity Center server is also not supported on SUSE Linux.) You must upgrade your operating systems to a supported level before attempting to upgrade to DB2 Version 9.7.
- For DB2 9.7 on AIX, a 64-bit AIX kernel is required.

The Windows platform is enabled for both IPv6 and IPv4 addressing. If the Windows platform is used in IPv6–only scenarios, all communication must

be done over IPv6. Also, DB2 V9.1, DB2 V9.5, or DB2 V9.7 has to be installed. In this scenario, Windows still has an IPv4 stack active, even if no IPv4 address is configured.

When you install Tivoli Storage Productivity Center, a single, shared database instance is created on DB2 for the Data server and Device server. The default database name is TPCDB.

At installation time, you can specify whether the DB2 database is system-managed space (SMS) or database-managed space (DMS). In SMS, the file system manager for the operating system allocates and manages the space where the table is stored. In DMS, the database manager controls the storage space. The administrator decides which devices to use, and DB2 manages the space on those devices.

Note: Any version of Agent Manager 1.3.2 supports DB2 9.1. For DB2 9.5 support, you need to use Agent Manager version 1.3.2.30 which is shipped with Tivoli Storage Productivity Center 4.1.1. If you are planning to use DB2 9.7, you must install Agent Manager 1.4.x or later.

Web browser support

IBM Tivoli Storage Productivity Center starts a web browser when you access web pages from items in the Help menu, launch another application using its launch-in-context feature, or start IBM Tivoli Storage Productivity Center for Replication from its user interface.

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 Tivoli Storage Productivity Center GUI, you must use the IBM Java.

You can use the following web browsers with Tivoli Storage Productivity Center:

AIX

- Mozilla Firefox 2.0
- Mozilla Firefox 3.0

Linux

- Mozilla Firefox 2.0
- Mozilla Firefox 3.0

Windows

- Internet Explorer 7
- Internet Explorer 8 (in Compatibility View)

  In Internet Explorer, the Computability View option allows you to display a web page that was designed for a version older than version 8 to be displayed correctly in Internet Explorer 8. To display your web page in Internet Explorer 8.0, click **Tools** > **Compatibility View Settings**.
- Mozilla Firefox 2.0
- Mozilla Firefox 3.0

Note:

- To get Firefox 2 or 3 for Linux or Windows, go to ftp://ftp.mozilla.org/pub/mozilla.org/firefox/releases/.
- To get Firefox version 2 on AIX, you can use a previous version of IBM Firefox Browser for AIX CD (if you have it). Or you can download the source code for version 2 or 3 from ftp://ftp.mozilla.org/pub/mozilla.org/firefox/releases/.

- If you do not have a web browser configured for use with Tivoli Storage Productivity Center when you access a web page from its user interface, the Configure Element Launcher window is displayed. Use this window to configure a web browser for use with the product.
- For information about the web browsers supported for the DS8000, see the DS8000 Information Center at http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp. Search for **Internet browser support**.
- For information about starting Java Web Start using Internet Explorer 8, see Starting the Tivoli Storage Productivity Center GUI using a Web browser.

The pop-up blocker 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 V7000
- Storwize V7000 Unified

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 you receive a message that content has been blocked because it was not signed by a valid security certificate, click the Information Bar at the top of the window and select **Show blocked content**.

Web browser support for Tivoli Storage Productivity Center for Replication

When you start Tivoli Storage Productivity Center for Replication from within the Tivoli Storage Productivity Center, its user interface appears in a web browser. Depending on the configuration settings for the web browser on your computer, Tivoli Storage Productivity Center for Replication appears within a tab of an existing browser session or it appears within a new browser session.

In Firefox, you can configure how Tivoli Storage Productivity Center for Replication appears when launched from Tivoli Storage Productivity Center. To configure Firefox, complete the following steps:

1. Start Firefox.
2. Select **Tools** > **Options** from the menu bar.
3. Select **Tabs**.
4. Select **a new window** to indicate that you want start Tivoli Storage Productivity Center for Replication in a new web browser session. Select **a new tab** to indicate that you want to start Tivoli Storage Productivity Center for Replication in a new tab of an existing web browser session.

   Note: For Internet Explorer, Tivoli Storage Productivity Center for Replication is always started in a new web browser session.

Note: If you start Tivoli Storage Productivity Center for Replication from within Tivoli Storage Productivity Center, you might be logged out of the Tivoli Storage Productivity Center for Replication user interface unexpectedly. This action occurs

when you use Tivoli Storage Productivity Center to start a session of Tivoli Storage Productivity Center for Replication in an Internet Explorer web browser and then open and close a wizard or secondary window in that Tivoli Storage Productivity Center for Replication session.

To work around this problem, use one of these solutions:

- Start Tivoli Storage Productivity Center for Replication from Tivoli Storage Productivity Center using Firefox
- Start Tivoli Storage Productivity Center for Replication from Tivoli Integrated Portal
- Start Tivoli Storage Productivity Center for Replication by entering its web address directly in a Firefox or Internet Explorer web browser.

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 launches 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 configured to successfully access these web pages.

Software requirements for LDAP servers

For information about the LDAP servers supported, see http://www-01.ibm.com/support/docview.wss?rs=180&uid=swg27007642. Click the relevant operating system and then navigate to the row labeled "LDAP Servers using Federated Repository Configuration."

Software requirements for CIM agents

CIM agents for DS3000, DS4000, DS5000, DS6000, and tape

IBM Tivoli Storage Productivity Center requires CIM agents for some IBM storage devices such as the DS3000, DS4000, DS5000, DS6000, and tape.

For information about installing and configuring the CIM agent for the DS3000, DS4000, and DS5000, contact your CIM agent (LSI SMI Provider) at http://www.lsi.com/storage_home/products_home/external_raid/management_software/smis_provider/index.html.

The LSI CIM agent must be installed and configured before you can collect data for the DS3000, DS4000, and DS5000 through.IBM Tivoli Storage Productivity Center When you configure IBM Tivoli Storage Productivity Center to connect to the CIM agent, you must provide the interoperability namespace for the device.

For information about the LSI CIM agent level required for each device, see http://www-01.ibm.com/support/docview.wss?rs=40&context=SSBSEX&q1=subsystem&uid=swg21384734&loc=en_US&cs=utf-8&lang=en. At the bottom of the page is a subsystem support matrix Excel spreadsheet that provides the device level support and Interoperability Namespaces.

Also see *TPC Hints and Tips* document at this Web site http://www-01.ibm.com/support/docview.wss?uid=swg27008254 for some useful information about installing the LSI CIM agent.

Also note the following information about the CIM agents from LSI:

- Do not install the CIM agent on the same server as the IBM Tivoli Storage Productivity Center server.
- Do not have the CIM agent monitor more than five storage systems.
- Use the LSI SMI-S Provider V10.06.GG.33 or later for the DS4000 and DS5000 MC level V7.60.x. Otherwise, performance monitoring does not work.
- Do not use the LSI SMI-S agent for the DS4000 and DS5000 with an MC level lower than V7.50.x. A lower level of microcode can cause a reboot of the controller device support and SMI-S agent download.

CIM agents for switches and directors

CIM agents (SMI-S Providers) are also required for switches and directors. For a complete list of switches and directors supported and the level of CIM agents required for each device, see http://www-01.ibm.com/support/docview.wss?rs=40&context=SSBSEX&uid=swg21384219&loc=en_US&cs=utf-8&lang=en.

CIM agents for non-IBM storage systems

IBM Tivoli Storage Productivity Center also supports some non-IBM storage systems. For example, IBM Tivoli Storage Productivity Center supports HP, TagmaStore, CLARiiON, Symmetrix, and other storage systems.

For a complete list of storage systems supported and the level of CIM agents required for each device, see http://www-01.ibm.com/support/docview.wss?rs=40&context=SSBSEX&q1=subsystem&uid=swg21384734&loc=en_US&cs=utf-8&lang=en. At the bottom of the page is a subsystem support matrix Excel spreadsheet that provides the device level support and Interoperability Namespaces.

Running IBM Tivoli Storage Productivity Center in a 64-bit environment

This section provides information for running Tivoli Storage Productivity Center in a 64-bit environment.

If you are running Tivoli Storage Productivity Center in a 64-bit environment, note the following conditions:

- All Tivoli Storage Productivity Center programs that run in a 64-bit environment is run in 32-bit compatibility mode for Tivoli Storage Productivity Center.
- The Tivoli Storage Productivity Center agents must be running in a 32-bit native mode or 32-bit compatibility mode environment.
- The databases to be monitored by Tivoli Storage Productivity Center must be in a 32-bit native mode or 32-bit compatibility mode environment.
- Tivoli Storage Productivity Center can monitor DB2 in 64-bit native mode on Windows and AIX as long as the DB instance is created in 32-bit mode.
- When installing Tivoli Storage Productivity Center on a 64-bit Windows system, the default directory appears like this directory path:

  ```
  C:\Program Files (x86)\IBM\TPC
  ```

  This default location must be changed for the installation to be successful. The suggested installation folder name is:

  ```
  C:\IBM\TPC or
  c:\program files\ibm\tpc
  ```

When installing DB2 on a 64-bit Windows system, change the default installation path so that it is NOT:

```
c:\Program Files (x86)\IBM\SQLLIB or
c:\Program Files\IBM\SQLLIB
```

The Windows installer puts back the (x86) into the directory path and this action prevents the Tivoli Storage Productivity Center Device Server from starting.

Once installed, the Tivoli Storage Productivity Center program runs in Windows 32-bit compatibility mode.

Planning for IBM Tivoli Storage Productivity Center for Replication

IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication, previously separated products, are now integrated. You can launch IBM Tivoli Storage Productivity Center for Replication from the IBM Tivoli Storage Productivity Center user interface.

Integration highlights

This integration enables you to:

- Start the IBM Tivoli Storage Productivity Center for Replication GUI from within the IBM Tivoli Storage Productivity Center GUI.
- Use the IBM Tivoli Storage Productivity Center GUI to set up IBM Tivoli Storage Productivity Center for Replication SNMP alerts and Tivoli Enterprise Console events.
- Provide a IBM Tivoli Storage Productivity Center super user role that has authority over all IBM Tivoli Storage Productivity Center commands. IBM Tivoli Storage Productivity Center for Replication has a replication administrator role that has authority to all IBM Tivoli Storage Productivity Center for Replication commands. IBM Tivoli Storage Productivity Center for Replication honors the IBM Tivoli Storage Productivity Center super user role giving the super user role authority over all IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication commands.

For more information about IBM Tivoli Storage Productivity Center for Replication, see http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp. Click **Tivoli Storage Productivity Center > Product Overview**.

Disaster recovery

When you install IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication, you can specify to use LDAP authentication or OS authentication.

When planning for disaster recovery, the IBM Tivoli Storage Productivity Center for Replication server must be accessible in the event of a disaster. Even if there is a IBM Tivoli Storage Productivity Center for Replication standby server for recovery purposes, the server still requires authentication. If you use LDAP authentication, the LDAP services must be available in the event of a disaster. You can set up the LDAP environment so that you have clustered servers for this purpose.

Tivoli Storage Productivity Center for Replication licenses

When you install Tivoli Storage Productivity Center for Replication, there are several licenses that you can purchase. See Table 8.

Table 8. Tivoli Storage Productivity Center for Replication licenses

Tivoli Storage Productivity Center for Replication license	**Includes**
Basic (installed when you install Tivoli Storage Productivity Center)	• FlashCopy sessions • Snapshot sessions
Two Site Business Continuity	Includes everything in basic license plus: • All Metro Mirror and Global Mirror session types • High availability functions
Three Site Business Continuity	Includes everything in basic license, Two Site Business Continuity license, plus: • Metro Global Mirror session types

Planning for Open HyperSwap replication

Open HyperSwap® replication is a special Metro Mirror replication method designed to automatically failover I/O from the primary logical devices to the secondary logical devices in the event of a primary disk storage system failure. This function can be done with minimal disruption to the applications that are using the logical devices.

Overview

Open HyperSwap replication applies to both planned and unplanned replication sessions. When a session has Open HyperSwap enabled, an I/O error on the primary site automatically causes the I/O to switch to the secondary site without any user interaction and with minimal application impact. In addition, while Open HyperSwap is enabled, the Metro Mirror session supports disaster recovery. If a write is successful on the primary site but is unable to get replicated on the secondary site, IBM Tivoli Storage Productivity Center for Replication suspends the entire set of data consistency checking, thus ensuring that a consistent copy of the data exists on the secondary site. If the system fails, this data might not be the latest data, but the data should be consistent and allow the user to manually switch host servers to the secondary site.

You can control Open HyperSwap from any system running IBM Tivoli Storage Productivity Center for Replication (AIX, Windows, Linux, or z/OS). However, the volumes that are involved with Open HyperSwap must be attached to an AIX system. The AIX system is then connected to Tivoli Storage Productivity Center for Replication.

Software and hardware requirements

There are several requirements for Open HyperSwap support:

AIX requirements

Open HyperSwap support requires AIX version 5.3 or 6.1. (You can find the supported AIX version for each Tivoli Storage Productivity Center for Replication release in the support matrix at http://www-01.ibm.com/

support/docview.wss?rs=40&context=SSBSEX&context=SSMN28
&context=SSMMUP&context=SS8JB5&context=SS8JFM&uid=swg21386446
&loc=en_US&cs=utf-8&lang=en. 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 version 5.3 or 6.1)

DS8000 hardware requirements
Only DS8000 storage systems are supported. Open HyperSwap requires DS8000 5.1 or later.

Note: Open HyperSwap does not support PowerHA (previously known as High Availability Cluster Multi-processing (HACMP)).

General tasks

Before you can use Open HyperSwap, you must set up your environment for this function. The general steps are:

1. Prepare the AIX system for Open HyperSwap. Use the AIX configuration manager (**cfgmgr**) to identify all volumes that are involved with the Open HyperSwap session.
2. Set up the host connection of Tivoli Storage Productivity Center for Replication to the AIX system. Use the Tivoli Storage Productivity Center for Replication user interface to manually set up the connection to the AIX system. Use the Host Systems page to enter the IP address and port number for the AIX system.
3. Set up the Tivoli Storage Productivity Center for Replication Metro Mirror Failover/Failback session, selecting the function **Manage H1-H2 with Open HyperSwap**.
4. Add the copy sets to the session where all the volumes in the copy sets are volumes that are on the AIX system that is connected to Tivoli Storage Productivity Center for Replication.
5. You can now start your Open HyperSwap session the same as other sessions.

General planning

Use this information to plan your IBM Tivoli Storage Productivity Center environment. Topics include general configuration guidelines, TCP/IP ports used, user IDs and user rights, and SMI-S support.

Licenses for Tivoli Storage Productivity Center

The following product licenses are available for IBM Tivoli Storage Productivity Center: Basic, Data, Disk, Select, and Standard Editions. Each license determines the functions that are accessible in the user interface. This topic lists the functions that are included with each license to help you determine which one best meets your storage management needs.

IBM Tivoli Storage Productivity Center Standard Edition
Contains all the functions for data management, disk management, fabric management, and tape management, plus analytical functions:

- Configuration Analysis
- Configuration History

- SAN Planner
- Storage Optimizer

IBM Tivoli Storage Productivity Center Select
Contains all the functions for data management, disk management, fabric management, and tape management, plus the following analytical functions:

- Configuration Analysis
- Configuration History
- SAN Planner
- Storage Optimizer

These analytical functions are available for the following storage systems:

- System Storage DS3000
- System Storage DS4000
- System Storage DS5000
- XIV system
- SAN Volume Controller
- Storwize V7000

IBM Tivoli Storage Productivity Center for Data Edition
Includes data management, basic tape, disk, and fabric management, database monitoring, and chargeback functions, but does not include performance monitoring functions.

IBM Tivoli Storage Productivity Center for Disk Edition
Includes basic disk, fabric, tape, and data management functions and storage system performance monitoring, but does not include fabric performance monitoring, chargeback, and database monitoring functions.

IBM Tivoli Storage Productivity Center for Disk Select

Provides storage device configuration, performance monitoring, and management of Storage Area Network (SAN) attached devices from a single console.

This product provides basic disk functions and performance monitoring capabilities for the following storage systems:

- System Storage DS3000
- System Storage DS4000
- System Storage DS5000
- XIV system
- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

IBM Tivoli Storage Productivity Center Basic Edition
Includes basic disk, fabric, tape, and data management functions, but does not include chargeback, database monitoring, and performance monitoring functions.

Use the following tables to view the specific functions, reports, and elements of the user interface that are available in each license.

Table 9. Available functions in each Tivoli Storage Productivity Center license

Function	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
Administrative Services						
Device discovery	Yes[1]	Yes[1]	Yes[1]	Yes	Yes	Yes
Data Sources	Yes[2]	Yes[2]	Yes[2]	Yes[7]	Yes	Yes
Configuration	Yes[3]	Yes[3]	Yes[3]	Yes	Yes	Yes
IBM Tivoli Storage Productivity Center						
Configuration Utility	Yes	Yes	Yes	Yes	Yes	Yes
Configure Devices	Yes	Yes	Yes	Yes	Yes	Yes
Job Management	Yes	Yes	Yes	Yes	Yes	Yes
Reporting • Batch Reports • User ID reports • System Reports • Rollup Reports • Data Source Reports	Yes Not all reports are available in the Basic edition. See Table 10 on page 28 for a list of available reports.	Yes Not all reports are available in the Disk edition. See Table 10 on page 28 for a list of available reports.	Yes Not all reports are available in the Select edition. See Table 10 on page 28 for a list of available reports.	Yes Not all reports are available in the Data edition. See Table 10 on page 28 for a list of available reports.	Yes See Table 10 on page 28 for a list of available reports.	Yes See Table 10 on page 28 for a list of available reports.
Topology Viewer	Yes	Yes	Yes	Yes	Yes	Yes
Monitoring • Probes • TPC Server Probes	 • Yes[4] • No	 • Yes[4] • No	 • Yes[4] • No	 • Yes • No	 • Yes • Yes	 • Yes • Yes
Storage Resource Group Management	Yes	Yes	Yes	Yes	Yes	Yes
Analytics • Configuration History • Storage Optimizer • SAN Planner • Configuration Analysis	 • No • No • No • No	 • No • No • No • No	 • No • No • No • No	 • No • No • No • No	 • Yes • Yes • Yes • Yes	 • Yes • Yes • Yes • Yes
Alerting • Alert Log • Authentication Configuration Alerts	 • Yes[5] • Yes	 • Yes[5] • Yes	 • Yes[5] • Yes	 • Yes • Yes	 • Yes • Yes	 • Yes • Yes

Table 9. Available functions in each Tivoli Storage Productivity Center license (continued)

Function	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
Data Manager						
Monitoring • Groups • Pings • Scans • Profiles	 • Yes • Yes • No • No	 • Yes • Yes • No • No	 • Yes • Yes • No • No	 • Yes • Yes • Yes • Yes	 • Yes • Yes • Yes • Yes	 • Yes • Yes • Yes • Yes
Alerting	Yes[6]	Yes[6]	Yes[6]	Yes	Yes	Yes
Policy Management	No	No	No	Yes	Yes	Yes
Reporting	Yes Not all reports are available in the Basic edition. See Table 10 on page 28 for a list of available reports.	Yes Not all reports are available in the Disk edition. See Table 10 on page 28 for a list of available reports.	Yes Not all reports are available in the Select edition. See Table 10 on page 28 for a list of available reports.	Yes	Yes	Yes
Data Manager for Databases	No	No	No	Yes	Yes	Yes
Data Manager for Chargeback	No	No	No	Yes	Yes	Yes
Disk Manager						
Storage Subsystems • Create Volumes • Create Virtual Disks • Launch Element Manager	Yes • Yes • Yes • Yes	Yes • Yes • Yes • Yes	Yes • Yes • Yes • Yes	Yes • Yes • Yes • Yes	Yes • Yes • Yes • Yes	Yes • Yes • Yes • Yes
Storage Optimizer	No	No	No	No	Yes	Yes
SAN Planner	No	No	No	No	Yes	Yes
Monitoring • Groups • Jobs • Storage Subsystem Performance	 • Yes • Yes • No	 • Yes • Yes • Yes	 • Yes • Yes • Yes	 • Yes • Yes • No	 • Yes • Yes • Yes	 • Yes • Yes • Yes

Table 9. Available functions in each Tivoli Storage Productivity Center license (continued)

Function	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
Alerting	Yes	Yes	Yes	Yes	Yes	Yes
Profile Management	No	Yes	Yes	No	Yes	Yes
Reporting • Groups • Storage Subsystems • Storage Subsystem Performance	 • Yes • Yes • No	 • Yes • Yes • Yes	 • Yes • Yes • Yes	 • Yes • Yes • No	 • Yes • Yes • Yes	 • Yes • Yes • Yes
Fabric Manager						
Fabrics	Yes	Yes	Yes	Yes	Yes	Yes
Monitoring • Groups • Jobs • Switch Performance Monitors	 • Yes • Yes • No	 • Yes • Yes • No	 • Yes • Yes • No	 • Yes • Yes • No	 • Yes • Yes • Yes	 • Yes • Yes • Yes
Alerting	Yes	Yes	Yes	Yes	Yes	Yes
Reporting: Switch Performance	No	No	No	No	Yes	Yes
Tape Manager	Yes	Yes	Yes	Yes	Yes	Yes
Element Manager	Yes	Yes	Yes	Yes	Yes	Yes
Replication Manager	Yes	Yes	Yes	Yes	Yes	Yes

Table 9. Available functions in each Tivoli Storage Productivity Center license (continued)

Function	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition

Notes:

1. You cannot run the following discovery jobs in the Basic, Disk, and Disk Select editions: Netware Filer, Windows Domain, NAS, SAN FS, and VMware VI Data Source.
2. You cannot associate VMware VI Data Sources or TPC Servers with Tivoli Storage Productivity Center in the Basic, Disk, and Disk Select editions.
3. The following nodes under **Configuration** are not available in the Basic, Disk, and Disk Select editions.
   - License Keys
   - Quota and Constraint email Address Rules
   - Scan/Probe Agent Administration
   - Manual NAS/Netware Server Entry
   - NetWare Tree Logins
   - Resource History Retention for Databases
   - Removed Resource Retention for Databases
   - Configuration History Settings (also not available in the Data edition)
4. Use probes and Storage Resource agents to collect asset and capacity information about storage entities in the Disk, Basic, and Disk Select editions. To collect specific file level information (for example, usage statistics), you must use the Data or Standard editions to run scans on your file systems.
   Note: You cannot probe hypervisors in the Basic, Disk, and Select editions.
5. The following nodes under **Alert Log** are not available in the Basic, Disk, and Disk Select editions:
   - Directory
   - User
   - OS User Group
   - Configuration Analysis (also not available in the Data edition)
   - Hypervisor
6. The following nodes under **Data Manager > Alerting** are not available in the Basic, Disk, and Disk Select editions:
   - Directory Alerts
   - Hypervisor Alerts
   - Other NAS Alerts
7. Data Sources for Data Edition does not support TPC Server Probes.

Table 10. Available reports in each IBM Tivoli Storage Productivity Center license

Reports	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
Batch reports	No	Yes[1, 2]	Yes[1, 2]	Yes[2]	Yes	Yes
System Reports: • Data • Fabric • Disk System Reports:	 • Yes[5] • Yes[3] • No	 • Yes[5] • Yes[3] • Yes	 • Yes[5] • Yes[3] • Yes	 • Yes • Yes[3] • No	 • Yes • Yes • Yes	 • Yes • Yes • Yes
Rollup Reports	No	No	No	No	Yes	Yes

Table 10. Available reports in each IBM Tivoli Storage Productivity Center license (continued)

Reports	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
Data Source Reports	Yes	Yes	Yes	Yes	Yes	Yes
Asset reports • By Cluster • By Computer • By Hypervisor • By Storwize V7000U • By OS Type • By Storage Subsystem • System-wide	 • Yes • Yes • No • No • Yes • Yes • Yes[4]	 • Yes • Yes • No • No • Yes • Yes • Yes[4]	 • Yes • Yes • No • No • Yes • Yes • Yes[4]	 • Yes • Yes • Yes • Yes • Yes • Yes • Yes	 • Yes • Yes • Yes • Yes • Yes • Yes • Yes	 • Yes • Yes • Yes • Yes • Yes • Yes • Yes
Availability reports	Yes	Yes	Yes	Yes	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	 • Yes • Yes • Yes • Yes	 • Yes • Yes • Yes • Yes	 • Yes • Yes • Yes • Yes	 • Yes • Yes • Yes • Yes
Usage reports:	No	No	No	Yes	Yes	Yes
Usage Violation reports:	No	No	No	Yes	Yes	Yes
Backup reports:	No	No	No	Yes	Yes	Yes
Monitored Computer Storage Space reports	Yes	Yes	Yes	Yes	Yes	Yes
Storage Subsystem reports	Yes	Yes	Yes	Yes	Yes	Yes
Storage Subsystem Performance reports	No	Yes	Yes	No	Yes	Yes
Switch Performance reports	No	No	No	No	Yes	Yes
Tape Library report	Yes	Yes	Yes	Yes	Yes	Yes

Table 10. Available reports in each IBM Tivoli Storage Productivity Center license (continued)

Reports	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
Notes: 1. You must deploy Storage Resource agents to enable batch reports for these licenses. 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 in the following editions: • Basic edition: No batch reports are available • Disk edition: Rollup Reports, Usage, Usage Violations, Backup, Switch Performance • Disk Select edition: Rollup Reports, Usage, Usage Violations, Backup, Switch Performance • Data edition: Storage Subsystem Performance, Switch Performance 3. The following Fabric system reports are not available in the Basic, Disk, Disk Select, and Data editions: Top switch Ports Data Rate Performance, Top Switch Ports Packet Rate Performance, Switch Port Errors, Switch Performance. 4. The following Asset > System-wide reports are not available in the Basic, Disk, and Disk Select editions: Unmanaged Virtual Machines, Virtual Machines with no VMWare Agent, Monitored Directories, Unmanaged Computers, Users, OS User Groups. 5. The following Data system reports are available with Basic Edition, Disk, and Disk Select Edition: • Disk Space Summary • Disk Defects • Computer Storage Availability • Computer Disk Space • Available File System Space						

If you have purchased a license for IBM Tivoli Storage Productivity Center, Table 11 lists the products for which you can purchase a license upgrade for your system.

Table 11. License upgrade scenarios for IBM Tivoli Storage Productivity Center

IBM Tivoli Storage Productivity Center license	First license upgrade	Second license upgrade
Basic Edition	Disk Edition	Standard Edition
Basic Edition	Data Edition	Standard Edition
Basic Edition	Disk Select Edition	Standard Edition
Basic Edition	Standard Edition	
Disk Select Edition	Disk Edition	Standard Edition
Disk Select Edition	Data Edition	Standard Edition
Disk Select Edition	Standard Select Edition	Standard Edition
Disk Select Edition	Standard Edition	
Data Edition	Disk Select Edition	Standard Edition
Data Edition	Standard Edition	
Disk Edition	Disk Select Edition	Standard Edition
Disk Edition	Standard Edition	

Table 11. License upgrade scenarios for IBM Tivoli Storage Productivity Center (continued)

IBM Tivoli Storage Productivity Center license	First license upgrade	Second license upgrade
Standard Select Edition	Standard Edition	
Standard Edition		

General configuration guidelines

IBM Tivoli Storage Productivity Center must be configured properly for monitoring. This section provides some guidelines for configuring Tivoli Storage Productivity Center.

The following guidelines are recommended for environments that exceed any of the following characteristics:

- 100 Data agents or Storage Resource agents or a mixture of both
- 1000 Disk volumes
- 5 Disk subsystems
- 10 Fibre channel switches
- 100 Fibre channel ports

Here are some general guidelines to follow:

Database repository

For best performance, install the database repository across multiple physical disks, either through operating system striping or hardware RAID. A minimum of three disks is recommended. It is recommended that these disks be separate from the Tivoli Storage Productivity Center host server operating system location and Tivoli Storage Productivity Center product installation directory location.

IBM Tivoli Storage Productivity Center server host

The Tivoli Storage Productivity Center server host must meet the following requirements:

- The Tivoli Storage Productivity Center server must have dual 3.2 GHz processors with 8 GB of RAM.
- The Tivoli Storage Productivity Center server must be a dedicated computer for Tivoli Storage Productivity Center operations and not shared with other applications.
- Paging space: On AIX, the paging space must be increased to 1 GB.

IBM Tivoli Storage Productivity Center GUI

For large environments, the Tivoli Storage Productivity Center GUI must be installed on a separate system from the Tivoli Storage Productivity Center server.

For a stand-alone GUI, a minimum of 2 GHz processor (single), 1 GB RAM workstation for the GUI, and up to 2 GB of RAM is required if a large server is used.

An environment is considered large if you have 50% or more of the maximum for any Tivoli Storage Productivity Center recommended maximums. For example, an environment is considered large if you have 500 or more disk volumes.

CIM agents

CIM agents are not required for storage systems that Tivoli Storage

Productivity Center communicates with using the native interface, such as DS8000, IBM XIV Storage System, SAN Volume Controller, IBM Storwize V7000 Unified, and IBM Storwize V7000. For devices that require CIM agents to be set up, review the CIM agent documentation to understand how many devices the CIM agent can be configured to manage. Note that the memory consumed by the CIM agent includes the time it takes for Tivoli Storage Productivity Center to probe a device through the CIM agent. The CIM agent is typically affected by the number of devices that the CIM agent manages. If no guidelines are available in the disk array subsystem CIM agent documentation, it is recommended to limit three subsystems per CIM agent.

Storage Resource agents

With this release, Storage Resource agents are designed to reduce the load of management calls on the fabric. Here are some improvements that Storage Resource agents take into consideration:

- Storage Resource agents will not monitor HBAs for fabric events and will not notify the server regarding fabric events. (The HBA event monitoring and event scanner function of the in-band Fabric agent is not supported by the Storage Resource agents.) Instead, SNMP traps and CIM indications are used for reporting fabric events (and subsequently an automatic fabric probe if the fabric is being monitored by the Storage Resource agent). This reduces the number of event notifications while still monitoring fabrics using other mechanisms.
- An agent assignment algorithm is used to reduce the number of Storage Resource agents that are used to probe a fabric. Depending on the mix of agents capable of reporting on a fabric, fewer Storage Resource agents might be used if multiple agents can collect similar information. For example, if Storage Resource agents are deployed on many systems connected to the same fabrics, then the algorithm automatically determines the number of agents to be used or not used to probe the fabric.
- Tivoli Storage Productivity Center can selectively probe specific fabrics after an event notification if there is sufficient information to determine which fabrics triggered the events. Because Tivoli Storage Productivity Center can determine the likelihood of which fabrics triggered the events, the number of fabrics being probed can be reduced.
- For Storage Resource agents, there is an option to selectively disable fabric functions for one or more Storage Resource agents. If a Storage Resource agent is disabled for fabric functions, no management calls will be made to the fabric. You can use this option to complement agent assignment to control the number of agents being used for fabric functions.
- Storage Resource agents only probe fabrics that you are interested in monitoring (fabrics that are part of a probe definition).
- If there are many calls made by the Storage Resource agents, Tivoli Storage Productivity Center can reduce the number of fabric probes being performed within a time window to reduce the probe activity.
- Tivoli Storage Productivity Center monitors the Device server filters for SNMP traps. Tivoli Storage Productivity Center filters out traps that are not related to fabric events so that fabric probes are not automatically triggered as a result of SNMP traps received.

Storage subsystems

When probing storage subsystems that are registered with the same

Storage Resource agent, do not probe more than three storage subsystems within the same probe job. The increased load on the agent would increase the likelihood of timeouts. Instead, spread the storage subsystems across multiple probe jobs with different starting times.

LSI SMI-S Provider
Tivoli Storage Productivity Center supports up to five storage systems with the LSI SMI-S Provider 1.3 or later.

Planning for IBM Tivoli Storage Productivity Center authorization

When a user ID is authenticated to IBM Tivoli Storage Productivity Center through the GUI, CLI, or APIs, group membership (either an LDAP-compliant repository group or an operating system group) determines the authorization level of the user in IBM Tivoli Storage Productivity Center. This topic describes the user roles and authorization levels.

If you want to install IBM Tivoli Storage Productivity Center so that it authenticates its users against an LDAP-based repository, then you must ensure that your LDAP-based repository contains at least one group and at least one user who is a member of that group (and who also has a value in the userPassword attribute) before installation. You use this LDAP user and LDAP group as the LDAP TPC Administrator username and LDAP TPC Administrator group values during IBM Tivoli Storage Productivity Center installation. The IBM Tivoli Storage Productivity Center installation program automatically maps the LDAP TPC Administrator group value to the IBM Tivoli Storage Productivity Center Superuser role.

If you want to install IBM Tivoli Storage Productivity Center so that it authenticates its users against the operating system repository, then you do not have to create any groups or users before installation. The IBM Tivoli Storage Productivity Center installation program automatically maps the IBM Tivoli Storage Productivity Center Superuser role to the "Administrators" group (which contains the user "Administrator") on Windows, to the "system" group (which contains the user "root") on AIX, or to the "root" group (which contains the user "root") on Linux.

For more information about user roles, see “Role-to-Group Mappings” on page 291.

TCP/IP ports used by the IBM Tivoli Storage Productivity Center family

This topic lists the default ports that must be opened through the firewall when you install the IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication servers. You must disable the firewall program or open the ports to allow incoming requests to the IBM Tivoli Storage Productivity Center ports. Review these ports before installing the IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication.

TCP/IP ports used by IBM Tivoli Storage Productivity Center

See Table 12 for the IBM Tivoli Storage Productivity Center ports.

Table 12. TCP/IP ports used by IBM Tivoli Storage Productivity Center

Component	Session initiator (server perspective)	Inbound/ outbound (server perspective)	Port	Inbound/ outbound (agent perspective)	Session initiator (agent perspective)
Data server		Both	Default: 9549 (see notes 1 on page 35 and 3 on page 35)		
Device server		Both	9550 (see notes 2 on page 35 and 3 on page 35)		
Remote installation of UNIX agent	Yes	Outbound	SSH (22)	Both	No
Remote installation of Windows agent	Yes	Outbound	NetBIOS Sessions Service (139)		
Remote installation of UNIX agent	Yes	Outbound	RSH (514)	Both	No
Remote installation of UNIX agent	Yes	Outbound	REXEC (512)	Both	No
Remote installation of UNIX agent	Yes	Inbound	601		
Remote installation of all agents	Yes	Inbound	Data Server 9549		
Device server to CIM agent	Yes	Outbound	Default: HTTP: 5988 HTTPS: 5989		
VMware VI Data Source to Virtual Center or ESX server	Yes	Outbound	Default: HTTP: 80 HTTPS: 443		
Tivoli Integrated Portal server. See note 6 on page 35.	Yes	Outbound	389		

Table 12. TCP/IP ports used by IBM Tivoli Storage Productivity Center (continued)

Component	Session initiator (server perspective)	Inbound/ outbound (server perspective)	Port	Inbound/ outbound (agent perspective)	Session initiator (agent perspective)
Tivoli Integrated Portal server	Yes	Inbound	16310, 16311, 16312, 16313, 16315, 16316, 16318, 16320, 16322, 16323		
Device server. See note 6.	Yes	Outbound	389		
IBM Tivoli Storage Productivity Center for Replication server. See note 6.	Yes	Outbound	389		

1. Data server inbound server port plus another port that is 10 greater than the Data server port (9549). For example, if the Data server port is 9549, then another port would be 9549+10 or 9559.
2. If you select the default port (9550) for the Device server during installation, the following ports are also used: 9551, 9552, 9556, 9557. If you use a port other than the default port for the Device sever, make sure that port and the ports above it are available for use. For example, if N represents the open port you want to use, ensure that ports N+1, N+2, N+6, and N+7 are open as well (for example, if you select 9570, then the following ports must be open as well: 9571, 9572, 9576, 9577).
3. You must select different ports for the Data server and Device server. For example, if N is the device server port, and Z the data server port, make sure that N+1, N+2, N+6 and N+7 are not the same ports as Z and Z+10.
4. When you install the Storage Resource agent, you can specify a listener port for a remote probe. This port is used by the program to communicate back to the installation program.
5. You can find the port numbers used on your system by running the following command:

   ```
   netstat -an
   ```

6. If you intend to install IBM Tivoli Storage Productivity Center so that it authenticates users against an LDAP-compliant repository, then the embedded WebSphere Application Servers for Tivoli Integrated Portal, IBM Tivoli Storage Productivity Center Device server, and IBM Tivoli Storage Productivity Center for Replication need to be able to initiate an outbound connection to the LDAP-compliant repository, which typically listens on port 389.

Other TCP/IP ports used by IBM Tivoli Storage Productivity Center

Table 13. Other TCP/IP ports used by IBM Tivoli Storage Productivity Center

Port	Default
DB2	50000

Table 13. Other TCP/IP ports used by IBM Tivoli Storage Productivity Center (continued)

Port	Default

• SAN Volume Controller
• IBM Storwize V7000
• IBM Storwize V7000 Unified
	For native interface:
	`Inbound: none`
	`Outbound (on SVC or Storwize V7000 cluster):`
	`22 for SSH`

CIM agent for IBM TotalStorage Enterprise Storage Server (ESS)	5989

DS8000	For native interface:
	`Inbound for events: 1750`
	`Outbound (on HMC): 1750`
	`Inbound and outbound on HMC for offloading logs: 1755`

IBM XIV Storage System	For native interface:
	`Inbound: none`
	`Outbound (on XIV storage system): 7777, 7778`

IBM Tivoli Storage Productivity Center with DS8000 GUI	`8451` `8452`
IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication	162 (default SNMP listening port)
WebSphere® Application Server	9043, 9060, 9353, 5060, 5061

TCP/IP ports used by IBM Tivoli Storage Productivity Center for Replication

Table 14 shows the ports used for incoming communication. Table 15 on page 37 shows the ports used for outgoing communication.

Table 14. TCP/IP ports used by IBM Tivoli Storage Productivity Center for Replication for incoming communication

Port	Description
3080 and 3443	A Web browser typically communicates with the IBM Tivoli Storage Productivity Center for Replication GUI using HTTP ports 3080 and 3443. These ports are configurable from the IBM Tivoli Storage Productivity Center for Replication installer during installation.
5110	The IBM Tivoli Storage Productivity Center for Replication command line and GUI typically communicate with the IBM Tivoli Storage Productivity Center for Replication server using port 5110.
5120	IBM Tivoli Storage Productivity Center for Replication uses port 5120 for communication with other IBM Tivoli Storage Productivity Center for Replication servers for high-availability purposes.
Note: If you changed ports 3080 or 3443 during the installation or changed other port settings, note the values to which these settings were changed.	

Table 15. TCP/IP ports used by IBM Tivoli Storage Productivity Center for Replication for outgoing communication

Port	Description
2433	For communication with ESS/DS (direct connection).
1750	For communication with DS8000 (HMC connection).
443 and 22	For communication with SAN Volume Controller, Storwize V7000 Unified or Storwize V7000 clusters.
9930	For communication with the AIX host for the Open HyperSwap.

Note:

- If you changed the port configuration of your storage controller, the ports would be different.
- Your network configuration must allow for IBM Tivoli Storage Productivity Center for Replication to send outgoing TCP/IP packets to the storage controllers. It is possible when adding the storage controllers to IBM Tivoli Storage Productivity Center for Replication to set a specific port number for your storage controller.
- Since there are typically multiple applications running on a server, it is possible that port conflicts might arise if other applications attempt to use the same ports that IBM Tivoli Storage Productivity Center for Replication is using. You can find the port numbers used on your system by running the following command:

  ```
  netstat -an
  ```

- If firewalls are being used in your configuration, make sure that none of these ports are being blocked. Ensure that not only is the IBM 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 IBM Tivoli Storage Productivity Center for Replication server.
- If you are running Windows Server 2008, you need to configure the inbound and outbound rules for IBM Tivoli Storage Productivity Center for Replication. To create an outbound rule, launch the New Outbound Rule wizard from the Windows Firewall with Advanced Security menu.
- If you are installing or deploying Storage Resource agents on a system, you must first disable the firewall on that system before installing or deploying an agent.

User IDs and passwords for installing and using Tivoli Storage Productivity Center

Use this information to plan for the user IDs and passwords required for installing and using Tivoli Storage Productivity Center.

Use the following tables for information about the user IDs and passwords required for installation and using Tivoli Storage Productivity Center.

Table 16. Installing the Device server, Data server, or Tivoli Storage Productivity Center server

User ID and password	When created	When used
User ID and password for installing: • For Windows – Must be a local user (do not use a domain user) – Add user ID to DB2 Administrative group or assign the following user rights: - Log on as a service - Act as part of the operating system - Adjust memory quotas for a process - Create a token object - Debug programs - Replace a process level token • For UNIX or Linux - must have root user authority	Must be created before installing Device server, Data server, or Tivoli Storage Productivity Center server	When installing the servers The user ID and password must follow the Tivoli Storage Productivity Center and DB2 rules for user IDs and passwords. For DB2 rules, see **Note** below. For Tivoli Storage Productivity Center rules, see **Note** below.
Host authentication password - used to authenticate the Device server with the Data server.	Created when you install the Device server. You must provide this password at installation time.	Used to authenticate the Device server with the Data server. For the Data agent and Fabric agent, used to authenticate the agent with the Device server.
Data server account password - used for the Data server account. Creates a new service account for running the Data server.	Created when you install the Data server. You must provide this password at installation time. The default is TSRMsrv1.	Used when the operating system runs the Data server
WebSphere Application Server administrator user ID and password. This user ID has the administrative rights to administer the Device server, Tivoli Integrated Portal, and Tivoli Storage Productivity Center for Replication WebSphere instances.	Created when Tivoli Storage Productivity Center is installed. You must provide this user ID and password.	Used to administer the Device server, Tivoli Integrated Portal, and Tivoli Storage Productivity Center for Replication WebSphere instances.

Table 16. Installing the Device server, Data server, or Tivoli Storage Productivity Center server (continued)

User ID and password	When created	When used
User ID for Tivoli Storage Productivity Center for Replication server. This user ID has the administrative rights to administer the Device server, Tivoli Integrated Portal, and Tivoli Storage Productivity Center for Replication WebSphere instances.	Must be created before starting the Tivoli Storage Productivity Center for Replication server installation. You must provide this user ID and password.	Used to log into the Tivoli Storage Productivity Center for Replication server. This user ID can be the same user ID used for the WebSphere Application Server administrator user ID.
LDAP user ID and password (if you are using LDAP).	The LDAP TPC Administrator user ID is used by Tivoli Storage Productivity Center to configure the Tivoli Storage Productivity Center superuser during installation.	Used by Tivoli Storage Productivity Center to administer LDAP-related information.

Note:

For DB2 rules, see http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.admin.dbobj.doc/doc/c0007248.html.

For Tivoli Storage Productivity Center, see "Valid characters for user IDs and passwords" on page 428.

Table 17. Installing the Storage Resource agent, Data agent, or Fabric agent

User ID requirements	When created	When used
• For Windows – Add user ID to DB2 Administrative group or assign the following user rights: - Act as part of the operating system - Log on as a service • For UNIX or Linux - must have root user authority	Must be created before starting the Storage Resource agent, Data agent, or Fabric agent installation	When installing the agents.

Table 18. Agent service logon for Windows

User ID description	When created	When used
Agent service logon user ID and password .	Created when the agent is installed locally or can using existing account.	Used when running the agent locally.

There are no special user ID requirements to install the GUI or CLI.

User IDs and passwords for using Tivoli Storage Productivity Center

For Windows, you can use a Windows domain user ID to log into the Tivoli Storage Productivity Center GUI.

Table 19. DB2 administration server

User ID description	When created	When used
Used to run the DB2 administration server on your system. For DB2 rules, see **Note** below.	Created when DB2 is installed.	Used by the DB2 GUI tools to perform administration tasks.

Note:

For DB2 rules, see http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.admin.dbobj.doc/doc/c0007248.html.

For Tivoli Storage Productivity Center, see "Valid characters for user IDs and passwords" on page 428.

Work sheet for user IDs and passwords

This section provides a work sheet to help plan your user IDs and passwords during the installation of IBM Tivoli Storage Productivity Center.

Table 20. User IDs and passwords for Tivoli Storage Productivity Center installation

Item	Description	Your input
DB2 administrator user ID and password	This user ID and password are required to install Tivoli Storage Productivity Center. On Windows, this user ID must be a member of the DB2ADMNS group and Administrators group. On UNIX, the user ID must be the instance owner of the instance you want to use. This user ID and password are created when you install DB2.	

Table 20. User IDs and passwords for Tivoli Storage Productivity Center installation (continued)

Item	Description	Your input
Storage Resource agent service logon user ID and password	This user ID and password are for Windows only and is optional. This user ID and password creates a new service account for the Storage Resource agent to run under. If your enterprise has a security policy that requires long passwords, the installation can fail. This user ID and password allow you to run under a local account user name and password that meets the security policy for your enterprise. This user ID and password are created if you do not have one.	No default. If this user ID and password do not exist, they are created at installation time.
Host authentication password	This password is required for: • the Storage Resource agent to communicate with the Device server • the subordinate server when added to the master server (for enterprise-wide reporting) Specified when you install the Device server. You must specify this password when you install the Storage Resource agent.	This password is created when you install the Device serverr.
NAS filer login user ID and password	This user ID and password are for Windows only. This user ID and password are the default user name and password used by Storage Resource agents during NAS discovery.	
WebSphere administrator user ID and password	This user ID and password are required by the Device server to communicate with embedded WebSphere. Embedded WebSphere runs as a Windows Service. The Windows Service runs under the authority of this user ID and password. This user ID and password are only used when you install the Device server.	If this user ID and password do not exist, they are created at installation time.

Table 20. User IDs and passwords for Tivoli Storage Productivity Center installation (continued)

Item	Description	Your input
LDAP Tivoli Storage Productivity Center Administrator username and password and LDAP Tivoli Storage Productivity Center Administrator group	These values are relevant if you install Tivoli Storage Productivity Center so that it authenticates users against an LDAP-compliant repository. The username, password, and group must exist in your LDAP-compliant repository in the "branches" specified by the "Relative Distinguished Name for usernames" and the "Relative Distinguished Name for groups" values. The user must also be a member of the group.	During the Tivoli Storage Productivity Center installation, the LDAP Tivoli Storage Productivity Center Administrator username value is used to set the WebSphere Administrative ID value for the Tivoli Integrated Portal server, the Tivoli Storage Productivity Center Device server, and the Tivoli Storage Productivity Center for Replication server. Also during the Tivoli Storage Productivity Center installation, the LDAP Tivoli Storage Productivity Center Administrator group value is mapped to the "superuser" role in Tivoli Integrated Portal, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication to allow initial access to the applications once the installation is complete.

For information about valid user IDs and passwords, see “Valid characters for user IDs and passwords” on page 428

Supported storage systems, devices, file systems, databases, and switches

This section provides information about the storage systems, tape libraries, file system formats, and databases that IBM Tivoli Storage Productivity Center supports.

Storage systems

IBM Tivoli Storage Productivity Center supports IBM and independent software vendor systems through the native interfaces or with Storage Management Interface Specification (SMI-S) compatible interfaces. This support includes storage provisioning, as well as asset and capacity reporting.

The IBM storage systems supported are as follows:

- IBM System Storage SAN Volume Controller (with native interface)
- IBM TotalStorage Enterprise Storage Server
- IBM storage systems (DS3000, DS4000, DS5000, DS6000 series)
- DS8000 series (with native interface)
- FAStT
- IBM XIV Storage System (with native interface)
- IBM Storwize V7000 (with native interface)

- SMI-S certified storage systems. For information about SMI-S certified storage systems, see http://www.snia.org/ctp/conformingproviders#10provider.

Note:

- For the Storage Optimizer, the DS4000 supports up to four fiber channel host connections and up to 112 SATA drives.
- SMI-S agents (CIM agents) are not required for DS8000, XIV, SAN Volume Controller or Storwize V7000 storage systems.
- A specific SAN Volume Controller should not be managed by more than one Tivoli Storage Productivity Center server at the same time.

For the most current information about storage systems supported, firmware and provider levels, see the Tivoli Storage Productivity Center support Web site at http://www-01.ibm.com/support/docview.wss?rs=40&uid=swg21386446 &loc=en_US&cs=utf-8&lang=en.

Switches

IBM Tivoli Storage Productivity Center supports IBM and independent software vendor switches that are Storage Management Interface Specification (SMI-S) compatible.

Tivoli Storage Productivity Center supports the following switches:

- Brocade

  Note: The Brocade 8000 switch is supported in toleration mode only. Tivoli Storage Productivity Center does not support the Converged Enhanced Ethernet (CEE) or Fibre Channel over Ethernet (FCoE) connectivity functions for the Brocade 8000 switch. Only WWN zoning is supported with the internal FCoE ports. The domain-port zoning for Brocade 8000 switch is limited to the eight FC ports.Tivoli Storage Productivity Center displays both the FCoE and FC ports in the switch port lists.
- Cisco

  Note: The Cisco Nexus 5000 switch is supported in toleration mode only. Tivoli Storage Productivity Center does not support the Converged Enhanced Ethernet (CEE) or Fibre Channel over Ethernet (FCoE) connectivity functions for the Cisco Nexus 5000 switch. Tivoli Storage Productivity Center also does not support domain-port zoning for Cisco Nexus switches. Tivoli Storage Productivity Center displays the FCoE and FC ports in the switch port lists.
- IBM
- McDATA Intrepid
- McDATA Sphereon
- QLogic

For the most current information about switches and directors supported, see the Tivoli Storage Productivity Center support Web site at http://www-01.ibm.com/support/docview.wss?rs=40&uid=swg21386446&loc=en_US&cs=utf-8&lang=en.

Tape libraries

This section provides information about the tape libraries that IBM Tivoli Storage Productivity Center supports.

For the most current information about tape libraries supported, see the Tivoli Storage Productivity Center support Web site at http://www-01.ibm.com/support/docview.wss?rs=40&uid=swg21386446&loc=en_US&cs=utf-8&lang=en.

Tivoli Storage Productivity Center supports the following tape libraries:

- IBM System Storage TS3500 Tape Library (formerly Tivoli Storage Productivity Center 3584 Tape Library) - The TS3500 Tape Library is designed to provide a highly scalable, automated tape library for mainframe and open systems backup and archive in midrange to enterprise environments.
- IBM System Storage TS3310 Tape Library is a modular, scalable tape library designed to grow as your needs grow.

Tivoli Storage Productivity Center supports IBM and non-IBM tape libraries through the SMI-S (CIM) agent. This support assumes that the tape library supports the SMI-S 1.1. profile for tape libraries.

Note:

- The media changers displayed by Tivoli Storage Productivity Center for the TS3500 tape libraries are logical partitions of the given library, and not physical accessors.
- Depending on the size of the tape library and network latency between the SMI-S agent host and the tape library, probes of the tape library might take a long time. The probes might also fail because of time outs.
- The SMI-S agent does not display an exception to the SMI-S client such as Tivoli Storage Productivity Center when communication between the SMI-S agent and the tape library fails. Instead, empty result sets are returned. This limits Tivoli Storage Productivity Center's capabilities in detecting connection failures relating to tape libraries.
- When probing tape libraries that are registered with the same IBM SMI-S agent for tape, do not probe more than two or three tape libraries within the same probe job. The increased load on the agent would increase the likelihood of timeouts. Instead, spread the libraries across multiple probe jobs with different start times.
- After the tape libraries are registered, and then a change is made to the IBM SMI-S Agent for Tape, a condition can occur where not all the tape cartridges are returned to the CIM client. In this case, the CIM client is Tivoli Storage Productivity Center. To resolve this situation, restart the IBM SMI-S Agent for Tape. For information about how to restart the agent, refer to the documentation for the SMI-S Agent for instructions.

  The same issue might occur if one out of a set of libraries registered with the same agent is unavailable. For example, there might be a network problem. To work around this problem, unregister the affected library from the agent (or fix the communication problem).

Databases for monitoring

This section lists the databases that the Data Manager can monitor.

The Data Manager can monitor these databases:

- DB2 9.1 with fix pack 2 or later (32-bit and 64-bit)

  Note: Windows 2008 requires DB2 9.1 with fix pack 5 or later.
- DB2 9.5 (32-bit and 64-bit)
- DB2 9.7 (32-bit and 64-bit)

- Microsoft SQL Server 8
- Microsoft SQL Server 2005 (32-bit and 64-bit)

  Note: There are some configuration steps required before you can monitor the Microsoft SQL Server 2005 database. For information about configuration, see "Configuring Microsoft SQL Server 2005 or 2008" on page 297.
- Microsoft SQL Server 2008 (32-bit and 64-bit)

  Note: There are some configuration steps required before you can monitor the Microsoft SQL Server 2008 database. For information about configuration, see "Configuring Microsoft SQL Server 2005 or 2008" on page 297.
- Oracle 10g (32-bit)
- Oracle 11g release 1 and 2 (32-bit and 64-bit)
- Sybase Adaptive Server version 15 (32-bit and 64-bit)

Note: Databases installed on clustered servers or clustered database servers (for example, MSCS, PowerHA for AIX (formerly HACMP), Veritas clusters, Oracle Real Application Cluster (RAC) environment) are not supported for monitoring.

Use the **Data Manager for Databases** node in the Tivoli Storage Productivity Center user interface to help you monitor the databases within RDBMS instances. The Storage Resource agents are used to monitor the RDBMS instances.

You can monitor the storage assets associated with Oracle, Sybase, SQL Server, and DB2 throughout your enterprise, watching for predefined events to occur and alerting you to potential situations before they occur.

For example, it can notify you when an Oracle tablespace is reaching a critical shortage of free space or when a Sybase table is dropped. By alerting you to these and other issues related to your stored data for the databases within your environment, it enables you to prevent unnecessary system and application downtime.

Before you can use the Storage Resource agents to manage the storage for your instances, you must do the following steps:

1. Assign Data Manager and Data Manager for Databases licenses to the agents that are monitoring RDBMS instances.
2. Register the instances on the machines that contain licensed agents.

You can do these steps by expanding **Administrative Services** > **Configuration** > **License Keys** in the navigation tree. For more information about this topic, see the online help or "License Keys" on page 293.

For information about the operating systems supported for the agent, see "Software requirements for operating systems" on page 5.

See the online help for more information about the **Data Manager for Databases** node.

File systems

This section lists the file systems that are supported for monitoring and reporting by IBM Tivoli Storage Productivity Center.

The Data Server supports monitoring and reporting of the following file systems:

- AIX JFS

- Data ONTAP V7 for Network Appliance, including flexible volumes (FlexVol). Flexvol enables you to create multiple flexible volumes on a large pool of disks.
- EXT2 and EXT3
- FAT, FAT32
- GPFS 3.2
- HP_HFS
- NTFS4, NTFS5
- NW_FAT, NW_NSS (NetWare)
- REISERFS
- SANFS
- TMPFS
- UFS
- VMFS
- Veritas Files System (VxFS) r4.5 and 5.x on the following operating systems:
  - AIX
  - HP-UX
  - Linux
  - Oracle Solaris

  Note: IBM Tivoli Storage Productivity Center supports file systems and volumes created through Veritas File System (VxFS) and Veritas Volume Manager (VxVM) release 4.x and 5.x, as long as those file systems and volumes conform to file systems and volumes created with version 3.x of VxFS and VxVM. IBM Tivoli Storage Productivity Center cannot work with volumes and files systems with new features that were introduced with VxFS and VxVM release 4.x and 5.x.
- WAFL
- Tivoli Storage Productivity Center SAN File System

Networked file systems

This section describes support for IBM General Parallel File System.

When using the Data server, you can use the monitoring and reporting of the IBM General Parallel File System (GPFS) Version 3.2 on AIX.

- The IBM Tivoli Storage Productivity Center Storage Resource agent must be installed on one node within a GPFS *nodeset*. A nodeset is a collection of computers that see the same file system. If multiple agents are installed per GPFS nodeset, the first agent that sees the file system owns the file system. You cannot change the owning (or scanning) agent. If the agent is deleted, another agent takes ownership.

  Because only one agent owns the file system and a file system cannot be scanned by more than one agent, there is no benefit to having more than one agent.

  If more than one agent is installed, the last agent that runs a probe job takes ownership of the volume group. If a GPFS file system exists in that volume group and the agent that owns the volume group is not the agent that owns the file system, the file system information is not displayed in the volume group. The volume group is part of the asset tree. However, when the agent that owns the file system probes again, the data is corrected.

  There must be physical access to the GPFS disk for IBM Tivoli Storage Productivity Center to gather hardware disk information. File system information is gathered.

Volume managers

This section lists the volume managers supported for monitoring by IBM Tivoli Storage Productivity Center.

The Data Server supports the monitoring of the following volume managers:

- Veritas Volume Manager is supported on the following operating systems:
  - AIX
  - HP-UX
  - Linux
  - Oracle Solaris
- AIX Logical Volume Manager (LVM)
- HP-UX Logical Volume Manager
- Linux Logical Volume Manager (LVM2)

Note: Tivoli Storage Productivity Center supports asset reporting on Linux LVM2 volumes only when the Linux LVM2 physical volumes of a disk all belong to the same LVM volume group. You can assign LVM physical volumes of a disk to different LVM volume groups, but they will not be reported accurately in asset reports.

Using these volume managers, you can create groups of logical volumes and disks. You can generate various reports for these disk and volume groups.

Planning for multipath subsystem device drivers

This topic lists the multipath subsystem device drivers (SDD) that you can use with IBM Tivoli Storage Productivity Center.

The subsystem device driver (SDD) is a software solution for multiple configuration environments in supported storage devices. It resides in a host system with the native disk-device driver and provides the following functions:

- Enhanced data availability
- Dynamic input/output (I/O) load balancing across multiple paths
- Automatic path failover protection
- You can download licensed machine code at the same time that applications are running

For the most current support for multipath subsystem device drivers, see the following Web site: http://www-01.ibm.com/support/docview.wss?rs=40 &uid=swg21386446&loc=en_US&cs=utf-8&lang=en. Select the appropriate platform for Agents, Servers and GUI.

With Tivoli Storage Productivity Center, you can use the subsystem device drivers (SDD) shown in Table 21 on page 48.

Table 21. Multipath subsystem device drivers supported by Tivoli Storage Productivity Center

Multipath product	Platform	Supported features in Tivoli Storage Productivity Center
SDD	AIX	Full support
	Red Hat 4	Full support
	SUSE Linux Enterprise Server 9	Full support
	Oracle Solaris	Full support
	Windows	Full support
SDDDSM	Windows	Full support
SDDPCM	AIX	Full support
DM_Multipath	Red Hat Enterprise Linux 5	Full support
	Red Hat Enterprise Linux 6	Full support
	SUSE Linux Enterprise Server 10	Full support
	SUSE Linux Enterprise Server 11	Full support
EMC PowerPath	AIX	Probe
	Linux	Probe
	Windows	Probe
HP-UX Native Multipath	HP-UX	Planner

Full support means that Tivoli Storage Productivity Center can:

1. Report on the disks accessed using the multipath subsystem device driver (probe).
2. Display multiple paths in the Data Manager asset tree (Data Manager Asset Tree).
3. Report on the multiple paths in the Data Path Explorer.
4. Configure paths as part of the SAN Configuration Planner.

Note:

1. The AIX SDD cannot coexist with SDDPCM on the same system.
2. The Linux SDD driver is no longer available. Starting with Red Hat Enterprise Linux Version 5 and SUSE Linux Enterprise Server Version 10, only the DM_Multipath is available. With DM_Multipath, you can only use the round-robin policy for the SAN Planner.
3. 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 SAN Planner.

For more information about multipath subsystem device drivers, see http://www-01.ibm.com/support/docview.wss?rs=540&context=ST52G7&dc=DA400&uid=ssg1S7001350&loc=en_US&cs=utf-8&lang=en.

For more information about how to install, configure, and use the subsystem device drivers, see http://www.ibm.com/support/docview.wss?rs=540&context=ST52G7&uid=ssg1S7000303.

Upgrading subsystem device drivers

SDD drivers cannot coexist on the same host with the SDDPCM, SDDDSM, or DM_Multiplath drivers. You must upgrade from the existing SDD drivers to the SDDPCM or SDDDSM driver. Part of the upgrade process is to unconfigure and remove all SDD vpath devices. After the upgrade and configuration of the devices, the device names might differ from the previous names. Each device is detected by Tivoli Storage Productivity Center as a new device. For information about how to upgrade SDD, see http://www.ibm.com/support/docview.wss?rs=540 &context=ST52G7&uid=ssg1S7000303.

EMC PowerPath Multipathing

EMC PowerPath Multipathing supports a wide range of servers including cluster servers connected to EMC storage systems. It tunes your storage area network and selects alternate paths for your data if necessary. It also integrates multiple path I/O capabilities, automatic load balancing, and path failover functions. For more information about EMC PowerPath Multipathing, see http://www.emc.com. Search for **EMC PowerPath**.

Disks provided by the EMC PowerPath driver are detected by the Storage Resource agents. The disks are visible in the topology viewer and in Data Manager asset reports. The correlation of EMC PowerPath provided hdisks to EMC storage systems is supported and the relation is visible in the topology viewer. Multipathing information is not available for those disks. The data reports are accurate because they do not double-count capacities. The Data Path Explorer does not show multipathing. The SAN Planner does not support configuring multipathing on hosts using EMC PowerPath Multipathing.

Planning for Storage Resource agents

Storage Resource agents perform the functions of the Data agent and Fabric agent. Storage Resource agents are designed to be more lightweight and easier to install or deploy than the Data agents and Fabric agents.

The Storage Resource agents perform the functions of the Data agents by collecting information from the system they are installed on. The agents collect asset information, file and file system attributes, and any other information needed from the computer system. 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. This topic provides planning information for installing the Storage Resource agents.

The first time you deploy a Storage Resource agent to a computer system, the information collected by that agent is added to the logs on that computer system. That information also displays in the Tivoli Storage Productivity Center GUI after the next default system scan completes.

The Storage Resource agents perform the functions of the Fabric agents by collecting information from the system they are installed on. The agents are installed on systems that have fiber connectivity (through HBAs) into the SAN fabrics that you want to manage and monitor.

Storage Resource agents use scanners to collect information. The scanners communicate through the HBA to collect fabric topology information, port state information, and zoning information. They also can identify other SAN-attached devices (if they are in the same zone).

You deploy or install the Storage Resource agent through the user interface rather than an installation wizard. You must have administrative privileges to deploy the Storage Resource agent.

To deploy a Storage Resource agent, expand **Administrative Services > Configuration**. Right-click **Storage Resource Agent Deployments**. Click **Create Storage Resource Agent Deployments**. In the content pane, the "Create Storage Resource Agent Deployments" window opens. From this window, you can enter host names from a Microsoft directory, enter host names manually, or import a host list file. You can schedule a job to add the Storage Resource agents at a time that is convenient for you. You can also specify the type of alerts to generate if a job fails. After you enter information for the Storage Resource agent, you can optionally validate the connection to the Data server. This step helps eliminate most of the possible failure cases after you have submitted the job for Storage Resource agent deployment.

For more information about Storage Resource agents, see "Deployment considerations for Storage Resource agents" on page 312.

Storage Resource agent protocol support

When installing the Storage Resource agent, Tivoli Storage Productivity Center uses specific protocols for connectivity between the server and agent.

When installing the Storage Resource agent, Tivoli Storage Productivity Center uses the following protocol for connectivity between the server and agent (listed in order):

1. Secure Shell protocol (SSH).
2. Windows server message block protocol (SMB protocol).
3. Remote execution protocol (REXEC).
4. Remote shell protocol (RSH).

At run time, the connectivity used between the server and agent is dependent on the type of service that is running: On-Demand service (non-daemon service) or run as a service (daemon service).

On-Demand service (non-daemon service)
: In this case, connectivity between the server and agent is established using the same protocols as for installation of the agent: SSH, SMB, REXEC, or RSH.

run as a service (daemon service)
: In this case, connectivity between the server and agent is established using the secured socket connection. The server and agent have their respective certificates and no additional information is needed besides the certificates and the security provided by the SSH protocol.

Note: Running the Storage Resource agent in daemon mode requires only one socket connection to be open and thus simplifies any firewall rules that might be in place. However, the daemon mode always consumes resources (although a very small amount) when idle. In stand-alone mode, the Storage Resource agent only

consumes resources when actively working for the Tivoli Storage Productivity Center server. Because the stand-alone mode uses the RXA protocol, additional firewall ports and security configuration are required.

The information required for these protocols are as follows:

SSH protocol
There are two cases where you can use the SSH protocol:

- You supply the user ID and password to connect to the server using the SSH protocol. This user ID must have administrative privileges.
- You supply a user ID, certificate, and passphrase. You create the certificate and use that certificate when you connect to the agent. This certificate must be accessible from the server.

  For information about how to create a certificate for SSH protocol connectivity, see "Creating a certificate for SSH protocol" on page 317. This information uses Cygwin for the Storage Resource agent.

Windows SMB protocol
You supply the user ID and password. The user ID must have administrative privileges. You must also enable **File & Printer Sharing for Microsoft Windows** under **Network Properties** for the connected network adapter. Make sure that the Server service is running on the system.

REXEC protocol
You supply the user ID and password. The user ID must have administrative privileges. The system must be enabled for remote execution of commands.

RSH protocol
You supply the user ID with administrative privileges. The system must be enabled for the user to be able to run commands through the remote shell.

To enable RSH, edit the .rhosts file in the login directory of the user. The rhosts file provides access for the user to connect from a remote system.

If you elect to use a Microsoft directory, you can also get a list of domain computers. You are required to enter the domain controller, user ID, password, and certificate location. After the list is displayed, you can select a list of computers on which to deploy the agent.

After the agent has been deployed, a probe job is automatically run.

The agent deployed on a Windows system has its registry information in the Windows registry. For UNIX or Linux, the registry information is stored in the following file: `/etc/Tivoli/TSRM/registryNA` .

For more information about setting up the environment for installing Storage Resource agents, see "Deployment considerations for Storage Resource agents" on page 312.

Planning for Internet Protocol Version 6

IBM Tivoli Storage Productivity Center supports Internet Protocol Version 6 (IPv6) for communication between its components. The key IPv6 enhancement is the expansion of IP address spaces from 32 bits (up to 15 characters in length) to 128 bits (up to 45 characters in length).

Overview

You can install and run Tivoli Storage Productivity Center on computers that are enabled for IPv4, IPv6, or dual stack. *Dual stack* indicates that a machine has both the IPv4 and IPv6 stacks enabled and both addresses configured. The following table shows the internal components of Tivoli Storage Productivity Center that can communicate with each other using IPv6.

Table 22. IPv6 communication for internal Tivoli Storage Productivity Center components

Components in this column can communicate over IPv6 with the components in the second column	**Tivoli Storage Productivity Center components**
User interface	Data server
Data server	Device server database repository
Device server	database repository tpctool (command-line tool) SRMCP
database repository	Data server Device server Repocopy tool
tpctool (command-line tool)	Device server
SRMCP	Device server
Repocopy tool	database repository

Tivoli Storage Productivity Center can communicate with the following external components over IPv6:

- SMI-S CIM agents
- SLP directory agents
- Out of band SNMP agents
- Storage devices (for example, the IBM XIV Storage System does not support IPv6)
- SMTP server

Note: These external components must be IPv6 enabled to communicate with Tivoli Storage Productivity Center using the IPv6 protocol.

Not all components and their related functions of Tivoli Storage Productivity Center are enabled for IPv6. Any functions that are not enabled for IPv6 are unavailable through the user interface when you install Tivoli Storage Productivity Center on an IPv6-only system. See "IPv6 considerations in IBM Tivoli Storage Productivity Center" on page 54 for information about the functions that are currently not supported under IPv6 communication.

Installation options

This section provides an overview of the installation options for installing and upgrading Tivoli Storage Productivity Center depending on the IP protocol enabled in your environment:

Installing Tivoli Storage Productivity Center on an IPv6-only machine
You can install Tivoli Storage Productivity Center on a system configured only for IPv6. This installation of the product can communicate with other Tivoli Storage Productivity Center servers (for roll up probes and reporting), SMIS CIM agents, out of band SNMP agents, and storage devices operating in IPv6 or dual stack mode. The product cannot communicate with storage entities located on IPv4-only systems.

See "Installation scenarios for Tivoli Storage Productivity Center" on page 56 for more information about this installation option. See "IPv6 considerations in IBM Tivoli Storage Productivity Center" on page 54 for information about the limitations of installing the product on IPv6-only systems.

Installing Tivoli Storage Productivity Center on a dual stack machine
You can install Tivoli Storage Productivity Center on a dual stack system where both IPv6 and IP4 addresses are enabled. This installation of the product can communicate with Storage Resource agents, other Tivoli Storage Productivity Center servers (for roll up probes and reporting), SMI-S CIM agents, out of band SNMP agents, and storage devices operating in IPv6-only, dual stack mode, or IPv4-only modes.

See "Installation scenarios for Tivoli Storage Productivity Center" on page 56 for more information about this installation option.

Installing Tivoli Storage Productivity Center on an IPv4-only machine
You can install Tivoli Storage Productivity Center in an IPv4-only environment. However, IPv6 communication for the Data and Device servers are disabled and you cannot enter IPv6 addresses in the user interface.

Upgrading Tivoli Storage Productivity Center for use with IPv6
You must upgrade to Tivoli Storage Productivity Center v3.3.2 or later if you want to enable the product to use the IPv6 protocol. You can upgrade an earlier version of Tivoli Storage Productivity Center for use on a machine that supports both IPv4 and IPv6 (dual stack). You cannot upgrade an earlier version of Tivoli Storage Productivity Center for use on a machine that is IPv6 only. If you want to use Tivoli Storage Productivity Center on an IPv6-only machine, you must perform a new install of the product on that machine.

See the Tivoli Storage Productivity Center Information Center for more information about how to install and upgrade the product. Click **Tivoli Storage Productivity Center > Installing** and **Upgrading and migrating**.

Using IPv6 addresses in the user interface

Support for IPv6 enables you to enter IPv6 addresses anywhere in the installation process and user interface that supports the entering of IP addresses. For example, you can enter and display 128–bit IP address spaces on the following panels in the user interface:

Logon panel
Enter the host name (or IPv6 address) and port number of an Tivoli Storage Productivity Center server to which you want to connect. See the Logon panel for more information about logging in to Tivoli Storage Productivity Center and using an IPv6 address for the server name.

Administrative Services > Data Sources > CIMOM agents
Specify IPv4 or IPv6 addresses for the CIMOMs that are located on an IPv6 or dual stack machine.

Administrative Services > Discovery > CIMOM Discovery
Specify IPv6 addresses when you run a discovery with SLP DA running on an IPv6-only or dual stack machine.

Administrative Services > Data Sources > Out of Band Fabric Agents
Specify IPv6 addresses for the Out of Band Fabric agents that are located on an IPv6 or dual stack machine.

Administrative Services > Data Sources > TPC Servers
Specify the names or IPv6 addresses of other Tivoli Storage Productivity Center servers that you want to add as subordinate servers for roll up probes and reports.

Administrative Services > Discovery > Switches and Subsystems (IP Scan)
Specify a range of IPv6 addresses for an SNMP discovery on an IPv6-only or dual stack machine against IPv6-enabled target machines.

Administrative Services > Configuration > Alert Disposition
If the Tivoli Storage Productivity Center servers are running on an IPv6-only or dual stack machine enter the following information:

- Specify an IPv6 address for the target machine of SNMP traps that are generated as a result of triggered alerts.
- Specify an IPv6 address for a mail server that receives email notifications that are generated as a result of triggered alerts.

IBM Tivoli Storage Productivity Center > Configuration Utility > Element Manager tab
Use an IPv4 or IPv6 address when you specify the URL for an external tool that you want to start from Tivoli Storage Productivity Center.

IPv6 considerations in IBM Tivoli Storage Productivity Center

IBM Tivoli Storage Productivity Center provides limited support of some functions when using Internet Protocol Version 6 (IPv6) communication.

Keep in mind the following considerations when installing the product on computers configured for IPv6 only:

Using DB2 Database V9.1 fix pack 2 or later for the Linux and UNIX database repository
If you install Tivoli Storage Productivity Center or its database repository on an IPv6-only computer, you must use *DB2 Database for Linux, UNIX, and Windows V9.1 fix pack 2 or later* for the database repository. See the Tivoli Storage Productivity Center Information Center for information about how to install and configure DB2 to use the IPv6 protocol. You cannot configure a computer that is running Windows 2003 Server Enterprise Edition for IPv6-only.

Starting the Tivoli Storage Productivity Center user interface using Java Web Start
You cannot use Java Web Start to start the Tivoli Storage Productivity Center user interface on a Windows system if both of the following conditions are true:

- An IPv6 address or a host name that resolves to an IPv6 address is specified as the target host in the Java Web Start URL.
- Java Runtime Environment (JRE) is installed on the same system as the Tivoli Storage Productivity Center user interface.

Using IPv6 communication between master and subordinate servers (for roll up reporting)
If the master server is located on an IPv6-only computer, it can communicate with existing subordinate servers under the following conditions: the subordinate servers are upgraded to Tivoli Storage Productivity Center v3.3.2 or later and the IPv6 protocol is enabled on the computers where the servers are located.

Defining Tivoli Enterprise Console® events as triggered actions
You cannot use IBM Tivoli Enterprise Console events as triggered actions when defining alerts for Tivoli Storage Productivity Center jobs.

Selecting files for archive and backup
You cannot select files for Tivoli Storage Manager archive and delete jobs. For example, you cannot use the alerting function of a constraint to automatically invoke a Tivoli Storage Manager archive or backup job to run against the files that violate that constraint.

Installing Tivoli Storage Productivity Center
You can install Tivoli Storage Productivity Center on both IPv6 and IPv4 computers.

- You can enter IP addresses of computers in IPv6 format. The preferred IPv6 address representation is written as eight groups of four hexadecimal digits: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx, where each x is a hexadecimal digit representing 4 bits. For example: 2001:0DB8:0000:1234:0000:0000:5678:ABCD. You can also specify IPv6 addresses using shortened formats that omit leading zeros: 2001:DB8:0:1234:0:0:5678:ABCD or use double colons: 2001:DB8:0000:1234::5678:ABCD in place of a series of zeros. Only one "::" is allowed per IPv6 address. You can use IPv6 addresses if the computer where Tivoli Storage Productivity Center is installed is IPv6 or dual stack (IPv4 and IPv6) enabled.
- The installation program uses IPv6 protocol for configuration and communication tasks that are related to the Device server, Data server, graphical user interface (GUI), and command-line interface (CLI) on IPv6-only computers.
- The installation program uses IPv4 protocol for configuration and communication tasks that are related to the Device server, Data server, GUI, and CLI on systems that are enabled for these systems:
  - Both IPv4 and IPv6 (dual stack)
  - IPv4-only systems
- You can upgrade an earlier version of Tivoli Storage Productivity Center for use on a computer that supports both IPv4 and IPv6 (dual stack). You cannot upgrade an earlier version of Tivoli Storage Productivity Center for use on a computer that is IPv6 only. If you want to use Tivoli Storage Productivity Center on an IPv6-only computer, you must perform a new installation of the product on that computer.

Gathering data from XIV system

IBM XIV Storage System does not support IPv6.

Installation scenarios for Tivoli Storage Productivity Center

You can install IBM Tivoli Storage Productivity Center on computers configured for IPv4, IPv6, or dual stack protocol. Depending on the protocol enabled on the target computer, you can perform different tasks during and after the installation of the product.

IPv6-only installation scenario

These scenarios illustrate some of the general steps you might perform when installing the product on computers with different protocols.

In this scenario, install Tivoli Storage Productivity Center on a computer configured for IPv6-only. See IPv6 considerations in Tivoli Storage Productivity Center for information about what functions are not available when installing the product on a computer configured for IPv6-only.

1. Use either of the following methods when installing Tivoli Storage Productivity Center components:
   - **Method 1:** Install the graphical user interface (GUI), CLI, Data server, Device server, and database repository components on a local, IPv6-only configured computer.
   - **Method 2:** Install the database repository on a remote IPv6-only computer, and the GUI, CLI, Data server, and Device server on a local, IPv6-only computer.

   If you want to connect to a remote database repository using the IPv6 protocol, you must use DB2 Database V9.1 Fix Pack 2 or later for the database repository.
2. Run the user interface from an IPv6-only or dual stack configured computer. You cannot run the user interface from a computer configured for IPv4-only.
3. Configure communication to other Tivoli Storage Productivity Center servers (for roll up probes and reporting), Storage Resource agents, SMI-S CIM agents, out-of-band Fabric agents, and storage devices operating in IPv6 or dual stack mode. You cannot configure communication from Tivoli Storage Productivity Center that is installed on an IPv6-only system to storage entities on IPv4-only systems.
4. Launch element managers on IPv6-only computers from the Topology viewer, the Configuration Utility, and other points in the Tivoli Storage Productivity Center GUI.

IPv6 and IPv4 (dual stack) installation scenario

In this scenario, install Tivoli Storage Productivity Center on a computer where both IPv4 and IPv6 are configured:

1. Use any of the following methods when installing Tivoli Storage Productivity Center components:
   - **Method 1:** Install the GUI, CLI, Data server, Device server, and database repository components on a local, dual stack computer.
   - **Method 2:** Install the database repository on a remote IPv4-only computer, and the GUI, CLI, Data server, and Device server on a local, dual stack computer.
   - **Method 3:** Install the database repository on a remote IPv6-only computer, and the GUI, CLI, Data server, and Device server on a local, dual stack

computer. You must use the *DB2 Database V9.1 fix pack 2 or later for the Linux or UNIX* for the database repository. This database is IPv6 enabled and is used for the database repository.

2. Install Storage Resource agents remotely on computers that are running IPv4, IPv6, or dual stack mode.
3. Run the user interface from an IPv4-only, IPv6-only, or dual stack computer.
4. Configure communication to other Tivoli Storage Productivity Center servers (for roll up probes and reporting), Storage Resource agents, SMI-S CIM agents, out-of-band Fabric agents, and storage devices operating in IPv4, IPv6, or dual stack modes.
5. Start element managers on IPv6-only, IPv4-only, and dual stack computers from the Topology viewer, the Configuration Utility, and other points in the Tivoli Storage Productivity Center GUI.

Planning for native interface support

Tivoli Storage Productivity Center now communicates with DS8000, XIV system, SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 storage systems through their native interfaces. You no longer need to install and maintain CIM agents to collect data from these storage systems.

Support for the native interfaces

For the DS8000, Tivoli Storage Productivity Center communication with the storage system is through the ESSNI interface. Through this interface, Tivoli Storage Productivity Center collects device information, performance information, and can detect changes and perform provisioning operations. This support simplifies the Tivoli Storage Productivity Center configuration and the launch in context configuration required to manage the storage system. Tivoli Storage Productivity Center also supports dual HMC environments.

For the SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000, Tivoli Storage Productivity Center uses the CLI interface to communicate with the storage system. Through this interface, Tivoli Storage Productivity Center collects device information, performance information, and can detect changes and perform provisioning operations. This support simplifies the Tivoli Storage Productivity Center configuration and the launch in context configuration required to manage the storage system.

For the XIV system, Tivoli Storage Productivity Center uses the XIV system XML API interface to communicate with the storage system. The Tivoli Storage Productivity Center server communicates directly with the XIV system using XML over an SSL socket. Through this interface, Tivoli Storage Productivity Center collects device information, performance information, and can detect changes and perform provisioning operations. This support simplifies the Tivoli Storage Productivity Center configuration and the launch in context configuration required to manage the storage system.

Use the Configure Devices wizard to define native interface connection information for these storage systems as described in "Configure Devices" on page 266. Use one of the following methods to access the wizard:

- In the navigation tree, go to **Administrative Services** > **Data Sources** > **Storage Subsystems** and click **Add** on the Storage Subsystems panel.
- In the navigation tree, go to **Disk Manager** > **Storage Subsystems** and click **Add Storage Subsystem** on the Storage Subsystems page.

- Click the **Configure Devices** wizard icon in the toolbar.

Upgrading Tivoli Storage Productivity Center

If you are upgrading from a previous version of Tivoli Storage Productivity Center and have a DS8000, XIV system, or SAN Volume Controller storage system that was monitored in the previous version, you must migrate your storage system credentials when upgrading. Tivoli Storage Productivity Center now communicates with these storage systems using their native interfaces instead of the CIM agents.

You can migrate your storage system credentials in one of the following ways:

- Use the migration tool that is shipped with Tivoli Storage Productivity Center
- Use the Tivoli Storage Productivity Center GUI to migrate the credentials (see "Update Storage Subsystem Credentials" on page 328)
- Run the Tivoli Storage Productivity Center upgrade installation program and specify that you want to use the migration tool. Run the migration tool during the Tivoli Storage Productivity Center upgrade procedure.

If Tivoli Storage Productivity Center can access the DS8000 using the existing credentials, you can continue to use the existing credentials. For XIV system and SAN Volume Controller storage systems, you must manually update the credentials.

The storage system credential migration applies to all DS8000, XIV system, and SAN Volume Controller systems.

If you have run a CIMOM discovery job for a storage system but have not run a probe job for that system, then upgrade Tivoli Storage Productivity Center, the IP address does not display in the GUI. You must manually enter the IP address for that storage system.

Planning for LDAP support

This section provides information about planning for LDAP support.

The Lightweight Directory Access Protocol (LDAP) is an application protocol for querying and modifying directory services running over TCP/IP. A directory is a set of objects with similar attributes organized in a logical and hierarchical manner. An LDAP directory tree often reflects various political, geographic, and organizational boundaries, depending on the model chosen. The directory might contain entries representing people, organizational units, printers, documents, groups of people, or anything else that represents a given tree entry (or multiple entries).

If you want to install IBM Tivoli Storage Productivity Center so that it uses an LDAP-compliant directory for user authentication and group authorization, then you need to know the following information regarding your directory configuration:

- What is the fully qualified domain name of the system that is running your LDAP-compliant directory?
- What network port does your LDAP-compliant directory listen to for communications?

- Does your LDAP-compliant directory allow for anonymous binding when performing user and group searches? If the answer to this question is no, then what are the Bind Distinguished Name and Password for your LDAP-compliant directory?

  Note: If you want the ability to create or modify LDAP users and groups from the IBM Tivoli Integrated Portal administrative panel, then the Bind Distinguished Name and Password are required.
- What is the Relative Distinguished Name for the IBM Tivoli Storage Productivity Center users in your LDAP-compliant directory? All IBM Tivoli Storage Productivity Center users must have the same Relative Distinguished Name (all IBM Tivoli Storage Productivity Center users must be in the same branch of the LDAP directory tree).
- What is the Relative Distinguished Name for the IBM Tivoli Storage Productivity Center groups in your LDAP-compliant directory? All IBM Tivoli Storage Productivity Center groups must have the same Relative Distinguished Name (all IBM Tivoli Storage Productivity Center groups must be in the same branch of the LDAP directory tree. IBM Tivoli Storage Productivity Center groups do not have to be in the same branch of the LDAP directory tree as the IBM Tivoli Storage Productivity Center users.
- What is the Naming Attribute for the IBM Tivoli Storage Productivity Center users in your LDAP-compliant directory? All IBM Tivoli Storage Productivity Center users must use the same naming attribute.
- What is the Naming Attribute for the IBM Tivoli Storage Productivity Center groups in your LDAP-compliant directory? All IBM Tivoli Storage Productivity Center groups must use the same naming attribute

If you want to install IBM Tivoli Storage Productivity Center so that it uses an LDAP-compliant directory for user authentication and group authorization, then you need to be aware of the following restrictions:

- If you want to install IBM Tivoli Storage Productivity Center on a Windows platform and have IBM Tivoli Storage Productivity Center authenticate its users against an LDAP-based repository, then the LDAP TPC Administrator username value must *not* contain a space in it. For information about this condition, see the WebSphere Application Server APAR PK77578. The WASService.exe fails when the WAS Admin ID contains a space or blank character; `stopserver.bat` does not stop the server1 service.
- If you want to install IBM Tivoli Storage Productivity Center so that it authenticates its users against an LDAP-based repository, you must first ensure that *all* of the LDAP attributes (for example, uid, cn, ou, o, and c) are configured to be not case sensitive within the LDAP-based repository. These attributes are used in the user and group Distinguished Names or as the user or group naming attribute.

Planning for the single sign-on feature

Single sign-on is an authentication process that enables you to enter one user ID and password to access multiple applications. For example, you can access IBM Tivoli Integrated Portal and then access Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication from Tivoli Integrated Portal using a single user ID and password. Single sign-on integrates with the launch and launch in context features to enable you to move smoothly from one application to another application.

How single sign-on works

A single sign-on environment requires a centralized authentication repository that is accessed by all applications within the environment. The user ID and other authentication information are passed between applications using Lightweight Third-Party Authentication (LTPA) tokens. 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 information. The authenticating service uses the corresponding private keys to decrypt the user information for authentication. As an additional security mechanism, LTPA tokens expire after a determined amount of time. By default, the tokens expire after 24 hours. You can change the LTPA token expiration time using the Tivoli Integrated Portal user interface.

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 IBM System Storage DS8000 element manager and other element managers that do not run within an IBM WebSphere Application Server instance use the authentication service that is provided with Tivoli Integrated Portal. The authentication service client is typically embedded in these other applications and the client communicates with the authentication service server in Tivoli Integrated Portal for all authentication requests.

Use of LTPA tokens with single sign-on

Tivoli Storage Productivity Center uses Lightweight Third Party Authentication (LTPA) tokens, which were developed by IBM WebSphere and Lotus organizations, for passing the user information between applications. To use LTPA tokens for single sign-on, each participating application must possess the same set of keys to encode and decode the user information contained in the token. As an additional security feature, the LTPA tokens expire after 24 hours by default (you can change the expiration time using the Tivoli Integrated Portal graphical user interface (GUI). After the LTPA token expires, the user ID and password information must be re-entered to continue accessing the participating applications.

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 IBM System Storage DS8000 element manager and other element managers that do not run within an IBM WebSphere Application Server instance use the authentication service that is provided with Tivoli Integrated Portal. The authentication service client is typically embedded in these other applications and the client communicates with the authentication service server in Tivoli Integrated Portal for all authentication requests.

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

During the installation of Tivoli Storage Productivity Center, you can specify whether to use LDAP or the operating system as the authentication and authorization repository (see the following description of each authentication type). If you do not specify LDAP, then Tivoli Storage Productivity Center uses the operating system (OS) users and groups on the server whereTivoli Storage Productivity Center is installed for authentication and authorization.

OS Authentication
This method authenticates the user against the users 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 using the Tivoli Integrated Portal GUI.

During installation, you can specify that Tivoli Storage Productivity Center use an existing Tivoli Integrated Portal instance rather than the instance that is installed with Tivoli Storage Productivity Center. In this case, Tivoli Storage Productivity Center uses the authentication repository that is used by Tivoli Integrated Portal (LDAP or the operating system). However, if the existing Tivoli Integrated Portal instance uses a file-based authentication repository, you cannot use the existing Tivoli Integrated Portal instance. You must install a new Tivoli Integrated Portal instance when you install Tivoli Storage Productivity Center.

Additional considerations

- If you install Tivoli Storage Productivity Center so that users are authenticated against an LDAP-based repository, you must first ensure that all the LDAP attributes (for example, uid, cn, ou, o, and c) used in the user and group distinguished names or as the user or group naming attribute, are configured to be not case-sensitive within the LDAP-based repository. For example, in IBM Tivoli Directory Server version 6, you can adjust the attribute case-sensitivity through the Equality Matching Rule for each attribute. By default, the Equality Matching Rule for most attributes in IBM Tivoli Directory Server is "caseIgnoreMatch" or "caseIgnoreIA5Match." Both of these settings are usable for integration with Tivoli Storage Productivity Center.
- The single sign-on feature is not available for the IBM XIV Storage System element manager, XIV Storage Manager.
- The single sign-on feature is not supported by the Tivoli Storage Productivity Center command-line interface (CLI).

Planning for enterprise-rollup reports

This topic provides information about planning for enterprise-rollup reports.

IBM Tivoli Storage Productivity Center enables you to use a single interface to generate reports based on data collected by multiple Tivoli Storage Productivity Center servers. You can perform the following actions:

- View the storage information that is gathered by other servers from a single Tivoli Storage Productivity Center user interface using enterprise-wide rollup reports. These 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.

  Note: Performance-related reports are not rolled up.

- Partition a large number of agents across multiple Data servers. This alleviates performance issues when running aggregation reports and scan jobs for multiple agents on one Data server.

To enable rollup reporting, you must select one Tivoli Storage Productivity Center server to act as the "master" server that gathers enterprise-wide data for reports. You can then associate other Tivoli Storage Productivity Center servers to that master server as "subordinate" servers. These subordinate servers provide data to the master server about the entities they manage.

It is possible set up an rollup-reporting environment where fragmentation occurs. Fragmentation occurs when related entities are being monitored by different subordinate Tivoli Storage Productivity Center servers.

Note: Keep in mind the following information when working with master and subordinate servers:

- The master and subordinate servers can change over time.
- A master server can also manage entities just like any other server and report on these entities.
- We recommend that the master server should monitor no more than 500 unique data sources. This number includes subordinate servers, Storage Resource agents, in-band Fabric 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, in-band Fabric agents, CIMOM agents, and VM servers (VMWare). Once this threshold has been met for a server, a new server should be deployed and all new agents pointed to it.
- You must have Tivoli Storage Productivity Center superuser or Tivoli Storage Productivity Center administrator authority to perform administrative operations such as adding a subordinate server to the master server list to monitor.
- After upgrading to Tivoli Storage Productivity Center v4.2 or later from a previous version, you must run probes against a master server's monitored storage assets to have information about those assets appear in **Tivoli Storage Productivity Center > Reporting > Rollup Reports**.
- When you add a subordinate server to the master server list, you must provide the host authentication password of the subordinate server to the master server. The host authentication password is checked each time the master server is probed. 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.

Planning for enterprise roll up reports configuration

This topic provides information about planning for enterprise roll up reports configuration.

Follow these general steps to get enterprise-wide roll up reports:

1. Run discoveries, probes, and scans on the master server to gather information about locally managed entities.
2. Register the subordinate server with the master server. Click **Administrative Services > Data Sources > IBM Tivoli Storage Productivity Center Servers**. Click **Add Tivoli Storage Productivity Center Server**. Enter information for the subordinate server including the host authentication password. The host authentication password is checked each time the master server is probed.

3. Run discoveries, probes, and scans on the subordinate servers to gather information about their locally managed entities.
4. Define and run Tivoli Storage Productivity Center server probes from the master server to collect storage information that has been gathered by subordinate servers.
5. View the information gathered by the Tivoli Storage Productivity Center server probes in roll up reports under **Tivoli Storage Productivity Center > Rollup Reports**.

Planning for the SAN configuration tools

This topic provides information about planning for the SAN configuration tools.

System administrators can identify potential problems with storage area network (SAN) configurations by using the Configuration Analysis and Configuration History tools.

Configuration Analysis
This tool allows the system administrators to check whether existing configurations violate defined best practices and policies. The user selects a set of best practices (policies) and specifies the name of a target SAN (or zone set) through the GUI, and then the tool checks to see whether the selected best practices are being violated in the specified SAN. The policy violations are displayed as IBM Tivoli Storage Productivity Center alerts. The topology viewer displays these policy violations and the affected SAN entities involved in each violation.

Configuration History
This tool displays the changes that have occurred in the SAN environment over time. Using this tool, the user can periodically or on-demand take snapshots of the SAN configuration managed by Tivoli Storage Productivity Center. Using the configuration history viewer, the user can select any two snapshots and view the changes that have occurred in the SAN environment between those two time periods when snapshots were taken. Changes occur when SAN entities are added, removed, or modified.

To create and schedule jobs using these tools, you must have Tivoli Storage Productivity Center superuser or Tivoli Storage Productivity Center administrator authority.

You must have the IBM Tivoli Storage Productivity Center Standard Edition license to use these tools.

Planning for configuration analysis

Use the configuration analysis tool so that you can determine whether an existing SAN configuration complies with predefined best practices (policies).

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. The policy violations are displayed as Tivoli Storage Productivity Center alerts.

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.

Note: The configuration analysis feature is available only with the IBM Tivoli Storage Productivity Center Standard Edition.

The policies used by the configuration analysis tool are as follows:

1. Each connected computer and storage system port must be in at least one zone in the specified zone sets.
2. Each host bus adapter (HBA) accesses storage system ports or tape ports, but not both.
3. Each volume can be accessed only by computers running the same type and version of operating system.
4. Each zone contains only HBAs from a single vendor.
5. Each zone contains only a single model of storage subsystem.
6. Each zone is part of a zone set.
7. Each host must be zoned so that it can access all of its assigned volumes.
8. Each computer has only HBAs of the same model and firmware version.
9. For each host type and operating system, every HBA of a given model must have the same firmware version.
10. Every SAN switch of a given model must have the same firware version.
11. Every storage subsystem of a given model must have the same firmware version.
12. Each fabric can have a maximum of x zones.
13. Each zone can have a maximum of x zone members.
14. Replication plan is intact with respect to the storage resource groups (SRGs) and the replication session associated during planning through the SAN Planner (for replication planning).
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 (for replication planning).
16. Inter/intra site connectivity is valid for replication plan deployments (for replication planning).

To use the configuration analysis tool, follow these general 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 Device server and Data server are up and running.
3. Ensure that the Tivoli Storage Productivity Center agents are installed and configured.
4. Go to **IBM Tivoli Storage Productivity Center > Analytics**. Right-click **Configuration Analysis**. Use the Create Analyzer window to select the scope and policies for the configuration analysis.
5. Schedule how often you want the job to run and how to handle time zones.
6. Save and run the job.
7. After the configuration analysis job has run, you can view the alerts generated by IBM Tivoli Storage Productivity Center. In the navigation tree, expand **IBM Tivoli Storage Productivity Center > Alerting > Alert Log > Configuration Analysis**.

For more information about using the configuration analysis tool, see the Information Center. Search for **Using configuration analysis**. You can also find this information in the *IBM Tivoli Storage Productivity Center User's Guide*.

Planning for the Configuration History tool

The Configuration History view is a variation of the topology viewer. The Configuration History view shows the current configuration as well as what the configuration looked like in the past.

Use the snapshot selection panel in the Configuration History page to view changes that occurred in your storage environment between two or more points in time. After you define the configuration history settings, the system takes snapshots of your configuration so that you can compare the configuration at different points in time. The comparison can help you analyze system performance, as well as enhance the quality of provisioning and planning tasks. The Configuration History tool does *not* show the *current* configuration.

To use the Configuration History tool, follow these general steps:

1. You must have run discovery and probe jobs for the computers, fabrics, switches, storage, and other objects of interest. For computers, you must also have run scan jobs.
2. Make sure that the Device server and Data server are up and running.
3. In the GUI, go to **Administrative Services > Configuration > Configuration History Settings**. Enable the periodic scheduling of configuration history snapshot creation or create some on-demand configuration history snapshots.
4. In the GUI, go to **IBM Tivoli Storage Productivity Center > Analytics > Configuration History** for the following:
   - Computers
   - Fabrics
   - Switches
   - Storage
   - Other

   A Configuration History topology view is displayed of the item you selected.
5. Select the time range for the snapshots and click **Apply**.
6. You can view the SAN configuration changes in the configuration history view for the two time periods you selected.

For more information about how to use this tool, see *IBM Tivoli Storage Productivity Center User's Guide*.

SAN Planner overview

SAN Planner assists the user in end-to-end planning involving fabrics, hosts, storage controllers, storage pools, volumes, paths, ports, zones, zone sets, storage resource groups (SRGs), and replication. Once a plan is made, the user can select to have the plan implemented by SAN Planner.

Supported Systems

You can use the following products with SAN Planner:

- TotalStorage Enterprise Storage Server
- System Storage DS8000
- DS6000

- SAN Volume Controller
- Storwize V7000
- The XIV system

SAN Planner supports the **Space Only** workload profile option for any other storage system.

Requirements for SAN Planner

For planning, IBM Tivoli Storage Productivity Center must be managing the host system, subsystem, and the fabric interconnecting them. If the host, subsystem, or fabric information is not collected by performing subsystem probes and fabric discovery, SAN Planner is not able to generate the plan or execute portions of the plan and issues an error message.

For volume or virtual disk creation, Tivoli Storage Productivity Center must be managing the subsystem. You must have run the subsystem probe either through an active CIM object manager (CIMOM) or native API connection as applicable for the subsystem.

For volume or virtual disk assignment, Tivoli Storage Productivity Center must be managing the host, fabric, and subsystem. The host and subsystem must be in the same fabric. If zoning changes are required, the SAN planner can be used to configure the zoning.

For fabric zoning configuration or fabric zone control operations, Tivoli Storage Productivity Center must manage the fabric. Depending on fabric type and agent capabilities, Tivoli Storage Productivity Center uses a CIM agent, inband agent, or Storage Resource Agent (SRA).

For a plan that requires a virtual disk to be assigned to a host, the virtual disk must have access to two nodes within an I/O group on the SAN Volume Controller or Storwize V7000. These two nodes must be shared on the same fabric as the host.

A storage controller performance monitor must be run to select the **Workload Profile** options. For SAN Volume Controller or Storwize V7000, it is required that the performance data is collected for all its backend subsystems as well. If a performance monitor has not been run, the **Space Only** workload profile option is allowed for planning based on available capacity only.

The IBM Subsystem Device Driver (SDD) installed on a host is required for multipath planning.

Note:

- The Linux SDD driver is no longer available. Starting with Red Hat Enterprise Linux Version 5 and SUSE Linux Enterprise Server Version 10, only the DM_Multipath is available. The DM_Multipath is supported by SAN Planner
- 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 SAN Planner.
- If you are using the SAN planning tools which collect data from Oracle Solaris systems, the HBAs must be configured to use Solaris systems. This note applies to Storage Resource agents.

Run the following command on the Solaris system:

```
cfgadm -al
```

If the system displays this message:

```
fc-fabric connected unconfigured unknown
```

run this command:

```
cfgadm -c configure c#
```

where *c#* is your HBA controller ID. An example of an HBA controller ID is: c5.

You must reboot the Solaris system.

- The SAN Planner does not support Easy Tier (mixed pools) as input to a plan for automated storage provisioning. These pools are not displayed when you select resources for planning or request a list of recommended pools.
- When creating a plan with multipath provisioning, do not include computers that have EMC Powerpath installed. For example, do not include EMC Clariion volumes in a plan with multipath provisioning.

For replication planning, you must add all the subsystems to be used in planning to IBM Tivoli Storage Productivity Center for Replication. You must also ensure that connectivity is established between the subsystems.

Creating a SAN plan

You can create a SAN plan by right clicking the **IBM Tivoli Storage Productivity Centre** > **Analytics** > **SAN Planner** node and selecting **Create Plan**. The planner wizard is started.

You can also create a SAN plan from the **IBM Tivoli Storage Productivity Centre** > **Disk Manager** > **SAN Planner** node and selecting **Create Plan**. The planner wizard is started.

The type of plan that you select determines the content of the remaining pages in the wizard. The options are:

Provision storage only, including optional multipath and zone considerations
: Click to create a plan for provisioning new storage. The plan does not include provisioning for storage replication.

Provision storage with replication, including optional multipath and zone considerations on the source resources
: Click to create a plan for provisioning new storage and storage replication.

Provision replication to existing storage
: Click to create a plan for provisioning replication for existing storage (provision target volumes for existing source volumes).

Configure multipathing or zoning only
: Click to create a plan for provisioning multipathing and zoning for existing storage.

To print plan information from the interface, click **File** > **Print** on the Tivoli Storage Productivity Center menu bar, or click  on the toolbar.

Provision storage only, including optional multipath and zone considerations

Use this option to provision storage volumes or virtual disks with multipathing and zoning. The input resources for planning can be storage subsystems, volumes, hosts, or storage resource groups (SRGs). If you are provisioning storage for SAN Volume Controller or Storwize V7000, you must select **Virtual Disks**.

An SRG is a set of entities managed by Tivoli Storage Productivity Center. These entities can be servers, switches, storage subsystems, fabrics, storage pools, and storage volumes. Using SRGs provides you with a means of separating your application data from your pool of unused resources. You can specify policies for provisioning (volume creation and selection, workload profiles, zoning, and multipathing configuration) and associate them with SRGs. These policies are used by SAN Planner to populate default settings.

You can use this option to plan and select appropriate storage controllers, storage pools, and storage volumes (when using unassigned volumes) that satisfy your inputs. This option enables you to select controller type preference, whether storage requests can be satisfied by multiple controller types, and the Redundant Array of Independent Disks (RAID) level. For SAN Volume Controller or Storwize V7000, the RAID level represents the RAID level for the backend controllers.

When creating new volumes on an extent pool, SAN Planner uses one of the following methods:

- If an encryption group is selected, SAN Planner creates volumes on pools that belong to the selected encryption group. Only pools that belong to that encryption group are used for planning. This option might not be enabled if the storage system does not support encryption.
- If an encryption group is not selected, SAN Planner uses the current performance utilization of storage resources to determine whether a new volume should be allocated on particular pool in a particular storage controller. If multiple storage pools from different controllers can potentially satisfy your provisioning request, then the SAN Planner uses the rated utilization of the pools (the sum of the previous provisioning performance requirements, which might be greater than the current utilization) to break the ties and select a candidate storage pool.

You can also select other options in the SAN Planner such as selecting a workload profile as an input to the SAN Planner's analysis, excluding solid-state drives from planning, or using thin provisioning for creating thin provisioned volumes based on the thin provisioning profile selected.

See the **Configure multipathing or zoning only** section for details on SAN Planner's multipathing and zoning options.

Provision storage with replication, including optional multipath and zone considerations on the source resources

In addition to provisioning storage with multipathing and zoning options, when you select replication options in the wizard, SAN Planner provisions replication for your storage resources. Resilient resource planning is available only for devices which are supported by IBM Tivoli Storage Productivity Center for Replication.

Depending on the storage resources you are planning for, you can select from the following copy services session types:

- FlashCopy®
- VDisk mirroring
- Synchronous Metro Mirror failover/failback
- Synchronous Metro Mirror failover/failback with practice
- Synchronous Metro Mirror single direction
- Asynchronous Global Mirror either direction with two-site practice
- Asynchronous Global Mirror failover/failback
- Asynchronous Global Mirror failover/failback with practice
- Asynchronous Global Mirror single direction
- Three-site Metro Global Mirror with practice
- Three-site Metro Global Mirror
- Snapshot

For example, if you are planning for the XIV system, you can only select from Snapshot, Metro Mirror, or Global Mirror session types. Practice sessions are not supported for the XIV system.

Provision replication to existing storage

This option allows you to extend existing replication sessions by adding more storage resources and protecting them. The new volumes are added to existing replication sessions. SAN Planner ensures that source volumes are added to the source side of a copy relationship and target volumes are added to the target side. For Snapshot copy sessions, SAN Planner does not provision or manage target volumes. The XIV system selects and manages the targets for Snapshot replication.

Configure multipathing or zoning only

SAN Planner enables you to set up multipath options. The supported multipath driver on a host is the IBM SDD.

Note:

- The Linux SDD driver is no longer available. Starting with Red Hat Enterprise Linux Version 5 and SUSE Linux Enterprise Server Version 10, 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 SAN Planner.

SAN Planner enables system administrators and storage area network (SAN) administrators to plan and implement storage provisioning for hosts and storage systems with multipath support in fabrics managed by Tivoli Storage Productivity Center. Planning the provisioning of storage to hosts with multipath drivers requires knowing which storage systems are supported by the host multipath driver and the multipath modes that are supported by the driver and storage subsystem. Planning the paths between the hosts and storage controller requires designing paths between hosts and storage subsystems which are implemented through zones in the fabric.

SAN Planner is used for specifying multiple paths options between selected hosts and storage systems. There should be at least two paths available to a storage

system and at least four paths available to SAN Volume Controller. Although not required, it is beneficial to have at least two host ports.

Path planning in SAN Planner does not directly interact with zone planning. The path planning option provides the path information which the zone planning option uses. Each path is represented by a host port WWPN, target port WWPN, and a Volume ID for the volume on the target that is mapped to the host port. These paths are created when the zone planning option is implemented.

SAN Planner specifies automatic zoning between ports on the selected hosts and subsystems in a fabric. All zoning is done with WWPN-based zoning. SAN Planner plans zoning configuration for new storage that is provisioned for use by a host. For example, the zoning option can be used when a new storage volume is created and assigned to a host. The zoning option can also be used with a volume that has already been created and is assigned to a host that requires more storage. In these cases, SAN Planner determines which host and storage need to be zoned together, which provides the zoning option with the exact set of ports that need to be zoned together. SAN Planner then uses the zoning inputs for the planning.

If the subsystem/host are within the same fabric and the zoning option is not checked, then existing zones or zone sets are used. If the zoning option is checked, a new zone or zone set is created.

To perform zoning, the host and subsystem being zoned by SAN Planner must be within the same fabric. For the case where the host and subsystem reside in more than one of the same fabrics, SAN Planner creates identical zones in all the fabrics.

The guidance policies used for zone planning are as follows:

- One zone per host bus adapter (HBA)
- One zone per host
- One zone per host port
- Auto zone: largest zone that satisfies the validation policies

The validation policies used for zone planning are as follows:

- Maximum number of zone members in a zone = N.
- No two controllers of different types should be in the same zone.
- Maximum number of zones in a fabric = N.

The second policy in the preceding list has a higher priority than the third policy.

Planner and user inputs interdependencies

The planner options depend on the inputs that you select. For example, selecting a host that does not have an agent installed limits the planning that can be done. The planning process reveals inconsistencies in your preferences. For example, it might not be possible to have your specified number of paths and have a different zone for each host. In such cases, the planner makes a best effort to satisfy your inputs. If the planner cannot satisfy conditions, it issues warnings. Errors discovered before plan execution are displayed in the planner. You can adjust the inputs to correct the problem. Errors discovered during plan execution are in the job execution log.

Storage Optimizer overview

The Storage Optimizer uses data in the IBM Tivoli Storage Productivity Center database to analyze your storage subsystems to identify performance bottlenecks, and recommend changes to improve performance. This topic lists the supported subsystems and describes the general steps for using Storage Optimizer.

The Storage Optimizer helps you develop storage migration or storage consolidation plans, and helps you plan for the growth of your storage infrastructure. For example, you can use the Storage Optimizer to create a weekly report that provides details on the performance of your storage subsystems and makes recommendations for how to improve performance.

Note: The Storage Optimizer does not actually perform any migrations or make any modifications to subsystem configurations. Its primary purpose is to provide you with a performance analysis and optimization recommendations that you can choose to implement at your discretion.

Important: Storage Optimizer does not take into account any established replication relationships or sessions. Migration recommendations should be followed with care to ensure continuity of all replication relationships.

To use the Storage Optimizer, you must have a IBM Tivoli Storage Productivity Center Standard Edition license.

The following IBM storage subsystems or applications are supported:

- DS8000
- DS6000
- DS5000
- DS4000
- TotalStorage Enterprise Storage Server
- SAN Volume Controller
- Storwize V7000

Before running Storage Optimizer, you must set up performance monitors to collect performance data at least once a day for all storage subsystems that you want Storage Optimizer to analyze. You must also collect performance monitoring data for the back-end subsystems for SAN Volume Controller or Storwize V7000. The performance data for the subsystems produce the most accurate Storage Optimizer analysis.

The analysis is less accurate if there are any gaps in the data collection for the time interval being analyzed. It is recommended that you collect at least one week's worth of performance monitoring data before using Storage Optimizer. Providing a longer time interval for data collection will increase the accuracy of the Storage Optimizer analysis and recommendations.

You must provide the Storage Optimizer with the following input:

- Configuration and performance data that you collect for the supported storage subsystems before running the Storage Optimizer
- One or more selected storage subsystems that you specify as input to the Storage Optimizer.
- Connected fabric topology and zoning information will be used in the analysis if it is available.

The Storage Optimizer produces the following output:

- An analysis report that displays performance heat maps and tables that graphically illustrate the performance utilization of the storage subsystems that you specified as input.
- An optimization report that provides migration and consolidation recommendations for improving performance.

The Storage Optimizer produces more accurate results if you provide both the required and optional input data requested by the Storage Optimizer, and if you collect storage subsystem performance monitoring data over a longer time interval before using the Storage Optimizer.

Note: Since running the Storage Optimizer is a processor-intensive task for the IBM Tivoli Storage Productivity Center server, schedule a time to run the Storage Optimizer when processor demand on the IBM Tivoli Storage Productivity Center server is expected to be at a minimum.

The general steps for using the Storage Optimizer are as follows:

1. In the Create Analysis window, enter information in the required fields and choose **File > Save** to submit the analysis report job.
2. Review the Storage Optimizer analysis report.
3. Create an optimization report based on this analysis.
4. View the optimization report.
5. At your discretion, implement the recommendations suggested in the optimization report, or re-run the analysis using different inputs.

Planning for the Data Path Explorer

This topic provides information about planning for the Data Path Explorer view in the topology viewer.

The Data Path Explorer is a type of view in the topology viewer. Data path explorer combines several of the typical topology views into a visualization that allows users to follow 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. This action allows you to use the Data Path Explorer to check the health and performance statistics of the storage components along the I/O path.

Different agents play different roles when discovering information and what can be shown in the topology viewer. The Data Path Explorer view is also subject to those limitations as well. Different amounts of information can be shown by the Data Path Explorer view depending on the agents deployed throughout the fabric.

To display data paths from a host to a subsystem in this view, it is necessary to have the following conditions:

- CIM agent installed to monitor the storage subsystem (CIM agents are not required for DS8000, the XIV system, SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000)
- Storage Resource agent installed to monitor the host or fabric

The Data Path Explorer view does not display the data path from a host to a subsystem if any one of these conditions is not met.

To show the host disk assignment from a given storage subsystem, it is necessary for the IBM Tivoli Storage Productivity Center Data server to have access to the CIM agent (SMI-S agent) or native interface for the subsystem. For the subsystem, the data path is defined by the host masking or mapping information maintained by the CIM agent or native interface.

Analysis on data paths involving SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 would require the native interface pointing to the SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 and the disk drives for the storage system.

To show accurate information, it is important to have the most current information from the CIM agents, native interface, or Storage Resource agents. You must probe all the components within the data path (or a set of agents strategically placed throughout the fabric).

To show physical fabric connectivity, you must have a Storage Resource agent, out-of-band Fabric agent, or a CIM agent for the fabric. (CIM agents are not required for DS8000, the XIV system, SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000.)

To show performance overlay in the Data Path Explorer view, you must first create and run performance monitors on the switches and storage subsystems of interest. To see the overlays you must also go to the topology viewer. Right-click the topology viewer. Click **Global Settings**. On the Settings window under Active Overlays, click the **Performance** check box.

To show the zone overlay, you need an in-band Fabric agent or Storage Resource agent connected to the fabric along the data paths defined or a CIM agent or Storage Resource agent for the fabric. To see the overlays you must also go to the topology viewer. Right-click the topology viewer. Click **Global Settings** and check the **Show zone tab** check box. Click **OK**.

If you want to launch the Data Path Explorer view for a host, that host must have a Storage Resource agent running on it.

These are the general steps to follow to display the data path explorer view:

1. Install and configure the SMI-S agents (CIM agents) for the subsystems to collect subsystem data. CIM agents are not required for DS8000, the XIV system, SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 to collect subsystem data. Install and configure the CIM agents for the fabric to collect fabric data. Contact your vendor for information about installing and configuring the CIM agent for the subsystem or fabric.
2. Install and configure the out-of-band Fabric agent to monitor a given fabric.
3. Install the Storage Resource agent on the hosts.
4. Run a discovery job for the CIM agents, Storage Resource agents, and storage systems.
5. Run a probe job for the subsystem, fabric and Storage Resource agents.
6. Run a performance monitoring job for the subsystem and switches if you want to see performance information.
7. You can now see the data paths from a host to a storage subsystem in the Data Path Explorer view.

Planning for the Data Manager

This section provides planning information for the 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 utilization, 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 platforms to help you 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 platforms to help you identify and reclaim space 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 lower storage costs by:

- Improving storage utilization.
- 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 collected by Data Manager helps you understand what is really going on with data on your servers and in your storage environment. You can 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.

In addition to understanding the current consumption and usage of data within the enterprise, Data Manager tracks the information over time. Not only does this historical view of storage consumption and 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 need to purchase additional capacity in a planned proactive manner rather than reacting to out-of-space emergencies.

Use the 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 that can detect potential problems and automatically adjust parameters based on the policies and actions you have established. Use the capability to provision storage based upon storage policies to expand a file system, and to allocate storage to a volume (LUN).

Planning for the Storage Resource agents on Oracle Solaris

You can skip automount maps (automaps) for a discovery job on Solaris. To skip automount maps, specify the **skipAutoFS=1** parameter in the server section of the TPCD.config file. The change globally affects all Solaris Storage Resource agents managed by the Data server.

To change the TPCD.config file, follow these steps:

1. Stop the Data server.
2. Modify the TPCD.config file in this directory:

   TPC_Data_Server_install_directory/config

   The default directory for Windows is:

   C:\Program Files\IBM\TPC\data\config

   The default directory for UNIX or Linux is:

   /usr_or_opt/IBM/TPC/data/config
3. Modify the TPCD.config file with the **skipAutoFS=1** parameter added:

```
[server]
   threadPoolSize=3
   abbreviatedProbe=1
   maxThreads=8
   pingReceiveTimeout=10
   skipAutoFS=1    <== Set to 1 for Discovery
        on Solaris Storage Resource agent to skip AutoMounts
        process.By default, discovery will always
        process AutoMounts on the  Solaris Storage Resource agent.

[gui]
   threadPoolSize=3
   maxThreads=10
   reportRowLimit=5000
   keepCachedReport=120
```

   Save the file.
4. Restart the Data server.

Planning for the Disk Manager

This section provides information about planning for Disk Manager.

Because Tivoli Storage Productivity Center provides an abstraction for the storage subsystems managed to make the heterogeneous environment look more homogenous, Tivoli Storage Productivity Center uses a number of terms, which map to different things when looking at the actual storage subsystems. Most of these terms are derived from SMI-S, which already provide a common model for storage subsystems.

These terms are:

storage pool
: A collection of storage capacity with a given set of capabilities. A pool has certain storage capabilities, which indicate the range of quality of service requirements that can be applied to objects created from the pool.

primordial pool
: A type of storage pool. This pool might contain unformatted or unprepared capacity. This type of pool might simply contain unassigned

capacity. Storage capacity is drawn from the primordial storage pool to create concrete storage pools. The primordial storage pool aggregates storage capacity that has not been assigned to a concrete storage pool. Storage volumes are allocated from concrete storage pools. For the TotalStorage Enterprise Storage Server, DS8000, and DS6000, the primordial pool is the disk groups or array sites which are installed in the machine but have not yet been configured into RAID arrays. Primordial pools for SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 are those MDisks which are available, but have not yet been configured to any MDisk Group (one primordial pool per back-end controller). The Primordial pool for the XIV system is a virtual concept that represents the aggregation of system-wide deallocated storage capacity which is available but unassigned to XIV system storage pools.

storage volumes
Allocations of storage capacity that is exposed from a system through an external interface. In SCSI terms, they are logical units.

The following table shows the mapping of the Tivoli Storage Productivity Center terms to the device-specific terms.

Tivoli Storage Productivity Center terms:	Storage pool	Primordial pool	Storage volume	Disk
Device terms				
DS6000, DS8000, TotalStorage Enterprise Storage Server	Extent pool	not applicable	Volume	Disk drive module (DDM)
DS4000	Volume group	not applicable	Volume	Disk
DS5000	Volume group	not applicable	Volume	Disk
• SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified	MDisk group	not applicable	Volume	MDisk
XIV system	Storage pool	not applicable	Volume	Disk

Many disk arrays provide an interface for the administrator to specify which initiators can access what volumes through which target ports. The effect is that the given volume is only visible to SCSI commands that originate from the specified initiators through specific sets of target ports. There might also be a capability to select the SCSI Logical Unit Number as seen by an initiator through a specific set of ports. The ability to limit access is called *device masking*. The ability to specify the device address seen by particular initiators is called device mapping. For SCSI systems, these terms are known as LUN masking and LUN mapping. In Tivoli Storage Productivity Center, masking and mapping is handled through host assignment.

Note:

- For DS6000 and DS8000: The DS6000 and DS8000 tracks the number of spares still available after a spare disk has been used to replace a failing disk drive in the device by marking the failing drive with an operational status of "Predictive

Failure." If there are not enough spare disks left for a rank, then the operational status of the disk also shows a status of "Error."

If the DS6000 or DS8000 marks a disk drive with a status of "Predictive Failure" but not "Error", then Tivoli Storage Productivity Center shows a green icon with a consolidated status of "OK" in the health overlay, and does not create an alert. The topology viewer does show the real status of the disk drive as "Predictive Error" in the Operational Status column of the tabular view. The topology viewer lets you see the changed status.

- For EMC storage systems: The data that Tivoli Storage Productivity Center collects about EMC storage systems is based on SMI-S standards. Within SMI-S model different types of StoragePools 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 system might be represented as a pool within the SMI-S model. Because of this, Tivoli Storage Productivity Center reports and CLI commands might display more pools for EMC storage systems than what is reported through the EMC Navisphere agent.

Planning for the IBM CIM agents

The CIM agents are provided by the vendor of the storage device, fabric switch, or tape library. This topic provides planning information for the IBM CIM agents for the IBM storage subsystems.

For storage subsystems, the CIM agents are needed for storage asset information, provisioning, alerting, and performance monitoring. For tape libraries, the CIM agents are used for asset and inventory information. The CIM agents conform to the SNIA SMI-S specification to provide a communication transport between IBM Tivoli Storage Productivity Center and the managed devices.

Note: For DS8000, the XIV system, SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000 storage subsystems, you do not need a CIM agent.

The CIM agents can be referred to by various names, such as CIMOM (CIM Object Manager) or SMIS agent. A CIM agent consists of a CIMOM and an SMI-S provider for the managed device. The CIM agent can be a separate agent installation or can be embedded in the device itself, as is the case with Cisco fabric switches. In this case, there is no proxy agent to install and Tivoli Storage Productivity Center is configured to point to the managed device itself.

Once the CIM agent is installed and configured, Tivoli Storage Productivity Center can be configured to communicate with it.

Note:

- Do not use different CIM Agent releases (for example 5.1, 5.2.1, 5.3, and 5.4) to manage the same DS6000 or TotalStorage Enterprise Storage Server with IBM Tivoli Storage Productivity Center. Each CIM Agent release reports some information in a different way so this reporting can cause a reporting failure. Use the same release of CIM Agent for all CIMOMs managing the same device.
- If an Engenio or LSI SMI-S provider is being used by the Tivoli Storage Productivity Center server, then no other SMIS-enabled application should be using that Engenio or LSI SMI-S provider. This controlled environment is required to ensure that the Engenio provider receives synchronized CIM Client requests.

- For the TotalStorage Enterprise Storage Server 2105, disk drive information could be displayed twice after an upgrade. The first half of the information is new and is displayed online. The second half of information is no longer detectable. The second half of information is impossible to correlate.

Before installing Tivoli Storage Productivity Center check the Tivoli Storage Productivity Center support site for the latest CIM agent support. Go to the Tivoli Storage Productivity Center support website at http://www-01.ibm.com/support/docview.wss?rs=40&uid=swg21386446&loc=en_US&cs=utf-8&lang=en.

For information about certified SMI-S devices, see http://www.snia.org/ctp/conformingproviders#10provider.

For information about how to install and configure the CIM agent, see http://www.ibm.com/servers/storage/support/software/cimdsoapi/.

Planning for the 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 5.8. This version of the CIM Agent supports only the Array Profile and not the Storage Virtualizer Profile. However, Tivoli Storage Productivity Center supports the TagmaStore as a Storage Virtualizer. Tivoli Storage Productivity Center displays 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* > Managed Disks > LUN** (does not show the relation between the disk and LUN or volume in the tree)
- **Disk Manager > Reporting > Storage Subsystem > Volume to Backend Volume**
- **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 LUN and volume correlation, the host machine must have the Storage Resource agent installed and the TagmaStore device must be in the same SAN fabric. You must also have a zone configured in the active zone set between the ports of the host machine and the ports of the TagmaStore device. The Fabric agent or Storage Resource agent needs to 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.

Before you can use the CIM agent 5.8, you must uninstall pre-version 5.8 CIM agents. You cannot have a mixture of CIM agent 5.8 and pre-version 5.8 CIM agents. A mixture of CIM agents can cause failures.

The events returned by the CIM agent are:

- Generation of a volume
- Deletion of a volume
- Allocation of a path
- Cancellation of a path

When you delete a volume, you can only delete the latest created volume. For example, if you create five volumes, you need to delete the volumes in reverse order from the creation order.

General procedure

The general steps to follow to use the TagmaStore CIM agents are:

1. Install or upgrade Tivoli Storage Productivity Center.
2. Remove the pre-5.8 CIM agent from Tivoli Storage Productivity Center as a data source before adding the 5.8 CIM agent. You must add the CIM agent before running a discovery job.
3. Run a discovery job for the TagmaStore CIM agent.
4. Run probe jobs, scan jobs, and ping jobs for the TagmaStore CIM agent.
5. View the storage information gathered by the monitoring jobs through the topology viewer and reports that you can generate through the Fabric Manager, Data Managerr, and Disk Manager.

Planning for Disk Manager configuration

This section provides information for planning Disk Manager configuration.

Disk Manager data can only be collected from devices managed by a CIM agent or through the native storage system interface. To get Disk Manager data, follow these general steps:

1. Prepare for collecting data from Disk Manager.

   For those devices that need the CIM agent, make sure that the following items are done before adding the CIM agent:

   - The version of CIM agent and firmware for the device is supported. For information about the CIM agent and firmware supported, see http://www-01.ibm.com/support/docview.wss?rs=40&q1=support+matrix&uid=swg21386446&loc=en_US&cs=utf-8&lang=en. Select the appropriate version for **Storage**.
   - A CIM agent is installed on a different server from IBM Tivoli Storage Productivity Center. For some vendors, the device comes with an embedded CIM agent, and therefore, no separate CIM agent installation is required.
   - The CIM agent is configured to manage the intended device.

   For information about how to install and configure the CIM agent, contact your CIM agent provider. Also see *IBM Tivoli Storage Productivity Center Hints and*

Tips for information about installing and configuring CIM agents. See http://www-01.ibm.com/support/docview.wss?uid=swg27008254.

2. For devices that use CIM agents, follow these steps:
   a. Add the CIM agent connection to Tivoli Storage Productivity Center. For information about adding a CIM agent, use the Configure Devices wizard. For information about the Configure Devices wizard, see "Configure Devices" on page 266.
   b. Run a CIMOM discovery job for the storage system.
   c. Run a probe job for the storage system. The probe job gets all component level information into the database. The probe job also includes getting all the instances of controllers, arrays, extents, volumes, LUN to host mappings, FC ports, and so on.
3. For devices that use the native interface, add the device to Tivoli Storage Productivity Center. Use the Configure Devices wizard to add the device. For information about the Configure Devices wizard, see "Configure Devices" on page 266. On the Select Device Type panel, select **Storage Subsystem**. After you add the storage system to Tivoli Storage Productivity Center, an automatic discovery and probe jobs are run.
4. Once the device is added to Tivoli Storage Productivity Center, automatic discovery and probe operations are done for the storage system.
5. Create alerts. You can specify which alerts you want to be notified of when monitoring the storage systems.

   For the storage systems that use the native interfaces, Tivoli Storage Productivity Center gets events directly from the storage system. For the DS8000, Tivoli Storage Productivity Center receives events from the DS8000 enterprise storage server network interface (ESSNI). For the SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000, and the XIV system, Tivoli Storage Productivity Center polls the device every minute for events.

   For storage systems that use the CIM agent, Tivoli Storage Productivity Center receives CIM indications.
6. You can configure Tivoli Storage Productivity Center to send SNMP traps.

   if you configure Tivoli Storage Productivity Center to send SNMP traps or Omnibus events, you must first import the Tivoli Storage Productivity Center MIB and Baroc files into your SNMP manager or Omnibus manager. For more information about configuring SNMP traps, see "Configuring Tivoli Storage Productivity Center to send SNMP traps and IBM Tivoli Enterprise Console events" on page 299.
7. Run a probe job for storage systems other than DS8000, the XIV system, SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000. Probe jobs collect detailed information about the storage systems known to Tivoli Storage Productivity Center. In contrast to a CIMOM discovery job, the probe interacts with the CIM agent to get information about just the storage devices configured to be probed. You can set up a probe job to gather data for just one storage system, or you can set up probe jobs to collect data from more than one storage system. A probe job must be run before volumes can be created on a storage system or a storage system can be monitored for performance management.

   Running a probe job puts additional workload on the Tivoli Storage Productivity Center server, the repository database, and the CIM agents. For probe jobs, the additional workload is especially important if multiple devices are managed by the same CIM agent and are probed concurrently. As a result, the probe completion time is impacted by the number of devices probed in

parallel and the number of CIM agents involved in the probe job. This fact is especially important if the CIM agent machines are at the lower end of the hardware prerequisites.

To help keep the workload low, run the probe jobs for the storage systems at a time when Tivoli Storage Productivity Center as well as the network are not used heavily, which is typically at night time. To allow this workload distribution, run probe jobs for each storage system to run one after the other.

The following example shows how you might run your probe jobs. This environment assumes that you have four storage systems (SS1, SS2, SS3, and SS4), with three of them having a probe job duration of less than 30 minutes and one (SS3) taking 90 minutes to complete a probe job. You could arrange the probe jobs according to the following table:

Probe	Probe 1	Probe 2	Probe 3	Probe 4
Storage system	SS1	SS2	SS3	SS4
Start time	0:30	1:30	2:30	4:30

If you have a probe job that collects data from more than one storage system, then configure the probe job to contain only one storage system per CIM agent. The configured probe job is important because the CIM agent response time is a major contributor to the overall Tivoli Storage Productivity Center probe job duration.

The storage system probe jobs collect detailed information about the configuration of the storage systems and the properties describing those configuration elements. The elements gathered are:

- Component computer systems (like the SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 nodes)
- Storage pools
- Storage volumes
- FC ports
- Disks
- Host-to-volume assignments
- Relationships among the elements (for example storage pool to storage volume relationships)

This configuration is typically not changing that much, especially because Tivoli Storage Productivity Center updates the repository database whenever you do configuration changes using the Tivoli Storage Productivity Center GUI (for example when creating a volume). Therefore, the probe jobs for the storage systems can be scheduled to occur less often than the CIMOM discovery.

If you use the device-specific CLI for configuration changes, the probe is needed in case the CIM agent does not support all needed indications to inform Tivoli Storage Productivity Center about the change.

Probe storage systems once every one or two days. Other factors influencing this interval are:

- Is the device managed by Tivoli Storage Productivity Center or is an alternative tool used (for example device-specific CLI)?
- Does the CIM agent support CIM indications and how extensively are Tivoli Storage Productivity Center alerts used?

8. You can view the Disk Manager reports and topology in Tivoli Storage Productivity Center.

For information about how to configure Tivoli Storage Productivity Center for performance monitoring, see “Performance management configuration” on page 119.

Planning for TotalStorage Enterprise Storage Server FlashCopy reports

Use this information to plan for displaying TotalStorage Enterprise Storage Server FlashCopy reports in IBM Tivoli Storage Productivity Center.

If you want to display TotalStorage Enterprise Storage Server FlashCopy relationships in additional to the general information in Tivoli Storage Productivity Center, there are two general steps that must be performed. For example:

1. Add the following types to the CIM agent:

   ```
   dscimcli mkdev 10.10.10.100 -type ess -user user_ID
   -password password
   dscimcli mkdev 10.10.10.100 -type esscs -user user_ID
   -password password
   ```

2. You must also have a specific level of microcode. Before installing IBM Tivoli Storage Productivity Center, check the IBM Tivoli Storage Productivity Center support site for the latest microcode level. Go to the Tivoli Storage Productivity Center support Web site at http://www-01.ibm.com/support/docview.wss?rs=40&uid=swg21386446&loc=en_US&cs=utf-8&lang=en.

Note: These steps are only required for the ESS and not the DS8000.

Storage capacity of volumes

This section provides information about how IBM Tivoli Storage Productivity Center displays the storage capacity of volumes.

The disk storage that IBM Tivoli Storage Productivity Center supports is expressed in powers of two:

```
1 KB = 2 to the power of 10 bytes (1024)
1 MB = 2 to the power of 20 bytes (1 048 576)
1 GB = 2 to the power of 30 bytes (1 073 741 824)
1 TB = 2 to the power of 40 bytes (1.09951E+12)
1 PB = 2 to the power of 50 bytes (1.1259E+15)
```

This matches the convention of many storage systems but does not match the behavior of the TotalStorage Enterprise Storage Server server family of devices and the IBM XIV Storage System. Therefore, a volume created using the TotalStorage Enterprise Storage Server Server Specialist or the XIV GUI with a size of 17 GB might be displayed by Tivoli Storage Productivity Center as 16 GB in size. The DS6000 and DS8000 storage system GUIs display the storage as powers of two.

Planning for Tivoli Storage Productivity Center for Disk Select

Tivoli Storage Productivity Center for Disk Select provides storage device configuration, performance monitoring, and management of Storage Area Network (SAN) attached devices from a single console for midrange storage systems.

This product provides basic disk functions and performance monitoring capabilities for IBM System Storage DS3000, IBM System Storage DS4000, IBM System Storage DS5000, IBM System Storage SAN Volume Controller, IBM Storwize V7000 Unified, and IBM Storwize V7000.

Tivoli Storage Productivity Center for Disk Select provides the following features, functions, and benefits:

- Helps monitor and manage performance metrics of the IBM System Storage DS3000, DS4000, and DS5000
- Enables device configuration and management of SAN-attached storage devices from a single console
- Helps simplify the complexity of managing multiple SAN-attached storage devices
- Delivers proactive performance analysis by providing an IT administrator with a single, integrated console for the performance management of IBM System Storage devices
- Helps monitor performance metrics across multiple storage subsystems from a single console
- Helps administrators monitor metrics, such as input and output (I/O), data rates, and cache utilization
- Helps monitor and analyze performance statistics for storage systems to measure service levels by storing received performance statistics into database tables for later use
- Helps analyze and generate reports on monitored devices for display in the central administrative console
- Receives timely alerts that can enable event action based on customer policies by monitoring performance thresholds for the devices based on performance metrics and the generation of alerts when those thresholds are exceeded

Tivoli Storage Productivity Center for Disk Select follows the same procedures for installing and configuring as you would for Tivoli Storage Productivity Center for Disk.

Note: When you install Tivoli Storage Productivity Center for Disk Select , you will also be installing Tivoli Storage Productivity Center for Replication. However, Tivoli Storage Productivity Center for Replication does not support the DS3000, DS4000, or DS5000. If you want to disable Tivoli Storage Productivity Center for Replication, see "Disabling unlicensed or unused Tivoli Storage Productivity Center components" on page 368.

CIM agent

Tivoli Storage Productivity Center for Disk Select requires CIM agents for the DS3000, DS4000, and DS5000. For information about installing and configuring the CIM agent for the DS3000, DS4000, and DS5000, contact your CIM agent (LSI SMI Provider) at http://www.lsi.com/storage_home/products_home/external_raid/management_software/smis_provider/index.html.

The CIM agent must be installed and configured before you can collect data for the DS3000, DS4000, and DS5000 through Tivoli Storage Productivity Center.

See *TPC Hints and Tips* document at this website http://www-01.ibm.com/support/docview.wss?uid=swg27008254 for some useful information about installing the CIM agent.

Also note the following information about the CIM agents:

- Do not install the CIM agent on the same server as the Tivoli Storage Productivity Center server.

- The CIM agent should not monitor more than five storage systems.
- Use the LSI SMI-S Provider V10.06.GG.33 or later for the DS4000 and DS5000 MC level V7.60.x. Otherwise, performance monitoring does not work.
- Do not use the LSI SMI-S agent for the DS4000 and DS5000 with an MC level lower than V7.50.x because this could cause a reboot of the controller device support and SMI-S agent download.

When configuring Tivoli Storage Productivity Center for the CIM agent, you need to enter an Interoperability Namespace. Depending on which storage system you have determines which provider you have and the Interoperability Namespace to use. The table below provides information about namespaces:

You should always check this website for current information about namespaces: http://www-01.ibm.com/support/docview.wss?rs=40&uid=swg21386446 &loc=en_US&cs=utf-8&lang=en. Select the appropriate version under **Storage**. At the bottom of the page is a subsystem support matrix spreadsheet that provides the Interoperability Namespaces.

General procedure

To view the reports and topology for the Tivoli Storage Productivity Center for Disk Select , follow these general steps:

1. Prepare for collecting data from the Disk Manager.

   Install and configure the CIM agent for the DS3000, DS4000, or DS5000.

   For information about how to install and configure the CIM agent, contact your CIM agent provider. Also see *IBM Tivoli Storage Productivity Center Hints and Tips* for information about installing and configuring CIM agents. See http://www-01.ibm.com/support/docview.wss?uid=swg27008254.
2. For devices that use CIM agents, follow these steps:
   a. Add the CIM agent connection to Tivoli Storage Productivity Center. For information about adding a CIM agent, use the Configure Devices wizard. For information about the Configure Devices wizard, see "Configure Devices" on page 266.
   b. Run a CIMOM discovery job for the storage system.
   c. Run a probe job for storage systems other than the XIV system, SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000. The probe job gets all component level information into the database. The probe job also includes getting all the instances of controllers, arrays, extents, volumes, LUN to host mappings, FC ports, and so on.

      Probe jobs collect detailed information about the storage systems known to Tivoli Storage Productivity Center. In contrast to a CIMOM discovery job, the probe interacts with the CIM agent to get information about just the storage devices configured to be probed. You can set up a probe job to gather data for just one storage system, or you can set up probe jobs to collect data from more than one storage system. A probe job must be run before volumes can be created on a storage system or a storage system can be monitored for performance management.

      Running a probe job puts additional workload on the Tivoli Storage Productivity Center server, the repository database, and the CIM agents. For probe jobs, the additional workload is especially important if multiple devices are managed by the same CIM agent and are probed concurrently. As a result, the probe completion time is impacted by the number of devices probed in parallel and the number of CIM agents involved in the

probe job. This fact is especially important if the CIM agent machines are at the lower end of the hardware prerequisites.

To help keep the workload low, run the probe jobs for the storage systems at a time when Tivoli Storage Productivity Center as well as the network are not used heavily, which is typically at night time. To allow this workload distribution, run probe jobs for each storage system to run one after the other.

The storage system probe jobs collect detailed information about the configuration of the storage systems and the properties describing those configuration elements. The elements gathered are:

- Component computer systems (like the SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 nodes)
- Storage pools
- Storage volumes
- FC ports
- Disks
- Host-to-volume assignments
- Relationships among the elements (for example storage pool to storage volume relationships)

This configuration is typically not changing that much, especially because Tivoli Storage Productivity Center updates the repository database whenever you do configuration changes using the Tivoli Storage Productivity Center GUI (for example when creating a volume). Therefore, the probe jobs for the storage systems can be scheduled to occur less often than the CIMOM discovery.

Probe storage systems once every one or two days. Other factors influencing this interval are:

- Is the device managed by Tivoli Storage Productivity Center or is an alternative tool used (for example device-specific CLI)?
- Does the CIM agent support CIM indications and how extensively are Tivoli Storage Productivity Center alerts used?

3. For devices that use the native interface, Add the device to Tivoli Storage Productivity Center. Use the Configure Devices wizard to add the device. For information about the Configure Devices wizard, see "Configure Devices" on page 266. On the Select Device Type panel, select **Storage Subsystem**. After you add the storage system to Tivoli Storage Productivity Center, automatic discovery and probe jobs are run.
4. Create alerts. You can specify which alerts you want to be notified of when monitoring the storage systems.

   For storage systems that use the CIM agent, Tivoli Storage Productivity Center receives CIM indications.

   If you configure Tivoli Storage Productivity Center to send SNMP traps or Omnibus events, you must first import the Tivoli Storage Productivity Center MIB and Baroc files into your SNMP manager or Omnibus manager. For more information about configuring SNMP traps, see "Configuring Tivoli Storage Productivity Center to send SNMP traps and IBM Tivoli Enterprise Console events" on page 299.
5. To view performance reports for the storage system, create a performance monitor job. Performance data collection is performed through a performance monitor job. This job can be run immediately, can be scheduled to run one time only, or can be scheduled to run repeatedly as wanted. Only after the device

has been probed successfully is that device eligible for a performance monitor. If the probe is not successful, that device will not show up in the performance monitor creation panel.

Performance management for storage systems is done through the Disk Manager. Follow these general steps:

a. Create a performance monitor job. Data collection for the storage systems can be scheduled through **Disk Manager** > **Monitoring** > **Subsystem Performance Monitors**. Right-click **Subsystem Performance Monitors**. Click **Create Subsystem Performance Monitor**. In the content pane, enter the information schedule and run a performance monitor job.

b. Check the performance monitor status. When the performance monitor job starts to run, you begin to collect performance data for the device. Check the status of the performance monitor job to make sure that it runs and continues running. In the navigation tree, expand **IBM Tivoli Storage Productivity Center** > **Job Management**. In the content pane, click the **Refresh** button because refreshes are not automatic on the Job Management panel. See Jobs for Selected Schedule for a list of performance monitor jobs and the status of the jobs. For more information about the Job Management panel, see "Managing schedules, jobs, and runs" on page 283.

c. Specify the retention period for the performance data. After the monitor job is created, you can configure how long to keep performance data. You can specify a retention period for collected sample data, for aggregated hourly data, and for daily data. The retention period applies to all performance monitors. For more information about the retention period, see "Resource History Retention" on page 332.

   Sample data is the data that is collected at the specified interval length of the monitor job, for example, data is collected every 5 minutes.

   For storage systems, the most numerous components are usually volumes, therefore, the largest amount of performance data collected is for volumes.

d. To view the performance reports, expand **Disk Manager** > **Reporting** > **Storage Subsystem Performance**

6. You can also view the Disk Manager reports and topology in Tivoli Storage Productivity Center.

Planning for the DS8000

Use this information to plan for using the IBM DS8000 storage system with Tivoli Storage Productivity Center.

IBM System Storage DS8000 series is a high-performance, high-capacity series of disk storage that is designed to support continuous operations.

For more information about the DS8000, see the DS8000 Information Center at http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp.

Tivoli Storage Productivity Center support

The DS8000 uses the Tivoli Storage Productivity Center native interface for connecting to the device. CIM agents are no longer required for the DS8000. After you add the DS8000 to Tivoli Storage Productivity Center, an automatic discovery and probe jobs are run.

General procedure

To view the reports and topology for the DS8000, follow these general steps:

1. Open the Tivoli Storage Productivity Center user interface and add the DS8000 to Tivoli Storage Productivity Center. Click **Administrative Services** > **Data Sources** > **Storage Subsystems**. Click **Add**. This action opens the Configure Devices wizard. For information about adding the storage system, see "Configure storage subsystem connections page" on page 270.
2. After you add the DS8000, an automatic discovery and probe are performed.
3. View the reports and topology for the DS8000 in Tivoli Storage Productivity Center.

Planning for DS8000 performance

This section provides information about DS8000 performance.

Space-efficient volumes are not fully allocated when they are created, unlike the "normal" type of volumes. Their actual allocated space grows as the volume is used, until it has reached the full size of the volume. After the full size of the volume is reached, the allocated size does not change further. For performance management, the current size of the volume is unknown at any given point in time. Because the current size of the volume is unknown, it is impossible to know how the segments (extents) for the volume are distributed among the various ranks in a multi-rank extent pool. It is therefore impossible to accurately compute the performance statistics for the arrays associated with a multi-rank extent pool that contains one or more space-efficient volumes.

Planning for the SAN Volume Controller

Use this information to plan for using the SAN Volume Controller with Tivoli Storage Productivity Center.

You can view reports and the topology for your SAN Volume Controller environment. Use this information to understand the configuration that is required to use SAN Volume Controller with Tivoli Storage Productivity Center.

By using Tivoli Storage Productivity Center, you can view SAN Volume Controller performance statistics for the following reports:

- **Disk Manager > Reporting > Storage Subsystem Performance > By Storage Subsystem**
- **Disk Manager > Reporting > Storage Subsystem Performance > By I/O Group**
- **Disk Manager > Reporting > Storage Subsystem Performance > By Node**

You can also view the SAN Volume Controller environment through the topology viewer.

When you have hosts connected to storage subsystems which have multi-pathing enabled, you must have the multi-pathing subsystem device drivers (SDD) installed on the hosts. For more information, see the SAN Volume Controller Information Center and search for **multipath support**.

SAN Volume Controller Version 6.1 or later

SAN Volume Controller includes IBM System Storage Easy Tier, a function that responds to the presence of solid-state drives (SSDs) in a storage pool that also

contains hard disk drives (HDDs). The system automatically and nondisruptively moves frequently accessed data from HDD MDisks to SSD MDisks, thus placing such data in a faster tier of storage.

Easy Tier eliminates manual intervention when assigning highly active data on volumes to faster responding storage. In this dynamically tiered environment, data movement is seamless to the host application regardless of the storage tier in which the data resides.

SAN Volume Controller supports these tiers:

Generic SSD tier
The SSD tier exists when SSDs are in the storage pool. The SSDs provide greater performance than hard disk drives (HDDs).

Generic HDD tier
The HDD tier exists when HDDs are in the storage pool.

All MDisks belong to one tier or the other, which includes MDisks that are not yet part of a storage pool.

If you create a storage pool (managed disk group) with both generic SSD MDisks and generic HDD MDisks, Easy Tier is automatically turned on for pools with both SSD MDisks and HDD MDiasks. SAN Volume Controller does not automatically identify external SSD MDisks; all external MDisks are put into the HDD tier by default. You must manually identify external SSD MDisks and change their tiers.

Note:

- The terminology for this release of SAN Volume Controller has changed:
  - Error is now event
  - MDisk is now storage pool
  - Space-efficient is now thin provisioning
  - VDisk is now volume

  However, the terms have not been changed in Tivoli Storage Productivity Center reports and views.
- In SAN Volume Controller 5.1 or earlier, there was a one-to-one mapping between MDisks and internal SSD drives. For SAN Volume Controller 6.1 or later, there are now one-to-many mapping between MDisks and drives.
- For this release of Tivoli Storage Productivity Center, the performance monitor, SAN Planner, and Storage Optimizer do not support Easy Tier.

In addition to information already available for SAN Volume Controller, Tivoli Storage Productivity Center reports and topology views also provide information about the tier, tier capacity, and tier free capacity.

Tivoli Storage Productivity Center also supports launch in context and single sign-on for SAN Volume Controller.

Note: Launch in context and single sign-on is not supported for SAN Volume Controller Version 5.1 or earlier.

Tivoli Storage Productivity Center for Replication also supports SAN Volume Controller Version 6.1.

Upgrading SAN Volume Controller Version 5.1

Currently SAN Volume Controller Version 6.1 does not support the solid-state drives. This support will be added in a future fix pack.

If you are upgrading SAN Volume Controller Version 5.1 to Version 6.1, you have the following options:

- Do not upgrade SAN Volume Controller Version 5.1 to Version 6.1 until the solid-state drive support is available.
- Follow these steps to upgrade SAN Volume Controller Version 5.1:
  1. Migrate data off the solid-state drive MDisks.
  2. Unconfigure (unmanage) the solid-state drive MDisks.
  3. Run a Tivoli Storage Productivity Center for Replication probe for the SAN Volume Controller.
  4. Upgrade the SAN Volume Controller Version 5.1 to Version 6.1.

Note: SAN Volume Controller does not permit an upgrade of SAN Volume Controller Version 5.1 to Version 6.1 while there are managed solid-state drive MDisks. After the upgrade, the unmanaged solid-state drive MDisk objects are marked as unused drive objects. You cannot configure these drive objects into RAID arrays.

General procedure

To view the reports and topology for the SAN Volume Controller environment, follow these general steps:

1. Log in to the SAN Volume Controller console and create a user ID and password for use with Tivoli Storage Productivity Center.
2. Open the Tivoli Storage Productivity Center user interface and add the SAN Volume Controller to Tivoli Storage Productivity Center. Click **Administrative Services > Data Sources > Storage Subsystems**. Click **Add**. This opens the Configure Devices wizard. For information about adding the storage system, see "Configure storage subsystem connections page" on page 270.
3. Tivoli Storage Productivity Center supports both OpenSSH format or putty key (.ppk) format that is not password protected. You only need to provide the private SSH key.

   If you have a password protected putty .ppk key, you can use the PuTTY tool on Windows to convert that key to OpenSSH format by following these steps:

   a. Load the putty .ppk file into the PuTTY tool.
   b. Open the PuTTY tool.
   c. From the menu bar, select **Conversion** and select the **OpenSSH conversion**.

   This action saves the key in OpenSSH format.

   Important: If you upload a new key for the existing user ID using Tivoli Storage Productivity Center, that new key overrides the existing key. If you have another application using that SAN Volume Controller user ID (for example Tivoli Storage Productivity Center for Replication), the application cannot log into SAN Volume Controller because the key associated with the user ID is overwritten.

   If you have this situation, you should create a user ID specifically for Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication on the SVC cluster. Add the SAN Volume Controller to Tivoli

Storage Productivity Center for Replication using the new user ID and password (or change the user ID and password while in the disconnected state).

Note:

- IBM Tivoli Storage Productivity Center ships default public and private keys in OpenSSH format. The public and private keys can be found in `tpc_installation_directory/device/conf`. The file names are tpc_svc.pub (public key) and tpc_svc.pem (private key).

4. After you add the SAN Volume Controller, an automatic discovery and probe are performed.
5. Before you can view performance reports for the SAN Volume Controller, you must define and run a performance monitor job.
6. View the reports and topology for the SAN Volume Controller in Tivoli Storage Productivity Center.

Planning for Storwize V7000

Storwize V7000 is a hardware and software solution that provides unmatched performance, availability, advanced functions, and highly scalable capacity.

General concepts

Storwize V7000 offers IBM storage virtualization, SSD optimization and "thin provisioning" technologies built in to improve storage utilization. The storage system can be reconfigured to meet changing needs quickly and easily. This solution helps to reduce costs without performance degradation.

The Storwize V7000 hardware consists of a set of drive enclosures. Control enclosures contain disk drives and two node canisters. The two nodes within the canisters make an I/O group that is attached to the SAN fabric. A single pair of nodes is responsible for serving I/O on a given volume. Because a volume is served by two nodes, there is no loss of availability if one node fails or is taken offline.

Storwize V7000 can be used as a traditional RAID storage system where the internal drives are configured into arrays, and volumes are created from those arrays. Storwize V7000 can also be used to virtualize other storage systems.

Storwize V7000 supports both regular and solid-state drives (SSDs). A Storwize V7000 system without any internal drives can be used as a storage virtualization solution.

Each Storwize V7000 node has two Ethernet ports that can be used for management. Ethernet port 1 must be configured with a management IP address and must be connected on all nodes in the system. The use of Ethernet port 2 is optional. At any point in time, only one node in the system can operate as the focal point for configuration and monitoring requests. This node is called the configuration node and is the only node that activates the management IP addresses.

Each Storwize V7000 can have zero to four management IP addresses. You can assign up to two IPv4 addresses and up to two IPv6 addresses.

All configuration, monitoring, and service tasks are performed at the cluster level (a cluster is a Storwize V7000 system that consists of two nodes).

Storwize V7000 includes IBM System Storage Easy Tier, a function that supports solid-state drives (SSDs) in a storage pool that also contains hard disk drives (HDDs). The system automatically and nondisruptively moves frequently accessed data from HDD MDisks to SSD MDisks, thus placing such data in a faster tier of storage.

Easy Tier eliminates manual intervention when assigning highly active data on volumes to faster responding storage. In this dynamically tiered environment, data movement is seamless to the host application regardless of the storage tier in which the data resides. Manual controls exist so that you can change the default behavior, for example, such as turning off Easy Tier on storage pools that have both types of MDisks.

Using Tivoli Storage Productivity Center to monitor Storwize V7000

Tivoli Storage Productivity Center communicates with Storwize V7000 storage systems through a native interface. No CIM agents are required.

Use Tivoli Storage Productivity Center to monitor Storwize V7000 systems, including discovery, data collection (probes), performance monitoring, trending, storage optimization, and SAN planning. Tivoli Storage Productivity Center reports and topology views also provide information about the tier, tier capacity, and tier free capacity for Storwize V7000.

Note: For this release of Tivoli Storage Productivity Center, the performance monitor, SAN Planner, and Storage Optimizer do not support Easy Tier.

Similar to other systems, Storwize V7000 appears as an entity within data sources, reports, data collection schedules, the topology viewer, and so on. Tivoli Storage Productivity Center also supports launch in context and single sign-on for Storwize V7000.

General procedure

To view the reports and topology for the Storwize V7000 environment, follow these general steps:

1. Log in to the Storwize V7000 console and create a user ID and password for use with Tivoli Storage Productivity Center.
2. Open the Tivoli Storage Productivity Center user interface and add the Storwize V7000 to Tivoli Storage Productivity Center. Click **Administrative Services > Data Sources > Storage Subsystems**. Click **Add**. This opens the Configure Devices wizard. For information about adding the storage system, see “Configure storage subsystem connections page” on page 270.
3. Tivoli Storage Productivity Center supports both OpenSSH format or putty key (.ppk) format that is not password protected. You only need to provide the private SSH key.

   If you have a password protected putty .ppk key, you can use the PuTTY tool on Windows to convert that key to OpenSSH format by following these steps:

   a. Load the putty .ppk file into the PuTTY tool.
   b. Open the PuTTY tool.

c. From the menu bar, select **Conversion** and select the **OpenSSH conversion**.

This action saves the key in OpenSSH format.

Important: If you upload a new key for the existing user ID using Tivoli Storage Productivity Center, that new key overrides the existing key. If you have another application using that Storwize V7000 user ID (for example Tivoli Storage Productivity Center for Replication), the application cannot log into Storwize V7000 because the key associated with the user ID is overwritten.

If you have this situation, you should create a user ID specifically for Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication on the Storwize V7000 cluster. Add the Storwize V7000 to Tivoli Storage Productivity Center for Replication using the new user ID and password (or change the user ID and password while in the disconnected state).

Note:

- IBM Tivoli Storage Productivity Center ships default public and private keys in OpenSSH format. The public and private keys can be found in `tpc_installation_directory/device/conf`. The file names are tpc_svc.pub (public key) and tpc_svc.pem (private key).

4. After you add the Storwize V7000, an automatic discovery and probe are performed.
5. Before you can view performance reports for the Storwize V7000, you must define and run a performance monitor.
6. View the reports and topology for the Storwize V7000 in Tivoli Storage Productivity Center.

Planning for Storwize V7000 Unified

The IBM Storwize V7000 Unified system is a virtualizing redundant array of independent disks (RAID) storage system that supports both block protocols and file protocols. This unified system includes Storwize V7000 File Module and the Storwize V7000 storage system.

General concepts

Storwize V7000 offers IBM storage virtualization, SSD optimization and "thin provisioning" technologies built in to improve storage utilization. The storage system can be reconfigured to meet changing needs quickly and easily. This solution helps to reduce costs without performance degradation.

The Storwize V7000 hardware consists of a set of drive enclosures. Control enclosures contain disk drives and two node canisters. The two nodes within the canisters make an I/O group that is attached to the SAN fabric. A single pair of nodes is responsible for serving I/O on a specific volume. Because a volume is served by two nodes, there is no loss of availability if one node fails or is taken offline.

Storwize V7000 File Module is a clustered system that consists of two units that provide file systems for use by network-attached storage. The file module uses the Storwize V7000 storage system to provide the file module with volumes. Other volumes, which are block volumes, are provided on the SAN to be presented to hosts.

The Storwize V7000 File Module software within the Storwize V7000 Unified system contains the interface node and management node functions. The management node function provides for configuring, administering, and monitoring a system. The interface node function connects a system to an Internet Protocol (IP) network. You can connect by using any of the following protocols:

- Common Internet File System (CIFS)
- Network File System (NFS)
- File Transfer Protocol (FTP)
- Hypertext Transfer Protocol Secure (HTTPS)
- Secure Copy Protocol (SCP)

The Storwize V7000 Unified software also performs the following functions for the host systems that attach to the Storwize V7000 Unified system by using block protocols:

- Creates a single pool of storage
- Provides logical unit virtualization
- Manages logical volumes
- Mirrors logical volumes
- Large scalable cache
- Copy Services
  - IBM FlashCopy (point-in-time copy) function, including thin-provisioned FlashCopy to make multiple targets affordable
  - Metro Mirror (synchronous copy)
  - Global Mirror (asynchronous copy)
  - Data migration
- Space management
  - IBM System Storage® Easy Tier™ to migrate the most frequently used data to higher performing storage
  - Metering of service quality when combined with IBM Tivoli Storage Productivity Center
  - Thin-provisioned logical volumes

Storwize V7000 can be used as a traditional RAID storage system where the internal drives are configured into arrays, and volumes are created from those arrays. Storwize V7000 can also be used to virtualize other storage systems.

Storwize V7000 supports both regular and solid-state drives (SSDs). A Storwize V7000 system without any internal drives can be used as a storage virtualization solution.

Each Storwize V7000 node has two Ethernet ports that can be used for management. Ethernet port 1 must be configured with a management IP address and must be connected on all nodes in the system. The use of Ethernet port 2 is optional. At any point in time, only one node in the system can operate as the focal point for configuration and monitoring requests. This node is called the configuration node and is the only node that activates the management IP addresses.

Each Storwize V7000 can have zero to four management IP addresses. You can assign up to two IPv4 addresses and up to two IPv6 addresses.

All configuration, monitoring, and service tasks are performed at the cluster level (a cluster is a Storwize V7000 system that consists of two nodes).

Storwize V7000 includes IBM System Storage Easy Tier, a function that supports solid-state drives (SSDs) in a storage pool that also contains hard disk drives (HDDs). The system automatically and nondisruptively moves frequently accessed data from HDD MDisks to SSD MDisks, thus placing such data in a faster tier of storage.

Easy Tier eliminates manual intervention when assigning highly active data on volumes to faster responding storage. In this dynamically tiered environment, data movement is seamless to the host application regardless of the storage tier in which the data resides. Manual controls exist so that you can change the default behavior, for example, such as turning off Easy Tier on storage pools that have both types of MDisks.

Using Tivoli Storage Productivity Center to monitor Storwize V7000 Unified

Tivoli Storage Productivity Center communicates with Storwize V7000 Unified storage systems through a native interface. No CIM agents are required.

Use Tivoli Storage Productivity Center to monitor Storwize V7000 Unified systems, including discovery, data collection (probes), performance monitoring, and trending. Tivoli Storage Productivity Center reports and topology views also provide information about the tier, tier capacity, and tier free capacity for Storwize V7000 Unified.

Similar to other systems, Storwize V7000 Unified appears as an entity within data sources, reports, data collection schedules, the topology viewer, and so on. Tivoli Storage Productivity Center also supports launch in context and single sign-on for Storwize V7000 Unified.

General procedure

To view the reports and topology for the Storwize V7000 Unified environment, follow these general steps:

1. Log on to the Storwize V7000 Unified console and create a user ID and password for use with Tivoli Storage Productivity Center.
2. Open the Tivoli Storage Productivity Center user interface and add the Storwize V7000 Unified to Tivoli Storage Productivity Center. Click **Administrative Services > Data Sources > Storage Subsystems**. Click **Add**. This opens the Configure Devices wizard. For information about adding the storage system, see “Configure storage subsystem connections page” on page 270.
3. Tivoli Storage Productivity Center supports Secure Shell (SSH) keys that are in OpenSSH format or 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:

   a. Start the puttygen utility.
   b. Load the `.ppk` file. Click **File > Load private key**.
   c. From the menu bar, click **Conversion > Export OpenSSH key conversion**.

   This action saves the key in OpenSSH format.

Important: If you upload a new key for the existing user ID using Tivoli Storage Productivity Center, that new key overrides the existing key. If you have another application using that Storwize V7000 Unified user ID (for example Tivoli Storage Productivity Center for Replication), the application cannot log on to Storwize V7000 Unified because the key associated with the user ID is overwritten.

If you have this situation, create a user ID specifically for Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication on the Storwize V7000 Unified cluster. Add the Storwize V7000 Unified to Tivoli Storage Productivity Center for Replication using the new user ID and password (or change the user ID and password while in the disconnected state).

Note:

- Tivoli Storage Productivity Center ships default public and private keys in OpenSSH format. The public and private keys can be found in `tpc_installation_directory/device/conf`. The file names are `tpc_svc.pub` (public key) and `tpc_svc.pem` (private key).

4. After you add the Storwize V7000 Unified, an automatic discovery and probe are performed.
5. Before you can view performance reports for the Storwize V7000 Unified, you must define and run a performance monitor.
6. View the reports and topology for the Storwize V7000 Unified in Tivoli Storage Productivity Center.

To scan a Storwize V7000 Unified file system, follow these general 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 which has a Storage Resource agent 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. This is done so that a situation cannot occur where the product counts 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 located 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. Open the Tivoli Storage Productivity Center GUI and click **Administrative Services** > **Configuration** > **Scan/Probe Agent Administration**. Save the page.

5. A scan job for the Storwize V7000 Unified can be created and run from **Data Manager** > **Monitoring** > **Scans**.

Identifying Storwize V7000 Unified devices in the user interface

The way you configure a Storwize V7000 Unified device determines how the Storwize V7000 Unified device is represented in the Tivoli Storage Productivity Center user interface, as follows:

- Storwize V7000 File Module generally appears as a "computer" in the user interface. In **Data Manager** > **Alerting** > **Other NAS Alerts** > **NAS**, it appears as a "NAS cluster."
- Storwize V7000 storage system appears as a "subsystem" in the user interface.

Planning for the XIV system

Use this information to plan for using the XIV system with Tivoli Storage Productivity Center.

IBM XIV Storage System is a disk storage architecture designed to eliminate the complexity of administration and management of storage. The XIV system parallelized architecture of the system, optimal exploitation of all system components (including disks, processors, and switches), and unique caching architecture all translate into excellent performance.

The unique balancing of all data across system components prevents the occurrence of hot spots. With all components working under the same load, performance and reliability are exceptional.

The XIV system uses large capacity Serial Advanced Technology Attachment (SATA) disk drives which optimizes the use of disk capacity, resulting in outstanding power consumption without compromising performance.

The XIV system is designed to be scalable in storage, interfaces, cache, CPU power, and internal bandwidth. The architecture supports each aspect to grow independently, resulting in a scalable system in both capacity and performance.

The XIV system provides data protection and availability. All disk drives, modules, switches, and uninterruptible power supply units are fully redundant, ensuring high reliability and excellent performance.

The built-in thin provisioning of the XIV system helps reduce direct and indirect costs by allowing users to install capacity only for data written. You can grow your data capacity over time with minimal management effort.

For more information about the XIV system, see the XIV system Information Center.

Tivoli Storage Productivity Center support

Tivoli Storage Productivity Center uses the native interface for the XIV system to run discovery and probe jobs. Because of the dynamic XIV architecture, where volume data is spread over all disks, no data that describes the space allocation from volumes to a specific physical disk is available.

The XIV system includes a graphical user interface (GUI) for configuration and administration (XIV Storage Manager). XIV Storage Manager is the element manager for XIV system. XIV Storage Manager version 4.2 or later must be installed on the same computer as the Tivoli Storage Productivity Center GUI. If XIV Storage Manager is not installed on the same computer as the Tivoli Storage Productivity Center GUI or is configured incorrectly, an error message is displayed when you attempt to start XIV Storage Manager from Tivoli Storage Productivity Center.

The XIV system can define up to three administrative nodes, each with their own IP address. In Tivoli Storage Productivity Center 4.2.1 or later, if the XIV system is configured with multiple administrative nodes, Tivoli Storage Productivity Center detects the IP addresses for these nodes. If Tivoli Storage Productivity Center fails to connect to one of the IP addresses, then an attempt is made to connect to the XIV system using one of the other IP addresses.

Note: The XIV system does not support IPv6.

Considerations for data collected for the XIV system

The following considerations apply to values shown in the XIV system and Tivoli Storage Productivity Center GUIs.

- The XIV Storage System GUI and the XIV Storage 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 Storage 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.
- XIV system differentiates thin provisioned pools from non-thin provisioned (regular) pools and volumes in a thin provisioned pool and a non-thin provisioned pool. When hard size is equal to soft size, the pool is represented as a regular pool, otherwise the pool is represented as thin provisioned. However, the XIV system CLI and CIM do not distinguish between regular and thin provisioned pools and all pool volume names in Tivoli Storage Productivity Center are prefixed with an asterisk (*) to indicate that the volume is thin provisioned.

General procedure

To view the reports and topology for the XIV system environment, follow these general steps:

1. Open the Tivoli Storage Productivity Center user interface and add the XIV system to Tivoli Storage Productivity Center. In the navigation tree, expand **IBM Tivoli Storage Productivity Center** > **Configure Devices**. This opens the Configure Devices wizard. For information about adding the storage system, see “Configure storage subsystem connections page” on page 270.
2. After you add the XIV system, an automatic discovery and probe are performed.
3. View the reports and topology for the XIV system in Tivoli Storage Productivity Center.

Planning for the Fabric Manager

This topic provides planning information for the Fabric Manager.

The main features of Fabric Manager are shown in the following table.

Feature	Advantages	Benefit
Automatic device discovery	Enables you to see the data path from the servers to the switches and storage systems	Easily manage your SANs from a single console
Real-time monitoring and alerts	Monitors SAN events and alerts administrators of problems	Helps maintain the availability of your SAN
Zone control	Gives you the ability to add and remove devices from zones	Comprehensive fabric management from a single console
SAN reporting	Provides the ability to generate reports on the SAN devices	Quickly provides reports and inventories of your SAN
Enterprise scalability	Scale from SAN islands to enterprise SANs	Accommodates new and ever changing business needs
Switch performance management (1)	Provides the ability to monitor switches for performance metrics which allows for the triggering of switch threshold alerts and performance reporting. The performance reporting is at the switch and port levels and displays in the defined Switch Performance reports and in the topology viewer.	Helps you maintain high SAN availability

Note:

1. Tivoli Storage Productivity Center provides predictive fault analytics for the SAN infrastructure through reporting, monitoring, and alerting of link failure rates and error frame rates. These transient errors are indicative of potential pending SAN link failures.

Most heterogeneous fabrics, which are composed of switches from two or more vendors, are not supported. Tivoli Storage Productivity Center supports heterogeneous fabrics that are comprised entirely of McDATA and Brocade switches, when CIM agents are configured for these switches. All other heterogeneous fabrics are not supported. The following are known issues which prevent Tivoli Storage Productivity Center from supporting heterogeneous fabrics (except for Brocade and McDATA heterogeneous fabrics):

- If there is no CIM agent for a heterogeneous fabric, Tivoli Storage Productivity Center attempts to perform zone control operations through an available in-band Fabric agent or Storage Resource agent for the fabric.
- Some zone member types that can be chosen for use in zone control operations causes zone control operations to fail in heterogeneous fabrics. This behavior is a result of switch vendors providing only limited support for zone member types in heterogeneous fabrics.

- Zone control capabilities of the fabric are assumed by Tivoli Storage Productivity Center to be the capabilities of the principal switch vendor. However, the agent used for zone control operations might be for a switch vendor that differs from the principal switch in heterogeneous fabrics. For this reason, zoning operations such as creation of "orphan" zones and use of non-alphanumeric characters in zoning entity names might fail at the switch rather than be caught by Tivoli Storage Productivity Center.
- Heterogeneous fabrics with Cisco might be represented with duplicate fabrics in the topology viewer if only CIM agent data sources are configured.

Planning for Brocade or McData fabric management agents

IBM Tivoli Storage Productivity Center uses a Storage Management Initiative (SMI) agent to manage Brocade and McData switches. If you are using Brocade Data Center Fabric Manager (DCFM) 10.4 or later (which manages Brocade and McData switches) the SMI agent is embedded in the DCFM. If you are not using Brocade DCFM 10.4 or later, you must use external SMI agents for Brocade and McData switches.

Supported Brocade and McData models

For the Brocade and McData models, firmware versions, and SMI agents that are supported by Tivoli Storage Productivity Center, see the support matrix at http://www-01.ibm.com/support/docview.wss?rs=40&q1=switch&uid=swg21384219&loc=en_US&cs=utf-8&lang=en.

Discovering fabrics and switches that use an SMI agent

When you configure Tivoli Storage Productivity Center to connect to the SMI agent (embedded or external), you can run a CIM object manager (CIMOM) discovery job. The discovery job discovers the fabric and switches and the status of the fabrics and switches. The fabric can be a Fabric OS (FOS) fabric, Enterprise OS (EOS) fabric, or a heterogeneous fabric (FOS and EOS fabrics).

When you run a probe job using the SMI agent, you can see the topology of the fabric, the properties of the hardware that make up the fabric, zoning information, and the endpoint devices. Zone aliases are supported for only for FOS fabrics.

When you run a fabric performance monitor job, you can see the performance data for the fabric.

Using other agents for discovery: legacy in-band fabric agents, out-of-band SNMP agents, or Storage Resource agents

The use of SMI agents provides the most benefits for fabric management, including:

- Probes and zone control do not occur within the data path.
- Zone alias control is supported by Tivoli Storage Productivity Center through the fabric SMI agents only.
- Propagation of alerts for real-time SAN events are not degraded by the use of fabric SMI agents.

However, the use of multiple agent types also has advantages, such as redundant collection of information in the event that one type of agent fails. Also, some types

of agents provide certain features the other agents do not. However, the use of multiple agent types can increase network traffic as well as management load on the switches.

The following table outlines the fabric management functions that are supported by agent type. Where multiple agents can be used, the preferred agent type is noted.

Table 23. Supported and preferred interfaces for fabric management

Function	Brocade	McData
Performance monitoring	Fabric SMI agent	Fabric SMI agent
Topology information collected	Preferred: Fabric SMI agent Still supported: Storage Resource agent or In-band fabric agent that uses FC-GS-3 Out-of-band SNMP agent	Preferred: Fabric SMI agent Still supported: Storage Resource agent or In-band fabric agent that uses FC-GS-3 Out-of-band SNMP agent
Zone information collected	Fabric SMI agent	Preferred: Fabric SMI agent Still supported: Storage Resource agent or In-band fabric agent that uses FC-GS-3
Zone control	Fabric SMI agent	Preferred: Fabric SMI agent Still supported: Storage Resource agent or In-band fabric agent that uses FC-GS-3
Zone control with zone aliases	Fabric SMI agent	Not supported
Switch and switch port information collected (for running performance monitor)	Preferred: Fabric SMI agent	Preferred: Fabric SMI agent
Fabric and switch events (see note below)	Preferred: Fabric SMI agent Switch SNMP trap configuration	Preferred: Fabric SMI agent Switch SNMP trap configuration
Launch of switch and fabric element managers	Preferred: Fabric SMI agent Still supported: Storage Resource agent or In-band fabric agent that uses FC-GS-3 Out-of-band SNMP agent	Preferred: Fabric SMI agent Still supported: Storage Resource agent or In-band fabric agent that uses FC-GS-3 Out-of-band SNMP agent

Table 23. Supported and preferred interfaces for fabric management (continued)

Function	Brocade	McData
Hosts, endpoint devices, device-centric and host-centric information collected	Storage Resource agent or In-band fabric agent that uses FC-GS-3	Storage Resource agent or In-band fabric agent that uses FC-GS-3
Switch sensors and events	Out-of-band SNMP agent	Out-of-band SNMP agent

Notes:

- The Storage Resource agent has two distinct functions:
  - **Discovery**

    When a discovery job is run, Tivoli Storage Productivity Center gathers information about the fabrics, fabric to switch relationships, and other key attributes about the switches in the fabrics.
  - **Probe**

    When a probe job is run, Tivoli Storage Productivity Center gathers additional details for topology and zoning information.

  If you deploy a Storage Resource agent and a fabric probe job has not been run, the connections from the Storage Resource agents to the switches are not displayed, and all the fibre channel (FC) ports are reported as unused.
- Zone information collected indicates that the active zone set and the inactive zoning library information is collected during a fabric probe. Zone control refers to the ability to actively configure zoning.
- Tivoli Storage Productivity Center does not get a transaction lock on the switch while making zone changes on the fabric.
- For fabric and switch events, the Storage Resource agent relies on both SNMP traps and CIM indications for event notification and automatic fabric discovery and probe, if needed.

  The preferred method includes both SMI agent and switch SNMP trap configuration. An out-of-band SNMP agent is not necessary for getting SNMP traps. You can configure the switch to send SNMP traps to the Device server. However, you can configure the out-of-band SNMP agent for other reasons.

Planning for embedded SMI agent support

An embedded Storage Management Initiative (SMI) agent is provided in Brocade Data Center Fabric Manager (DCFM) 10.4 or later. Brocade DCFM manages both the McData and Brocade switches. DCFM manages multiple fabrics within and across data centers. When you configure DCFM, you set up one switch to be the *seed* or master switch that interconnects to all the other switches in the fabric.

Configuring the Brocade DCFM

Before Tivoli Storage Productivity Center can discover the switches in a fabric through the embedded SMI agent, the Brocade DCFM must be configured to manage the fabric. This configuration is done through the seed switch. For heterogeneous Fabric OS (FOS) and Enterprise OS (EOS) fabrics, the seed switch must be one of the FOS switches. There are special requirements for the seed switch:

- For pure FOS fabrics, the seed switch must have a firmware of 5.3 or later
- For heterogeneous FOS and EOS fabrics, the seed switch must have a firmware of 6.0 or later. For pure M-EOS fabrics, the seed switch must be running M-EOS 9.6.X or later.

For more information about configuring the Brocade switch, see *Brocade SMI Agent for EOS Products User Guide* and *DCFM Migration Guide*.

Migrating from a non-embedded SMI agent to an embedded SMI agent

You can configure Tivoli Storage Productivity Center to manage the fabric in the following cases:

- Configure Tivoli Storage Productivity Center to connect to the embedded SMI agent in place of a non-embedded Brocade FOS SMI agent for managing the same fabric.
- Configure Tivoli Storage Productivity Center to connect to the embedded SMI agent in place of a non-embedded Brocade EOS SMI agent for managing the same fabric.
- Configure Tivoli Storage Productivity Center to connect to the embedded SMI agent in place of the non-embedded Brocade EOS and FOS SMI agents for managing the same fabric.

Migrating from a legacy agent to a SMI agent

You can configure Tivoli Storage Productivity Center to use the embedded SMI agent instead of the out-of-band Fabric agent or in-band Fabric agent for the EOS fabric or FOS fabric.

Limitations

These are some limitations using the embedded SMI agent:

- The embedded SMI agent has asset collection period logic for handling fabric events. If anything in the fabric topology changes and is restored within the asset collection period, the SMI agent does not send a CIM indication. In this situation, you must run a fabric probe to collect the most updated fabric topology information. For more information about the asset collection period, including the value for the collection period, refer to the DCFM documentation.
- The embedded SMI agent is not supported on IBM AIX.
- Zone control sessions through the SMI agent do not establish a session with the seed switch until the session from Tivoli Storage Productivity Center is committed.
- DCFM Professional Edition does not have the embedded SMI agent.
- DCFM supports either https or http protocol at one time (both protocols cannot be supported at the same time).
- Empty zones, zone sets, and zone aliases are not supported through FOS or EOS fabrics through the SMI agent.
- When you set up the seed switch, if the switch link is broken anywhere along the connection line, Tivoli Storage Productivity Center will not be able to report on the switches past the broken link.
- When a Brocade i10000 or McData i10000 switch has a faulty control processor card, the SMI agent managing the switch can hang indefinitely. This issue causes time-outs and probe failures from the Tivoli Storage Productivity Center server.

Planning for external SMI agent support

Brocade and McData switches that do not use the Brocade Data Center Fabric Manager (DCFM) 10.4.0 or later require a separate external Storage Management Initiative (SMI) agent for the Brocade and McData hardware. The external SMI agents must be installed on separate hosts and can manage one or more switches.

Installing and Configuring external SMI agents

You need to install and configure the external SMI agents before adding the SMI agents to IBM Tivoli Storage Productivity Center.

For more information about configuring the external SMI agents for Brocade and McData switches, refer to the documentation for the switch.

Considerations

These are some considerations for using external SMI agents:

- The McData external SMI agent can be configured to run in *proxy* mode or in *direct* mode.
  - In proxy mode, the SMI agent communicates through the Enterprise Fabric Connectivity Manager (EFCM) Server or the DCFM Server (versions earlier than 10.4.0). The zoning database visible from the SMI agent is synchronized with the zoning database stored on the EFCM or DCFM Server.
  - In direct mode, the SMI agent communicates directly with the switches and stores inactive zone sets on the SMI agent.

    ECFM and DCFM cannot monitor switches that are configured for direct connection with the SMI agent.
- Brocade fabrics that are not laid out as a *mesh*, where one inter-switch link (ISL) break can segment the fabric, manage the fabric with both a fabric SMI agent and an out-of-band Fabric agent (SNMP). The reason for this is that Brocade SMI agents do not automatically manage newly created fabrics in the event of a fabric segmentation.
- Brocade SMI agents do not correctly report whether switches and fabrics are running in Interop Mode or Native Mode. Because of this limitation, Tivoli Storage Productivity Center is unable to determine whether a fabric or switch is in Interop Mode or Native Mode. For homogeneous Brocade fabrics that are configured in Interop Mode, the following actions from Tivoli Storage Productivity Center will fail:
  - Zone control with domain/port members and node WWN members
  - Creating zone aliases for Brocade Fabric OS (FOS) fabrics

Planning for QLogic or Cisco fabric management agents

Depending on the function that you want to perform (topology discovery, zone discovery, performance monitoring, and so forth), IBM Tivoli Storage Productivity Center uses a Storage Management Initiative (SMI) agent, legacy in-band fabric agent, out-of-band Simple Network Management Protocol (SNMP) agent, or Storage Resource agent to manage QLogic and Cisco switches.

Supported QLogic and Cisco models

For the QLogic and Cisco models, firmware versions, and agents that are supported by Tivoli Storage Productivity Center, see the support matrix at http://www-01.ibm.com/support/docview.wss?uid=swg27019378.

The following table outlines the fabric management functions that are supported by agent type.

Table 24. Supported and preferred interfaces for fabric management

Function	QLogic	Cisco
Performance monitoring	Fabric SMI agent	Fabric SMI agent
Topology information collected	Storage Resource agent or In-band fabric agent that uses FC-GS-3 Out-of-band SNMP agent	Storage Resource agent or In-band fabric agent that uses FC-GS-3 Out-of-band SNMP agent SNMP is needed for physical infrastructure discovery.
Zone information collected	Storage Resource agent or In-band fabric agent that uses FC-GS-3	Storage Resource agent or In-band fabric agent that uses FC-GS-3 required in each VSAN
Zone control	Storage Resource agent or In-band fabric agent that uses FC-GS-3	Storage Resource agent or In-band fabric agent that uses FC-GS-3 required in each VSAN
Zone control with zone aliases	Not supported	Not supported
Switch and switch port information collected (for running performance monitor)	Preferred: Fabric SMI agent	Preferred: Fabric SMI agent
Fabric and switch events (see note below)	In-band fabric agent Out-of-band SNMP agent	In-band fabric agent Out-of-band SNMP agent
Launch of switch and fabric element managers	Storage Resource agent or In-band fabric agent that uses FC-GS-3 Out-of-band SNMP agent	Storage Resource agent or in-band fabric agent that uses FC-GS-3 Out-of-band SNMP agent
Hosts, endpoint devices, device-centric and host-centric information collected	Storage Resource agent or In-band fabric agent that uses FC-GS-3	Storage Resource agent or In-band fabric agent that uses FC-GS-3
Switch sensors and events	Out-of-band SNMP agent	Out-of-band SNMP agent

Table 24. Supported and preferred interfaces for fabric management (continued)

Function	QLogic	Cisco
Notes: 1. The Storage Resource agent has two distinct functions: • **Discovery** When a discovery job is run, Tivoli Storage Productivity Center gathers information about the fabrics, fabric to switch relationships, and other key attributes about the switches in the fabrics. • **Probe** When a probe job is run, Tivoli Storage Productivity Center gathers additional details for topology and zoning information. If you deploy a Storage Resource agent and a fabric probe job has not been run, the connections from the Storage Resource agents to the switches are not displayed, and all the fibre channel (FC) ports are reported as unused. 2. Zone information collected indicates that the active zone set and the inactive zoning library information is collected during a fabric probe. Zone control refers to the ability to actively configure zoning. 3. Tivoli Storage Productivity Center does not get a transaction lock on the switch while making zone changes on the fabric. 4. The following notes are specific to Cisco: • Each virtual storage area network (VSAN) is viewed as an individual storage area network (SAN). • The out-of-band SNMP agent and SMI agent gets the physical fabric information and can correlate the VSAN information to a physical infrastructure. • The out-of-band SNMP agent also collects some VSAN information. • The GS-4 function is not supported by the Storage Resource agent for Cisco. • For many Cisco proprietary zone members, the values are displayed as zeroes. • If a zoning configuration has Cisco proprietary zones, then you will not be able to modify the zoning for that configuration through Tivoli Storage Productivity Center. You must use the element manager to remove proprietary zone members or to perform zoning. • No aliases for Cisco are reported. • Orphan Cisco zones and aliases are not reported. 5. For fabric and switch events, the Storage Resource agent relies on both SNMP traps and CIM indications for event notification and automatic fabric discovery and probe, if needed. The preferred method includes both SMI agent and switch SNMP trap configuration. An out-of-band SNMP agent is not necessary for getting SNMP traps. You can configure the switch to send SNMP traps to the Device server. However, you can configure the out-of-band SNMP agent for other reasons.		

Discovering fabrics and switches

The type of discovery that perform for fabrics and switches depends on the agent that you are using. For information about discovery by agent type, refer to the topic about discovering storage resources in the IBM Tivoli Storage Productivity Center User's Guide .

Collecting data with the Fabric Manager

IBM Tivoli Storage Productivity Center uses different types of agents to gather data about switches. Different combinations of these agents are required to effectively enable the functions of Fabric Manager. In addition, the topology viewer is greatly

affected by the proper discovery of all the managed entities in the management scope of Tivoli Storage Productivity Center.

The following table shows the actions that trigger data collection and the data that is collected depending on the type of agent that is used.

Table 25. Information collected by agent

Action that triggers data collection	Storage Resource agent or in-band fabric agent	Out-of-band fabric agent (SNMP)	SMI agent (Brocade and McData)	SMI agent (Cisco and QLogic)
Administrative Services > Discovery > Out of Band Fabric	–	• Discovers switches and connections for switches with IP addresses in a specified range (if you selected this option) • Discovers switches and connections for switches that were previously added as out-of-band agents (either manually or switches with IP addresses in a specified range for which an out-of-band discovery was previously run) • Retrieves Cisco VSAN and physical infrastructure and correlates VSAN to physical infrastructure (see note 2 on page 108)	–	
Administrative Services > Discovery > CIMOM	–	–	• Discovers Storage Management Initiative (SMI) agents through Service Location Protocol (SLP) • Retrieves switch information • Retrieves fabric information	• Discovers SMI agents through SLP • Retrieves switch information • Retrieves fabric information
IBM Tivoli Storage Productivity Center > Probe (including fabric) (see notes 3 on page 108 and 10 on page 109)	• Retrieves zoning information, excluding zone alias information, for QLogic, Cisco, and McDATA • Retrieves switch and topology information • Retrieves attribute information for hosts and devices connected to the fabric	• Retrieves switch and topology information • Retrieves Cisco VSAN and physical infrastructure and correlates VSAN to physical infrastructure (see note 2 on page 108)	• Retrieves detailed fabric, switch, switch port, and blade information • Retrieves zoning information, including zone aliases for Brocade devices • Retrieves fabric topology connectivity information • Subscribes to CIM indications	• Retrieves detailed fabric, switch, switch port, and blade information

Table 25. Information collected by agent (continued)

Action that triggers data collection	Storage Resource agent or in-band fabric agent	Out-of-band fabric agent (SNMP)	SMI agent (Brocade and McData)	SMI agent (Cisco and QLogic)
Action of Storage Resource agent or in-band fabric agent (see note 1): • Turns on • Connectivity changes • Change event detected	• Retrieves zoning information, excluding zone alias information • Retrieves switch and topology information • Retrieves attribute information for hosts and devices connected to the fabric	• Retrieves switch and topology information • Retrieves Cisco VSAN and physical infrastructure and correlates VSAN to physical infrastructure (see note 2 on page 108)		
SNMP trap received from switch (see note 7 on page 108)	• Retrieves zoning information, excluding zone alias information • Retrieves switch and topology information • Retrieves attribute information for hosts and devices connected to the fabric	• Retrieves switch and topology information • Retrieves Cisco VSAN and physical infrastructure and correlates VSAN to physical infrastructure (see note 2 on page 108)		
CIM indication received from fabric SMI agent (see note 8 on page 108)			Performs a mini-probe to collect information that is relevant to the indication received, or sets entities as missing. Might do one of the following actions: • Retrieves switches and fabrics, or • Retrieves zoning data, or • Retrieves switch port status and connection to node, or • Retrieves switch blade status and all port connections for the blade, or • Sets fabric, switch, blades, connections, or nodes as missing	

Table notes:

1. The Storage Resource agent has two distinct functions:
   - **Discovery**

When a discovery job is run, Tivoli Storage Productivity Center gathers information about the fabrics, fabric to switch relationships, and other key attributes about the switches in the fabrics.

- **Probe**

  When a probe job is run, Tivoli Storage Productivity Center gathers additional details for topology and zoning information.

If you deploy a Storage Resource agent and a fabric probe job has not been run, the connections from the Storage Resource agents to the switches are not displayed, and all the fibre channel (FC) ports are reported as unused

2. SNMP traps received result in a fabric probe of fabrics for the switches (for which traps were received). Agent assignment determines the agents that need to be used for that probe. Event notifications from the event agent could result in any of the agent types to be used (in-band agent, Storage Resource agent, SNMP agent, and SMI agent) for the probe.
3. A fabric probe always results in a fabric discovery using Storage Resource agents and SMIS agents that can report on fabrics been probed.
4. With agent assignment, all Storage Resource agents capable of reporting on a specific fabric might not be run if Tivoli Storage Productivity Center does not deem it necessary to do so.
5. For Storage Resource agents:
   - Starting a storage resource agent will result in fabric discovery been performed using that storage resource agent. If the fabrics discovered are part of some fabric probe definition, then fabric probe will also be performed. Agent assignment is used to determine the actual agents needed for that fabric probe.
   - Stopping a storage resource agent does not cause any data collection
6. For in-band Fabric agents:
   - Starting an in-band Fabric results in a fabric probe being performed for all fabrics that the agent can report on. The same agent (which was started) is used for the fabric probe.
   - The in-band Fabric agent sending events in response to some fabric event notification results in a fabric probe of fabrics for which the events were received. Agent assignment determines the agents that need to be used for that probe. So, such event notifications from the in-band Fabric agent could result in any of the agent types being used (in-band Fabric agent, Storage Resource agent, SNMP agent, and SMIS agent).
7. SNMP traps result in a fabric probe of fabrics for which the traps were received. Agent assignment determines the agents that need to be used for that probe. So, SNMP traps could result in any of the agent types being used (in-band Fabric agent, Storage Resource agent, SNMP agent, and SMIS agent).
8. A CIM indication results in a Storage Resource agent fabric discovery for Storage Resource agents that can report on fabrics for which CIM indications were received. Also, if a mini-probe fails and a Storage Resource agent is capable of reporting on the fabric for which the indication is received, the Storage Resource agent is run to collect fabric information.
9. Starting or restarting a device or the Data server or both can result in a fabric probe of fabrics which are part of some probe definition defined by the user. Agent assignment determines the agents that must be used for that probe. This could result in any of the agent types being used (in-band Fabric agent, Storage Resource agent, SNMP agent, and SMIS agent).

Deleting or uninstalling a Storage Resource agent can result in agent assignment attempting to use an alternate agent to probe fabrics that the uninstalled Storage Resource agent could report on if such fabrics were part of a probe definition defined by the user.

10. If a fabric is in the process of rebuilding itself, a fabric probe job could fail. The fabric rebuilds itself if you add or remove a switch, or if you change the principal switch. If you experience a fabric probe job failure, suspect recent changes to the fabric which might have caused the fabric to rebuild. After the fabric rebuilds itself, the fabric probe job will likely succeed. Consider implementing a change-control process for the storage environment so that the previously mentioned changes are planned, approved, and recorded.
11. Your fabric probe or zone control job can fail on AIX with error message: AGT0430I.

    AIX has a constraint wherein the size of the response buffer passed in for GS-3 commands made to the switch cannot exceed 4 K. The Storage Resource agent passes 4 K buffer sizes for AIX.

    If the zone configuration for a fabric probe or zone control job exceeds the 4 K limitation for AIX, Tivoli Storage Productivity Center flags this condition. The fabric probe or zone control operation fails under this condition with error message AGT0430I.

    Consider using Storage Resource agents deployed on another platform that is connected to the same fabric. Use this Storage Resource agent for collecting zone information and for performing zone changes. In this case, the fabric functions for the Storage Resource agent that is failing can be disabled.

    If the fabric probe fails with the error message, and you try to make zone changes for the fabric, then you might get a warning message that the zoning has changed. Also, the zone control job log displays the AGT0432E error message.

For information about the supported agent types by switch vendor, refer to "Planning for Brocade or McData fabric management agents" on page 99 or "Planning for QLogic or Cisco fabric management agents" on page 103.

Storage Resource agent or in-band Fabric agent

The Storage Resource agent or in-band Fabric agent collects information about the SAN and sends that information to the IBM Tivoli Storage Productivity Center Device server. The Storage Resource agent or in-band Fabric agent can gather topology information for the entire fabric. To gather host-level and detailed Host Bus Adapter (HBA) information, the agent must be installed on each host where that information is wanted.

The Storage Resource agent or in-band Fabric agent performs the following functions:

- Gathers information about the SAN by querying switches and devices for topology and identification information.
- Gathers host-level information for the local system.
- Gathers information about the zoning and allows zone control (except for Brocade switches) of the fabric.
- Gathers information about the HBAs installed on the local system, including make, model, and driver versions.
- Gathers event information detected by the HBAs and forwards it to the Device server.

Note: This function applies only to in-band Fabric agents. Storage Resource agents do not gather event information detected by the HBAs.

Limitations to note are:

- A known limitation exists regarding zone control with previous in-band Fabric agents. For McDATA fabrics, Tivoli Storage Productivity Center users must use the McDATA SMIS agents, or must have all Storage Resource agents or in-band Fabric agents that are attached to the fabric to use the zone control feature of Tivoli Storage Productivity Center. For Cisco and QLogic fabrics, users must have all Storage Resource agents or in-band Fabric agents attached to the fabric to use the zone control feature of Tivoli Storage Productivity Center.
- Can only gather host-level and HBA information about the host containing the Storage Resource agent or in-band Fabric agent .
- Can only provide detailed identification for devices that are in the same zone as the Storage Resource agent or in-band Fabric agent. The Storage Resource agent or in-band Fabric agent identifies the device by worldwide name (WWN) if the Storage Resource agent or in-band Fabric agent is NOT in the same zone as the device rather than being able to uniquely identify the device by device type (for example, host, subsystem, and so on).
- Identification of endpoint devices is reliant on the endpoint device responding to RNID or other queries.
- Can only gather a subset of the switch attributes. There are some switch attributes that can only be collected using out-of-band Fabric agents.
- If the HBA that the in-band Fabric agent uses is not in the active zone set for a fabric, events from the fabric are not received by the in-band Fabric agent.

  Note: Storage Resource agents do not collect events from HBAs.
- Some removable media devices cannot handle command queuing causing long read or write commands to fail. To avoid this problem, see "Excluding devices from Fabric agent scans" on page 338.

  Note: This item applies only to the in-band Fabric agent.
- Tivoli Storage Productivity Center does not collect information about aliases through the Storage Resource agents or in-band Fabric agents.

Out-of-band Fabric agent

IBM Tivoli Storage Productivity Center uses SNMP queries to discover information about the SAN. Management Information Base (MIB) information is collected from the switches and directors by the out-of-band Fabric agent. Switches and directors are added as out-of-band Fabric agents and contacted from the Tivoli Storage Productivity Center Device server by SNMP.

The out-of-band Fabric agent performs the following functions:

- Gathers information about the fabric by querying the switch or director for topology information.
- Gathers virtual SAN information for Cisco switches.

Limitations to note are:

- Topology information is only gathered for the switch added as an out-of-band Fabric agent. The agent cannot gather the topology information for any connected switches unless they are also defined as out-of-band Fabric agents. If

you are performing out-of-band discovery only on a fabric with several switches, you must install an out-of-band Fabric agent for each switch in the fabric to discover the whole fabric.

- Device information is limited; most devices are unknown with a type equal to "Other" and identified by their WWN.
- There must be a TCP/IP connection between the switch and the Tivoli Storage Productivity Center Device server.
- To enable events from the switch to the Device server, the switch must be configured to send SNMP traps to the Device server. The SNMP traps are synonymous to the events that are received by the in-band Fabric agent.

Configuration guidelines for agent placement

This section provides configuration guidelines for various agent types required for fabric functions.

The placement of in-band agents, out-of-band agents, and CIM agents determines the information displayed in the topology. Switches must be configured to send SNMP traps to the server for fabric event notification. Follow these guidelines to determine the best agent placement strategy for these agent types.

- For a topology map, if no CIM agents are present, you must define Storage Resource agents and out-of-band agents to discover all of your topology. If CIM agents are present for McDATA or Brocade fabrics, Storage Resource agents are still required to collect detailed information about endpoints and HBA information. If no CIM agents are present, switch zoning and LUN masking might restrict access to some devices, and some zone control features are unavailable.
- For a complete topology map, including correct device icons, you must define Storage Resource agents and out-of-band agents, or Storage Resource agents and CIM agents (for Brocade and McDATA), on all servers and switches.
- For information about Brocade zoning, you must configure CIM agents for the Brocade fabric.
- For complete device centric and host centric views, you must place Storage Resource agents on all servers you want displayed.
- Cisco requires an out-of-band discovery for VSAN information.

Before implementing in-band agents, out-of-band agents, and CIM agents, formulate a clear plan of your environment and the information you want to collect.

For more information about what the agents collect and for the supported agent types for switch performance management and fabric zone configuration, see “Collecting data with the Fabric Manager” on page 105.

Planning for Fabric Manager configuration

You must configure the main SAN devices properly before using IBM Tivoli Storage Productivity Center.

There are two main components to setting up Tivoli Storage Productivity Center:

- SAN environment setup
- Tivoli Storage Productivity Center setup

The essential elements of the SAN environment needed for Tivoli Storage Productivity Center are:

- Switches and directors
- Host Bus Adapters (HBAs)
- CIM agents for switches

The Storage Resource agent collects the full zone database information from the same switch in the fabric during a fabric probe. This condition assumes that there are no topology changes.

When there are topology changes and the agent state changes, Tivoli Storage Productivity Center might switch to reporting on the full zone database from a different switch in the fabric. When zone control is done through Tivoli Storage Productivity Center, a check is made to see if the zoning on the switch is the same as zoning information in the Tivoli Storage Productivity Center database. If the information is not the same, a warning message is issued before proceeding with zone changes.

This function is similar to selecting a data source for zoning for CIM agents. For CIM agents, the agent managing the principal switch in the fabric is the default data source for zoning. You have the option of specifying an alternative data source for zoning using the Tivoli Storage Productivity Center user interface.

For any fabric, the active configuration information is the same regardless of which switch in the fabric is queried by Tivoli Storage Productivity Center. However, zoning entities in the inactive configuration (also referred to as full zone database) can be different depending on the switch in the fabric that was queried.

Switches have a configuration parameter that controls propagation of inactive zoning entities to all switches in the fabric. If the parameter is set to propagate inactive zoning entities to every switch in the fabric, then any time an inactive zone entity changes, this change is propagated to all switches in the fabric so all switches have the same full zone database.

For Storage Resource agents, the selected data source is automatically determined by Tivoli Storage Productivity Center. If the Storage Resource agent is connected to multiple switches in the fabric, then the agent collects zoning information using only one switch (so that redundant fabric information is not collected from multiple switches). Not all Storage Resource agents connected to a fabric is used to collect zoning information (only one switch is used).

You might see different inactive zone sets displayed because of the following conditions for the Storage Resource agent:

- A different Storage Resource agent is used to collect information
- There was a topology change
- The Storage Resource agent was stopped
- The Storage Resource agent was uninstalled
- A new Storage Resource agent was deployed

You might see different inactive zone sets displayed because of the following conditions for the switch:

- A different CIM agent is used
- The switch is not configured to propagate inactive zoning entities

If the switch configuration parameter is not set to propagate the information, then it is possible for switches in the fabric that have full zone database to have different zoning entities (active configuration is always the same for all switches in the fabric).

For fabric configuration, follow these steps:

1. Install and configure the SMI-S agent (CIM agent) for the switch. Contact your switch vendor for instructions on how to install and configure the SMI-S agent.
2. For every fabric you want to monitor, configure one or more switches in the fabric to send SNMP traps to the Device server. This action causes SNMP traps to be sent to Tivoli Storage Productivity Center whenever there are changes in the fabric. Tivoli Storage Productivity Center can then perform automatic fabric discoveries and probes in response to these traps.
3. For each fabric you want to monitor, install a Storage Resource agent on each host that is connected to the fabric of interest.
4. Run a CIMOM discovery job for the CIM agents. Wait for the job to complete successfully for the CIM agents.
5. Run a fabric probe job. Wait for the job to complete successfully for the fabric.
6. You can now view information for the fabric and create a switch performance monitor job.

These are limitations for switches:

SNMP V1 protocol
Switches that are being managed out-of-band such as SNMP agents must be configured to use the SNMP V1 protocol. Some switches use SNMP V2 or V3 by default and must be reconfigured before they can work correctly with Tivoli Storage Productivity Center. Refer to the switch vendor documentation for specific configuration information.

HBA API for HP-UX
The HBA API for the HP-UX platform does not support the API function that is required for the in-band Fabric agent to support the collection of SAN events through the HBA. If there are no other Fabric agents (out-of-band or in-band agent) that is covering a particular SAN, or no switch in the fabric is configured to send SNMP traps to the Device server, an automatic discovery that would have normally been triggered due to reporting of such events will not occur.

Users might still be able to have a Tivoli Storage Productivity Center agent report SAN events for a particular SAN and have automatic triggering of discovery if one of the following conditions is true:

- Installing a Storage Resource agent on another supported host (non-HP-UX platform) that is connected to the same SAN that the HP-UX Fabric agent or Storage Resource agent is connected to.
- Defining out-of-band agents that are configured to get SNMP traps from switches in that SAN. You must consider this option carefully because there can be performance implications for this option.

Performance statistics
Some switch vendors do not return all the performance statistics required by Tivoli Storage Productivity Center. In these cases, some performance statistics might report as zero.

Note: A CIMOM discovery and fabric probe are required for all switch vendors supporting SMIS 1.1 and higher.

In the Storage Management Initiative Specification (SMI-S), peak data rates for a port are reported in the CIM_FcPortRateStatistics class. For SMI-S Version 1.1, this class is an optional class, meaning that it might not be supported by a given vendor. This can be the case if the Tivoli Storage Productivity Center user interface shows zero values for the peak rates, even when there is non-zero traffic flowing through the port. When a vendor does support this class, the peak data rate represents the peak value of the associated metric counter (as reported by the CIMOM used to access the device) over a recent, vendor-defined measurement window (for example, 20 milliseconds) of the Tivoli Storage Productivity Center measurement window (for example, 15 minutes).

Index number of a fiber channel port on a blade
Currently there is no SMI-S standard property that holds the index value number of a fiber channel port on a blade. For this reason, the Tivoli Storage Productivity Center SMIS Fabric Topology Table Mappers must have vendor-specific code for calculating the index. At this time, only Cisco and Brocade implement the Blades Subprofile, so index calculations are done only for Cisco and Brocade.

Switch blade port index values
The switch blade port index values might be incorrect for Brocade switches where one of the following scenarios apply:

- When a port swap operation is performed
- When the Extended Edge PID mode has been enabled on the director. For more information about the Extended Edge PID mode, see the Brocade document *Fabric OS Administrator's Guide*.

Planning for Fabric Manager switches and directors

For the Fabric Manager to gather and display information about switches and directors, the switches and directors must be configured correctly. The configuration varies between vendors. Incorrect switch configuration can lead to missing information and misconceptions about the Tivoli Storage Productivity Center product not working properly with certain switches.

For the switch to successfully receive and respond to queries, there are some basic requirements:

- The FC Management MIB and FE MIB must be enabled on the switch.
- The switch must be configured to receive SNMPv1 queries and respond in SNMPv1. Some switches are configured to use SNMPv2 or SNMPv3 by default.
- The community string configured in Tivoli Storage Productivity Center must match one of the community strings configured on the switch with read access. Cisco switches must additionally have a community string match for write access. The default community strings in Tivoli Storage Productivity Center are "public" for read access and "private" for write access. Additional community strings can be defined on the switches, but are not used.
- SNMP access control lists need to include the Tivoli Storage Productivity Center Device Server. Some automatically include all hosts while others exclude all hosts by default.

Another aspect of the SNMP configuration includes trap notification. SNMP traps are generated by the switch and directed to the Tivoli Storage Productivity Center Device Server as an indication that something in the fabric has changed and that a discovery must occur to identify changes. The default configuration for handling switch traps is to send them from the switch to port 162 on the Tivoli Storage

Productivity Center Device Server. For the successful generation and reception of traps, there are some configuration requirements:

- The trap destination must be set. The trap destination is typically the Tivoli Storage Productivity Center Device server, but it can also be an intermediary SNMP manager that receives the trap and sends it to Tivoli Storage Productivity Center. The trap destination is set on the switch.
- The destination port must be set on the switch. Tivoli Storage Productivity Center listens on port 162 by default.
- The traps must be sent as SNMPv1. The SNMPv1 designation is set on the switch.
- The trap severity level must be set to generate traps for change conditions. The trap severity level indicates that error level traps and anything more severe are sent. The trap severity level is set on the switch.

Note:

- Storage Resource agents do not collect HBA events.
- If you want to monitor the fabric, you should have one or more switches configured to send SNMP traps to the server.

Configuring these settings differs between switch vendors and models. For information about configuring supported switches, see http://www.ibm.com/developerworks. In the upper right corner, search for "configure switches." Open the article "Configure switches successfully for Tivoli Storage Productivity Center."

Planning for host bus adapters (HBAs)

Proper installation of the HBA driver is an essential element for enabling data collection and is one of the most common problems when an agent appears to be healthy, but is not discovering any data other than the host information. Some HBA vendors require a separate installation of the HBA API package while others include it with the driver.

For the switches to successfully receive and respond to queries, there are some basic requirements:

- The switch must support the FC-GS-3 standard interface for discovery:
  - Name server
  - Configuration server
  - Unzoned name server
- For zone control functions, the fabric zone server must be supported, except in the case of Brocade switches.

There are some basic requirements on the HBA for proper agent operation:

- The HBA driver must be installed.
- The HBA API package must be installed. Some vendors include the HBA API package as part of their driver.
- HBA firmware is at a level that IBM Tivoli Storage Productivity Center supports.
- The current state of the HBA as reported by the operating system is good.
- The HBA has a good connection to the fabric.

Fabric events are automatically sent from the agent to the Tivoli Storage Productivity Center Device server. There is no need for configuration.

Note: This item applies only to the in-band Fabric agents. Storage Resource agents do not send fabric events to the server.

Planning for private switch networks

Some switch vendors recommend a private IP network for the fiber channel switches. It is important to understand that the IBM Tivoli Storage Productivity Center Device server cannot communicate with the switches if they are on a private IP network. It is important to note that the out-of-band Fabric agents require a TCP/IP connection from the the Device server to the switch. Also, SNMP traps from the switches cannot travel directly from the switches to the the Device server.

If you are using a private switch IP network you must rely on Storage Resource agents to gather your SAN information and to forward SAN events to the Device server.

Another option that is sometimes used with a private switch network is to allow the Tivoli Storage Productivity Center Device server to communicate on the private switch network using a second network interface card (NIC).

Planning for performance management

IBM Tivoli Storage Productivity Center can collect performance data for devices that use the native interfaces (DS8000, the XIV system, SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 storage systems), or for storage systems and fiber channel switches that use CIM agents. The devices that use CIM agents must be SMI-S 1.1 compliant.

There are three main functions for Tivoli Storage Productivity Center performance management:

- Performance data collection (performance monitoring)
- Performance thresholds and alerts
- Performance reports

You can define how long performance data is stored in the database. Tivoli Storage Productivity Center can set thresholds for important performance metrics. When any boundary condition is crossed, Tivoli Storage Productivity Center can notify the user through email, SNMP, or other alerting mechanisms. Tivoli Storage Productivity Center can generate reports, historic trend charts, and help analyze the bottleneck of a performance congestion by drilling down to threshold violated components and the affected hosts. A combination of these functions can be used to:

- Monitor a complicated storage network environment
- Predict warning signs of system failure
- Do capacity planning as overall workload grows

For more information about performance planning, see http://www-01.ibm.com/support/docview.wss?uid=swg27017573.

Performance metrics

IBM Tivoli Storage Productivity Center can report on many different performance metrics, which indicate the particular performance characteristics of the monitored devices.

Two important metrics for storage systems are throughput in I/O per second, and the response time in milliseconds. Throughput is measured and reported in several different ways:

- Throughput of an entire box (subsystem)
- Each cluster (TotalStorage Enterprise Storage Server)
- Each controller (DS6000 or DS8000)
- Each I/O Group (SAN Volume Controller or Storwize V7000

Throughputs are measured for:

- Each volume (or LUN)
- At the Fibre Channel interfaces (ports) on some of the storage boxes
- On fiber channel switches
- At the RAID array after cache hits have been filtered out

For storage systems, the performance statistics are separated into frontend I/O metrics and backend I/O metrics. Frontend I/O metrics are a measure of the traffic between the servers and storage systems. Backend I/O metrics are a measure of all traffic between the storage system cache and the disks in the RAID arrays in the backend of the subsystem. Most storage systems give metrics for both kinds of I/O operations: frontend and backend operations. It is important to know whether the throughput and response times are at the frontend (close to the system level response time as measured from a server) or backend (between the cache and disk).

Note: A CIMOM discovery and fabric probe are required for all switch vendors supporting SMIS 1.1 and higher.

The main frontend throughput metrics are:

- Total IO rate (overall)
- Read IO rate (overall)
- Write IO rate (overall)

The corresponding frontend response time metrics are:

- Overall response time
- Read response time
- Write response time

The main backend throughput metrics are:

- Total backend IO rate (overall)
- Backend read IO rate (overall)
- Backend write IO rate (overall)

The corresponding backend response time metrics are:

- Overall backend response time
- Backend read response time
- Backend write response time

It is desirable to track any growth or change in the rates and response times. It frequently happens that I/O rate grows over time, and that response time increases as the I/O rates increase. This relationship is what "capacity planning" is all about.

As I/O rates increase, and as response times increase, you can use these trends to project when additional storage performance (as well as capacity) is required.

Depending on the particular storage environment, it might be that throughput or response time times change drastically from hour to hour or day to day. There might be periods when the values fall outside the expected range of values. In that case, other performance metrics can be used to understand what is happening. Here are some additional metrics that can be used to make sense of throughput and response times:

- Total cache hit percentage
- Read cache hit percentage
- Write-cache delay percentage (previously known as NVS full percentage)
- Read transfer size (KB/operation)
- Write transfer size (KB/operation)

Low cache hit percentages can drive up response times, because a cache miss requires access to backend storage. Low hit percentages also tend to increase the utilization percentage of the backend storage, which might adversely affect the backend throughput and response times. High write-cache delay percentage (previously known as NVS full percentage) can drive up the write response times. High transfer sizes typically indicate more of a batch workload, in which case the overall data rates are more important than the I/O rates and the response times.

All these metrics can be monitored through reports or graphs in the Tivoli Storage Productivity Center GUI. Some examples of supported thresholds are:

- Total I/O rate and total data rate thresholds
- Total backend I/O rate and total backend data rate thresholds
- Read backend response time and write backend response time thresholds
- Total port I/O rate (packet rate) and data rate thresholds
- Overall port response time threshold
- Port send utilization percentage and port receive utilization percentage thresholds
- Port send bandwidth percentage and port receive bandwidth percentage thresholds

For fiber channel switches, the important metrics are total port packet rate and total port data rate, which provide the traffic pattern over a particular switch port. Port bandwidth percentage metrics are also important to provide an indicator of bandwidth usage based on port speeds. When there are lost frames from the host to the switch port, or from the switch port to a storage device, the dumped frame rate on the port can be monitored.

The important things are:

- Monitor the throughput and response time patterns over time for your environment
- Develop an understanding of expected behaviors
- Investigate the deviations from normal patterns of behavior to get warning signs of abnormal behavior
- Generate the trend of workload changes

Performance management configuration

Performance data can be collected from devices that are managed through the native interfaces or with devices that are managed by a CIM agent. For those devices that use the CIM agent, the devices must be SMI-S compliant.

To get performance data for devices, follow these general steps:

1. Prepare for performance monitoring.

   For those devices that need the CIM agent, make sure that the following items are done before adding the CIM agent:

   - The version of CIM agent and firmware for the device is supported. For information about the CIM agent and firmware supported, see http://www-01.ibm.com/support/docview.wss?rs=40&q1=support+matrix&uid=swg21386446&loc=en_US&cs=utf-8&lang=en. Select the appropriate version for **Storage**.
   - A CIM agent is installed on a different server from IBM Tivoli Storage Productivity Center. For some switch vendors, the switch comes with an embedded CIM agent, and therefore, no separate CIM agent installation is required.
   - For subsystems or switches on a private network, be sure to have the CIM agent installed on a gateway machine so that the Tivoli Storage Productivity Center server on a different machine can communicate with the CIM agent.
   - The CIM agent is configured to manage the intended device.

   For information about how to install and configure the CIM agent, contact your CIM agent provider. Also see *IBM Tivoli Storage Productivity Center Hints and Tips* for information about installing and configuring CIM agents. See http://www-01.ibm.com/support/docview.wss?uid=swg27008254.

2. Select the steps to follow for your device:

   - For the devices that use the native interface DS8000, the XIV system, SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000, follow these steps:
     a. Add the device to Tivoli Storage Productivity Center. Use the Configure Devices wizard to add the device. For information about the Configure Devices wizard, see "Configure Devices" on page 266. On the Select Device Type panel, select **Storage Subsystem**.
     b. Once the device is added to Tivoli Storage Productivity Center, automatic discovery and probe operations are done for the storage systems.
   - For devices that use the CIM agent (ESS, DS3000, DS4000, DS5000, DS6000, all non-IBM devices), follow these steps:
     a. Add the CIM agent to Tivoli Storage Productivity Center. For information about adding a CIM agent, use the Configure Devices wizard. For information about the Configure Devices wizard, see "Configure Devices" on page 266.
     b. Run a CIMOM discovery job for the storage system or switch.
     c. Run a probe job for the storage system or switch. The probe job gets all component level information into the database. The probe job also includes getting all the instances of controllers, arrays, extents, volumes, LUN to host mappings, FC ports, and so on. Only those storage systems for which configuration information has been collected successfully can have performance monitor jobs run on them.
     d. For fiber channel switches, run a fabric probe job for the fabric.

3. Create threshold alerts. A performance threshold is a mechanism by which you can specify one or more boundary values for various performance metrics, and can specify to be notified if any collected performance data violates these boundaries.

   Tivoli Storage Productivity Center has default thresholds defined (but no default alerts). If you explicitly define an alert for one of the default thresholds, that means that if the threshold is violated, the constraint violation report displays the violation but the alert log will not.

   When thresholds are applied during the processing of data collection, a performance monitor for a device must be actively running for a threshold violation to be recognized. Chronologically, you would need to define the threshold alert before the performance monitor job starts.

   There is a special report that can provide statistics about how often and when an alert threshold has been reached, even without setting up alert destinations like email receivers or SNMP trap destinations. This special report is called a constraint violation report. To get a constraint violation report, expand **Disk Manager** > **Reporting** > **Storage Subsystem Performance** > **Constraint Violations**.

   Note: Alerts can be created, deleted, or modified at any time, even if the performance monitor is already running. Changes are applied dynamically to the running job and take effect the next time data is collected for the job.
4. Create a performance monitor job. Performance data collection is performed through a performance monitor job. This job can be run immediately, can be scheduled to run one time only, or can be scheduled to run repeatedly as wanted. Only after the device has been probed successfully is that device eligible for a performance monitor. If the probe is not successful, that device will not show up in the performance monitor creation panel.

   Performance management for storage systems is done through the Disk Manager. Data collection for the storage systems can be scheduled through **Disk Manager** > **Monitoring** > **Subsystem Performance Monitors**. Right-click **Subsystem Performance Monitors**. Click **Create Subsystem Performance Monitor**. In the content pane, enter the information schedule and run a performance monitor job. To view the performance reports, expand **Disk Manager** > **Reporting** > **Storage Subsystem Performance**.

   Performance management for fiber channel switches is done through the Fabric Manager. Data collection for the switches can be scheduled through **Fabric Manager** > **Monitoring** > **Switch Performance Monitors**. Right-click **Switch Performance Monitors** and click **Create Switch Performance Monitors**. To view the performance reports, expand **Fabric Manager** > **Monitoring** > **Switch Performance Monitors**.
5. Check the performance monitor status. When the performance monitor job starts to run, you begin to collect performance data for the device. Check the status of the performance monitor job to make sure that it runs and continues running. In the navigation tree, expand **IBM Tivoli Storage Productivity Center** > **Job Management**. In the content pane, click the **Refresh** button because refreshes are not automatic on the Job Management panel. See Jobs for Selected Schedule for a list of performance monitor jobs and the status of the jobs. For more information about the Job Management panel, see "Managing schedules, jobs, and runs" on page 283.
6. Specify the retention period for the performance data. After the monitor job is created, you can configure how long to keep performance data. You can specify a retention period for collected sample data, for aggregated hourly data, and

for daily data. The retention period applies to all performance monitors. For more information about the retention period, see “Resource History Retention” on page 332.

Sample data is the data that is collected at the specified interval length of the monitor job, for example, data is collected every 5 minutes.

For storage systems, the most numerous components are usually volumes, therefore, the largest amount of performance data collected is for volumes.

7. View the performance reports in Tivoli Storage Productivity Center.

Tivoli Provisioning Manager 5.1 for storage workflows

IBM Tivoli Storage Productivity Center supports Tivoli Provisioning Manager 5.1 for storage workflows. For information about how to use storage workflows, see *IBM Tivoli Storage Productivity Center Workflow User's Guide*.

Microsoft Cluster Server

Data Manager can monitor and report on Microsoft Cluster Server (MSCS) clustered nodes and cluster resource groups.

Microsoft Cluster Server (MSCS) is a built-in feature of Windows. It is software that supports the connection of up to eight servers into a "cluster" for higher availability and easier manageability of data and applications. MSCS can automatically detect and recover from server or application failures. It can be used to move server workload to balance utilization and to provide for planned maintenance without downtime.

As with standard Windows clusters, you perform most of the configuration tasks, and also the management tasks, associated with Exchange clusters using the Cluster Administrator. Cluster Administrator is installed by default on servers that have Cluster Service installed.

You can also use Cluster Administrator to remotely administer a server cluster. Computers that are used to administer a server cluster remotely must be secure and restricted to trusted personnel.

When a cluster node is probed, there are normal errors in the logs for the disks belonging to the cluster resource groups that are not currently hosted on this node.

When looking at an MSCS cluster problem, the current state of the clustering environment is often needed. Microsoft provides a tool to dump the MSCS cluster configuration into a set of files.

Note: The term "virtual servers" has been replaced with "cluster resource groups" in IBM Tivoli Storage Productivity Center. Microsoft still refers to "virtual servers."

Microsoft Cluster Server environment

This topic provides information about IBM Tivoli Storage Productivity Center support in an MSCS environment.

Report information can include local resources for a clustered node as well as the resources in a cluster resource group. The Storage Resource agent is not cluster aware so the Storage Resource agent cannot run in the cluster resource group and fail over from one node to another. The Storage Resource agent supports failover of a clustered file system so that Tivoli Storage Productivity Center scan jobs can

continue to work when the clustered file system has moved from one node to another. The Storage Resource agent running on the node hosting the cluster resource group discovers the resources on that cluster resource group. The Storage Resource agents need to be installed on every node in the cluster.

The following information lists the requirements for using the Data server in an MSCS cluster environment.

Server requirement
Not applicable. The server component can be installed on a node in an MSCS cluster, but the Data Manager can only monitor and report on nodes that have an agent installed.

agent requirements
Each node in an MSCS cluster must have a Storage Resource agent installed. The Data Manager can only monitor and report on nodes that have a daemon Storage Resource agent installed. All agents installed in a cluster must use the same port number to communicate with the Tivoli Storage Productivity Center server. Agents in other clusters can be configured with a different port address. However, all agents in the cluster must use the port address used by the other agents in that cluster.

You can use either Storage Resource agents or Data agents on each node. You can have a mixed configuration when migrating Data agents to Storage Resource agents, but scan jobs of clustered file systems will fail until all the nodes in the cluster have the same type of agent.

client requirements
Not applicable. The client component can be installed on a node in an MSCS cluster, but the Data Manager can monitor and report on only the nodes with an agent installed.

Note: The Data server does not currently support the monitoring of clustered database application.

Note:

- Cluster resource group scanning is only supported by the Tivoli Storage Productivity Center for Data and Tivoli Storage Productivity Center Standard Edition licenses.
- If the Storage Resource agent port is behind a firewall, that port needs to be opened for all virtual server and cluster resource group addresses.

Microsoft Cluster Server support

Use this information to plan for running IBM Tivoli Storage Productivity Center with MSCS.

Cluster resource groups

A *cluster resource group* represents an MSCS entity.

You can configure MSCS into cluster resource groups. You can define resource group policies and attributes that dictate how MSCS manages resources to keep them highly available at different stages of cluster operation (startup, failover, and fallback). The types of resources you can put into clustered resource groups are:

Logical volumes
A set of logical partitions that MSCS makes available as a single storage unit that is, the logical volume is the “logical view” of a physical disk.

Filesystem
A file system is written to a single logical volume. Typically, you organize a set of files as a file system for convenience and speed in managing data.

Shared filesystems
A file system that resides entirely in a shared logical volume.

Physical disk
A physical disk.

Network name
A network name.

IP address
An IP address.

Reports

The asset reporting navigation tree includes a By Cluster report that shows the nodes and cluster resource groups in each cluster. The disks and file systems that are shared as clustered resources are associated with the cluster resource group and not with the node where they are hosted.

When generating reports, you can use the By Cluster report subtype wherever the By Computer subtype is available. The By Cluster report summarizes the results of the nodes and the cluster resource groups in the cluster.

Monitoring groups

When you create a monitoring group or create a data collection schedule, you can select individual nodes and cluster resource groups. You can create a file system monitoring group that includes clustered and non-clustered resources.

Computer alerts

You can register for cluster-specific alerts that notifies you that your cluster has gone through a configuration change such as a cluster resource being added or removed. You can also register for an alert that is triggered when a cluster resource group is moved. The same alert is used for a cluster resource group failover, cluster resource group fallback, and when the cluster resource group is manually moved to a new node.

Cluster alerts are not triggered until the next time the cluster is probed.

Upgrading agents in an MSCS cluster

Take the following steps when upgrading the Tivoli Storage Productivity Center agents in an MSCS cluster:

- Make sure that all agents in the cluster are configured to use the same port address.
- Use either the Tivoli Storage Productivity Center installation program or the graphical user interface to perform the agent upgrades.
- Do not probe or scan the agents in the cluster while the agent upgrade is in progress.

- Once all agents in the cluster have been upgraded successfully, create and run a new probe job on the cluster.

For more information about upgrading Tivoli Storage Productivity Center agents, see "Upgrading Storage Resource agents" on page 406.

Installing MSCS support

Complete the following general steps to install MSCS support:

1. Make sure that your MSCS environment is configured for Tivoli Storage Productivity Center support.
2. Install a Storage Resource agent on each node of the cluster. All agents in the cluster must be configured to use the same listening port.
3. Run a discovery job for the agent.
4. Add a Network Name resource with an IP address that can be contacted from the Tivoli Storage Productivity Center server.
5. Run a scan job for the agent.
6. Run a probe job for the agent.
7. View reports for the MSCS entities.
8. View the MSCS environment in the topology viewer.

Planning for PowerHA

Tivoli Storage Productivity Center supports Storage Resource agents installed on PowerHA nodes (previously known as HACMP). Use this information to configure the PowerHA environment before using it with Tivoli Storage Productivity Center.

Cluster resource groups

A *cluster resource group* represents a PowerHA entity.

You can configure PowerHA into cluster resource groups so that they can be highly available. You can define resource group policies and attributes that dictate how PowerHA manages resources to keep them highly available at different stages of cluster operation (startup, failover, and fallback). The types of resources you can put into clustered resource groups are:

Volume groups
A set of physical volumes that AIX treats as a contiguous, addressable disk region.

Logical volumes
A set of logical partitions that AIX makes available as a single storage unit that is, the logical volume is the "logical view" of a physical disk.

Filesystem
A file system is written to a single logical volume. Typically, you organize a set of files as a file system for convenience and speed in managing data.

Shared filesystems
A journaled file system that resides entirely in a shared logical volume.

Applications
Critical services that are accessible to users.

Service IP labels or addresses
A way to establish communication between client nodes and the server

node. Services, such as a database application, are provided using the connection mode over the service IP label.

Tape resources
You can configure a SCSI or a Fiber Channel tape drive as a cluster resource in a non-concurrent resource group, making it highly available to two nodes in a cluster.

Communication links
You can define communication links as resources in a PowerHA resource group.

Planning for PowerHA support

With PowerHA, you can use the following types of configurations:

- PowerHA with cluster resources groups that are nonconcurrent. Concurrent cluster resource groups are not supported.
- The following volume groups in a PowerHA environment:
  - Standard volume groups
  - Enhanced concurrent-mode volume groups
  - Scalable volume groups

When you create a Tivoli Storage Productivity Center configuration, include the Storage Resource agent. A Storage Resource agent must be installed on each node of the cluster. All agents in a cluster must be configured to use the same listening port. You must also have a cluster resource group (CRG) with at least one IP address that is accessible from the Tivoli Storage Productivity Center server if you plan to perform scans on the cluster resource group. Each Storage Resource agent collects information about the local node and the clustered resources that are currently hosted on the node.

The following operations are not supported:

- Probing and scanning of clustered databases
- Reporting on cluster information when the cluster is configured to have concurrent cluster resource groups

Note:

- Cluster resource group scanning is only supported by the Tivoli Storage Productivity Center for Data and Tivoli Storage Productivity Center Standard Edition licenses.
- If the Storage Resource agent port is behind a firewall, that port needs to be opened for all virtual server and cluster resource group addresses.

Reports

The asset reporting navigation tree includes a By Cluster report that shows the nodes and cluster resource groups in each cluster. The disks, volume groups, and file systems that are shared as clustered resources are associated with the cluster resource group and not with the node where they are hosted.

When generating reports, you can use the By Cluster report subtype wherever the By Computer subtype is available. The By Cluster report summarizes the results of the nodes and the cluster resource groups in the cluster.

Monitoring groups

When you create a monitoring group or create a data collection schedule, you can select individual nodes and cluster resource groups. You can create a file system monitoring group that includes clustered and non-clustered resources.

Computer alerts

You can register for cluster-specific alerts that notify you that your cluster has gone through a configuration change such as a cluster resource being added or removed. You can also register for an alert that is triggered when a cluster resource group is moved. The same alert is used for a cluster resource group failover, cluster resource group fallback, and when the cluster resource group is manually moved to a new node.

Cluster alerts will not be triggered until the next time the cluster is probed.

Installing PowerHA support

Complete the following steps to install PowerHA support:

1. Ensure that your PowerHA environment is configured for Tivoli Storage Productivity Center support.
2. Create a service IP label for the CRG with an IP address that can be contacted from the Tivoli Storage Productivity Center server.
3. Install a Storage Resource agent on each node of the cluster. All agents in the cluster must be configured to use the same listening port.
4. Run a discovery job for the agent.
5. Run a probe job on the agents.
6. Run a scan job on the agents.
7. Run a scan job on the cluster resource groups.
8. View reports for the PowerHA entities.
9. View the PowerHA environment in the topology viewer.

PowerHA environment

This topic provides information about Tivoli Storage Productivity Center support in a PowerHA environment.

Tivoli Storage Productivity Center supports Storage Resource agents installed on PowerHA nodes. The Storage Resource agent must be installed on every node in the cluster. You cannot configure the Storage Resource agent as a clustered application.

The Storage Resource agent collects information from the cluster when the node is probed. The information that is collected is available in the Data Manager and Disk Manager reports and the topology viewer. Resources that are not clustered are reported under the node. The resources that are associated with a cluster resource group are reported under a computer entity that represents the cluster resource group. For example, if a single node cluster that has one cluster resource group is probed it produces two computer entities in the Data Manager and Disk Manager reports and the topology viewer.

The following entities are reported under the associated clustered resource group and not the node:

- Volume groups
- NFS shares
- Service IP labels

Physical volumes, logical volumes, and file systems for clustered volume groups are also reported under the cluster resource group. Scan and ping jobs can be created for a cluster resource group and these jobs work regardless of which node is hosting the cluster resource group. Scan and ping requests for a cluster resource group are sent to the service IP label defined for the cluster resource group. Therefore, in order for these jobs to succeed the service IP label must be accessible from the Data server and the CRG must be online.

Restriction: You can use PowerHA only with cluster resource groups that are nonconcurrent. You cannot use PowerHA with concurrent cluster resource groups.

The following information lists the support requirements for a PowerHA cluster environment.

Agent requirement
Each node in a PowerHA cluster must have a Storage Resource agent installed. The Data Manager can only monitor and report on nodes that have a daemon Storage Resource agent installed. If you are using daemon Storage Resource agents, all agents installed in a cluster must listen on the same port number. Agents in other clusters can be configured with a different port address. However, all agents in the cluster must use the port address used by the other agents in that cluster.

You can use either Storage Resource agents or Data agents on each node. You can have a mixed configuration when migrating Data agents to Storage Resource agents, but scan jobs of clustered file systems will fail until all the nodes in the cluster have the same type of agent.

Probe requirements
Probes are not automatically executed in response to cluster events. Schedule probes to run as appropriate for the needs of the environment.

Scan requirements
To perform scans on a cluster resource group, the cluster resource group must be configured with an IP address that the Tivoli Storage Productivity Center server can communicate with. This IP address is displayed in reports associated with PowerHA clusters. If the server cannot contact this address once it has been configured, the server tries the next known address for the cluster resource group until it successfully contacts an IP address.

Note: Cluster resource groups are now listed as computers in scan job results, and now have log files specific to the cluster resource group.

NAS and NetWare support

This section provides information for the planning of NAS and NetWare support.

Network Attached Storage system requirements

This section provides information about the requirements for installing and running the Data Manager within a Network Attached Storage (NAS) environment.

IBM Tivoli Storage Productivity Center supports the monitoring of the NetApp Data ONTAP.

To be supported, NAS devices other than Network Appliance Filers must meet the following criteria:

- If scanned from a UNIX agent, a NAS device must support Network File System (NFS) queries.
- If scanned from a Windows agent, a NAS device must support Common Internet File System (CIFS) queries.
- A NAS device must support SNMP queries (sysName and sysOID).
- A NAS device must supply a unique sysName.
- If a NAS device hosts both Windows and UNIX file systems, you need at least two proxy Storage Resource agents to monitor the device. At least one Windows proxy Storage Resource agent is required to probe and scan the Windows CIFS shares. At least one UNIX proxy Storage Resource agent is required to probe and scan UNIX NFS shares.

For detailed information about configuring NAS and Tivoli Storage Productivity Center, see http://www.redbooks.ibm.com/. Search for **sg247490**.

Here are the requirements for installing and running the Data Manager within a NAS environment.

Server requirement
Not applicable. You cannot install the server on a NAS filer.

Agent requirements
You must install the agent on a machine that has access to the NAS filers within your environment that you want to monitor.

Windows
The agent that is logging in to and scanning the NAS filer is not required to be in the same domain as the user or the NAS filer. If you install the agent on a different domain from the NAS filer, the agent scans the NAS filer *if* the domain of the agent computer is a "trusted domain" by the domain of the NAS filer.

UNIX and Linux
The agent computer must import the NAS filer export files as NFS mounts (or automounts on Oracle Solaris).

Note:

- You do not install agents to the NAS filers themselves.
- For Tivoli Storage Productivity Center 5196 Network Attached Storage 300 G machines only, install the agent directly on those machines. If the 300 G is clustered, you need to install an agent on each local node.

Client requirements
Not applicable. You cannot install the client component to a NAS filer.

NAS requirements
The NAS filers within your environment must be visible to the machines where you install the agent or agents. If you want to monitor NAS filers from Windows, you must configure those NAS filers to be members of a Windows domain.

Planning for NetApp device support

NetApp devices provide unified storage for organizations with Network Attached Storage or storage area network environments that have file and block-level services. IBM Tivoli Storage Productivity Center provides the same file and block-level support for NetApp devices as is provided for other filers and storage subsystem devices.

Overview

Tivoli Storage Productivity Center provides the following support for NetApp devices:

- supports NetApp Data ONTAP SMI-S Agent 3.0
- supports the SMI-S 1.2 Array profile implemented by the NetApp SMI-S agent
- supports all filer models running NetApp Data ONTAP version 7.2 and 7.3

NAS Gateway and V Series are not supported with the Array profile.

The following licenses are required:

- For performance monitoring, a IBM Tivoli Storage Productivity Center for Disk license is required
- For SNMP discovery, a IBM Tivoli Storage Productivity Center Standard Edition license is required

Configuration considerations

The way you use Tivoli Storage Productivity Center to interact with NetApp devices, and the data you can collect for NetApp devices, depends on how you configure your NetApp devices. For example:

- If you configure a Storage Resource agent to be a Scan/Proxy agent, you use Data Manager to work with NetApp devices. The Data Manager user interface displays a NetApp filer as a "computer". For NetApp devices configured this way, Tivoli Storage Productivity Center collects file storage information.
- If you configure a NetApp SMI-S agent (CIMOM), you use Disk Manager to work with NetApp devices. The Disk Manager user interface displays a NetApp device as a "subsystem". For NetApp devices configured this way, Tivoli Storage Productivity Center collects block storage information.
- If you configure a NetApp device as both a filer and a subsystem, Tivoli Storage Productivity Center collects both file and block storage information. You can use Disk Manager and Data Manager to work with the NetApp device. Keep in mind that the Data Manager user interface displays a NetApp filer as a "computer," and the Disk Manager user interface displays a NetApp device as a "subsystem".
- If multiple Storage Resource agents are set up to probe or scan the same NetApp filer, 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. For example, if you add a Storage Resource Agent for the Windows operating system and then you add a Storage Resource Agent for the UNIX operating system, Windows operating system related data such as shares are shown, but not exports.

If you initially configure a NetApp device using a Storage Resource agent, you do not lose any function if you later decide to also configure the NetApp device using a CIM agent.

General procedure to install and configure NetApp devices

The general steps to install and configure a NetApp device to enable Data Manager function are:

1. Select a machine to act as the Storage Resource agent that meets the auto-discovery requirements for Tivoli Storage Productivity Center to auto-discover the NetApp filers.
2. Install a Storage Resource agent on that machine.
3. Set the license for the NetApp filers that are discovered.
4. Set the Scan/Proxy agent for the discovered filers.
5. Create and run a probe job for each of the filers to collect Data Manager information.

The general steps to install and configure a NetApp device to enable Disk Manager functionality are:

1. Install Tivoli Storage Productivity Center or upgrade to version 4.2 or later.
2. Install the NetApp SMI-S agent and add filers to its configuration using the NetApp SMI-S agent utility
3. Add the NetApp CIMOM by using **Administrative Services** > **Data Sources** > **CIMOM Agents**. In the content pane, click **Add CIMOM**. To automatically discover the NetApp SMI-S agent using SLP, skip this step.
4. Run a CIMOM discovery.
5. If you want to use SLP to discover the NetApp SMI-S agent, make sure that your SLP directory agents are configured or the **Scan local subnet** box is checked on the **Administrative Services** > **Discovery** > **CIMOM**. In the content pane, go to the Options page to enter the IP addresses or host names for the SLP directory agents to be used before you run the CIMOM discovery.
6. Create and run a probe job for the configured filers to collect Disk Manager information.
7. If wanted, use the command-line interface (CLI) **tpctool** command to view information for NetApp storage systems.

Refer to the *IBM Tivoli Storage Productivity Center User's Guide* for more information about how to collect and view information gathered by NetApp devices. Refer to the *IBM Tivoli Storage Productivity Center Command-Line Interface Reference* for information about using the **tpctool** command.

Removing a NetApp device

If you later decide to remove a NetApp device, the steps to remove it depend on how you configured it. For example, if the NetApp device is configured only as a filer, you can remove it using the **Administrative Services** > **Configuration** > **Scan/Probe Agent Administration** page. If the NetApp device is configured only as a subsystem, you can remove it using the **Disk Manager** > **Storage Subsystems** page. If the NetApp device is configured as both a filer and a subsystem, then you must perform both of the preceding steps.

Limitations of NetApp support

Keep in mind the following limitations for NetApp support:

- The NetApp Data ONTAP SMI-S Agent 3.0 or later implements the Block Server Performance subprofile. It provides volume performance data but, for now, does not provide performance data at the storage subsystem level.

- Tivoli Storage Productivity Center supports only the SMI-S Array profile of the NetApp Data ONTAP SMI-S Agent. Other SMI-S profiles, including the self-contained NAS profile and the NAS Head profile, are not supported.
- For the Data ONTAP SMI-S 3.0 agent or later, if a volume is offline, the performance monitor might fail with the message: `PM HWNPM2132W Performance data could not be collected for device device`.
- To work with Network Appliance quotas (using **Data Manager** > **Policy Management** > **Network Appliance Quotas**), the NetApp device must be configured as a filer. This function is not available if the NetApp device is configured only as a CIMOM.

Identifying NetApp devices in the user interface

The way you configure a NetApp device determines how the NetApp device is represented in the Tivoli Storage Productivity Center user interface, as follows:

- NetApp device configured as a filer: the NetApp device appears as a "computer" in the user interface.
- NetApp device configured as a storage subsystem using a CIMOM: the NetApp device appears as a "subsystem" in the user interface.
- NetApp device configured as both a filer and storage subsystem: the NetApp device appears as either a "computer" or a "subsystem" in different places in the user interface. For example, some reports under My Reports refer to the NetApp filer as a "computer" while others refer to the NetApp device as a "subsystem".

The NetApp device is represented either as a "computer" or a "subsystem", or both, depending on how you have configured the NetApp device, are:

- Dashboard
- My Reports
- Topology viewer
- Alerting
- Data Manager
- Disk Manager

For example:

- **Data Manager** > **Reporting** > **Asset** > **By Storage Subsystem** lists a NetApp device as a storage subsystem if the device was probed by the NetApp Data ONTAP SMI-S agent.
- **System Reports** > **Data** > **Disk Defects** displays a NetApp device name as either a "computer" or a "subsystem", depending on whether the device information was collected by the Storage Resource agent or the SMI-S agent.

In the topology viewer, the NetApp device appears as follows:

- If a Storage Resource agent is configured to probe the NetApp device, the NetApp device is displayed as a "computer"
- If a CIM agent is configured to probe the NetApp device, the NetApp device is displayed as a "storage subsystem"

Planning for System Storage N Series Gateway servers

Tivoli Storage Productivity Center supports the System Storage N Series Gateway servers as the **Other NAS** node.

Upgrade information

If you are upgrading Tivoli Storage Productivity Center from TotalStorage Productivity Center 3.3.x, there are specific steps to follow. You must delete and unlicense existing N Series Gateway servers first before adding the N Series Gateway servers. For information about upgrading Tivoli Storage Productivity Center with N Series Gateway servers, see “Upgrading System Storage N Series Gateway servers” on page 410.

New installation

If you are installing Tivoli Storage Productivity Center as a new installation, the N Series Gateway servers can be automatically discovered after the Storage Resource agent is installed. If the N Series Gateway servers are automatically discovered, they are listed under **Administrative Services** > **Configuration** > **License Keys**. In the content pane, click the magnifying glass icon to the left of the product to display the **Licensing** page. The **Licensing** page displays the OS Type as **NetApp Data ONTAP** with the Licensed box cleared.

To manually add the N Series Gateway servers as **Other NAS**, the Licensed box must be cleared. The OS Type and Licensed fields are updated after the N Series Gateway Server is manually added as **Other NAS**. For more information, see “Manually adding an N Series Gateway server ” on page 412.

Planning for VMware

Tivoli Storage Productivity Center supports the VMware Virtual Infrastructure which consists of the ESX Server and VMware VirtualCenter.

Overview

The ESX Server is a true hypervisor product which can host multiple virtual machines that run independently of each other while sharing hardware resources. The VirtualCenter is the management application that is the central entry point for the management and monitoring of the ESX Servers for a data center. Install a Storage Resource agent on each virtual machine that you want to monitor.

For more information about the ESX Server or VMware VirtualCenter, see http://www.vmware.com.

For information about permissions required by Tivoli Storage Productivity Center to access information from VMware, see the information center. Search for "VMware data source."

The storage systems and platforms supported on VMware are shown in Table 26 on page 133.

Table 26. Supported platforms and storage systems for VMware

Supported platform	Supported storage systems	Supported function
ESX Server 3.0.1 ESX Server 3.0.2 ESX Server 3.5 ESX Server 3.5i VirtualCenter 2.0.1 VirtualCenter 2.5	• 3PAR • EMC CLARiiON • EMC Symmetrix • Hitachi Data Systems 9000 series • Hewlett Packard Enterprise Virtual Arrays (EVA) • IBM DS4000 • IBM DS6000 • IBM DS8000 • IBM ESS • IBM SAN Volume Controller • IBM Storwize V7000 • IBM Storwize V7000 Unified • IBM XIV Storage System	Monitoring LUN correlation

For full function, both the Storage Resource agent and Virtual Infrastructure must be up and running. If one of the items is not present in a given environment, only a limited picture is presented to the user. Some virtual machines might not be recognized.

The hierarchical mapping of storage allocated to the virtual machine is available for the virtual machines on the ESX Server.

Note: Tivoli Storage Productivity Center now supports the mapping of storage from the ESX Server to the disk drives for the ESX Server 3.5.

Note the following limitations:

- No HBA virtualization is available for the VMware virtual machines.
- No events directly generated by the Virtual Infrastructure is supported.
- No VMware clusters are supported.
- Data path explorer is not supported for VMware ESX Server and virtual machines.

Tivoli Storage Productivity Center for Replication in a VMware environment

Use Tivoli Storage Productivity Center for Replication to maintain a consistent copy of customer data. You can install Tivoli Storage Productivity Center for Replication under a VMware image. However, there are some aspects of running in a virtualized environment that can interfere with the ability of the software to function correctly if a disaster occurs. Be aware of the following limitations:

- Virtual machines in a VMware environment can be moved from one physical server to another using VMotion or other tools. If you move your virtual machine with Tivoli Storage Productivity Center for Replication on it, unpredictable results can occur during the processing of certain critical

operations. This action might prevent Tivoli Storage Productivity Center for Replication from maintaining data consistency or prevent you from being able to quickly recover your data at the remote site.

- If the VMware image resides on the same storage that is managed by Tivoli Storage Productivity Center for Replication, Tivoli Storage Productivity Center for Replication might attempt to "freeze" I/O on the storage system. This freezing of I/O on the storage system can cause Tivoli Storage Productivity Center for Replication to stop responding. If this freeze occurs, the I/O is not released until the storage system freeze timeout occurs. Any applications running against that storage system might be affected by a delay, which can prevent Tivoli Storage Productivity Center for Replication from maintaining data consistency.
- Virtual images are controlled and managed by the ESX server. You must ensure that the internal VM network is set up correctly so users can access Tivoli Storage Productivity Center for Replication. Correctly setting up the network ensures that:
  - Tivoli Storage Productivity Center for Replication can access the storage systems that it is managing
  - Helps users quickly recover data on the remote storage systems if a disaster occurs

  Also ensure that the virtual environment is secure and that it is monitored to maintain data consistency.
- Because virtual environments can vary, ensure that you thoroughly test your environment with Tivoli Storage Productivity Center for Replication, and ensure that you test disaster recovery-related steps.

Software requirements

For information about the software requirements for your ESX Server or VMware VirtualCenter environment, see “Software requirements for operating systems” on page 5.

Installing Tivoli Storage Productivity Center server on a virtual machine

You can install the Tivoli Storage Productivity Center server on a virtual machine on VMware ESX server 3.5.x. The hardware and operating system requirements are the same requirements as for a physical machine.

In addition, the following requirements must be met:

CPU For the ESX server, do not have more virtual processors than there are physical cores in the system. Plan your system so that no processor scheduling is required by the VM kernel for the virtual machine.

RAM Ensure that you have enough RAM in the ESX server to service all the virtual machines with a maximum RAM usage. Plan your system so that the ESX server does not need to swap RAM for the virtual machine.

Disk Use the SAN-attached RDM with SCSI pass-through for DB2 data and log storage.

For more information about VMware, see the following publications:

- "Using VMware ESX Server with IBM WebSphere Application Server" at http://www.vmware.com/partners/vmware/ESX_WAS_WP.pdf.

- "Scaling IBM DB2 9 in a VMware Infrastructure 3 Environment" at http://www.vmware.com/pdf/db2_scalability_wp_vi3.pdf.

VMware configuration

Before you can display reports or see the topology for VMware Virtual Infrastructure, you must configure the VMware environment.

You must complete the following general steps:

1. If the VMware Virtual Infrastructure uses SSL certificates for communication, you have to use **keytool** to manually import the SSL certificates into a truststore. Each Virtual Infrastructure data source provides an individual certificate. There is a default truststore registered in the Device server's system properties file. **Keytool** is a tool shipped with the Java run-time environment. For information about how to import the SSL certificates, go to the Tivoli Storage Productivity Center information center. Search for **Importing SSL certificates for VMware**.
2. Add the VMware VI data source. The data source can be a hypervisor (ESX Server or VirtualCenter). Adding the data source is the first step in getting information from VMware Virtual Infrastructure. Adding a VMware data source is like adding a CIM agent or Storage Resource agent.
3. Test the connection to the VMware VI data source. This test ensures that you can access information from the VMware data source.
4. Run a discovery job for the VMware environment. The discovery is needed to retrieve every ESX Server instance that is part of the Virtual Infrastructure that has been added. The discovery mechanism is like a discovery for storage subsystems. Discovery jobs can be scheduled and are performed on the complete list of known VMware data sources.
5. Run a probe job for the ESX Server, hypervisor, and virtual machines. This step gets the detailed information from the hypervisors and virtual machines for IBM Tivoli Storage Productivity Center.
6. Configure alerts for VMware. You can create alerts for the following alert conditions:
   - Hypervisor discovered
   - Hypervisor missing
   - Virtual Machine added
   - Virtual Machine deleted
7. Install the Storage Resource agent on each of the virtual machines you want to monitor. For full functionality, you need two data sources.
8. You can now view VMware reports and VMware topology.

VMware capacity reports

You must probe both the ESX Server and the Data agent or Storage Resource agent on the virtual machines before you can generate accurate reports for disk and file system capacity.

If you have an ESX Server that has 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) have been probed. You also have a physical computer (PC1) that has been probed. The TOTAL capacity for the file system or disk capacity row includes everything - virtualized disks and virtual machines as well as non-virtualized disks and machines.

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 have probed the PC1 (physical computer) and VC1 (virtual computer) but have not probed the ESX Server (hypervisor), the capacity is shown as follows:

Column heading	Capacity	Used Space (calculated as capacity minus free space, any negative values are represented 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 have probed the hypervisor (H1) and PC1 (physical computer) but have not probed the VC1 (virtual computer), 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

Planning for the Virtual I/O Server

You can use the Storage Resource agent to gather information about Virtual I/O Servers. Before you can monitor Virtual I/O Servers, you must plan on how to install the agents in your environment.

Overview

The Virtual I/O Server is part of the IBM PowerVM® hardware feature. The Virtual I/O Server allows the sharing of physical resources between LPARs including Virtual SCSI and virtual networking. Sharing of physical resources allows more efficient utilization of physical resources through sharing between LPARs and facilitates server consolidation.

Note:

- For more information about installing the Data agent or Fabric agent on the Virtual I/O Server, see the IBM Tivoli Storage Productivity Center V4.1.1 Information Center.

Support for Virtual SCSI environment

The Virtual SCSI environment consists of a Virtual SCSI Server Adapter or Adapters that are created on the Virtual I/O Server and mapped to a Virtual I/O Client. If you have a Virtual SCSI environment and the virtual target device has a one to one mapping to a storage volume, Tivoli Storage Productivity Center can extract the storage system serial number and correlate the Virtual SCSI disk on the LPAR to a corresponding storage volume. This function means that the Storage Resource agent installed on the Virtual I/O Client (LPAR), can collect this information. The Linux system must be running on the IBM Power® System.

Note:

- Tivoli Storage Productivity Center only performs the correlation if the virtual target device is a one to one mapping to a single storage system volume.

The storage systems supported are:

- DS8000
- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- DS3000,
- DS4000
- DS5000

After the Storage Resource agent is installed on the Virtual I/O Server, you can see the following information:

- Storage system serial number
- Virtual SCSI Client Adapter
- Virtual SCSI disk
- Volume groups, logical volumes, and file system

Note:

- Currently, the correlation of Virtual SCSI disks that originate from the XIV system is not performed.
- Tivoli Storage Productivity Center can only perform the correlation of Virtual SCSI disks if AIX Multiple Path I/O (MPIO) or SDDPCM multipath driver is used on the Virtual I/O Server.
- If there is one or more Virtual SCSI disks present on an LPAR, Tivoli Storage Productivity Center considers the LPAR as a virtual machine. In this case, Tivoli Storage Productivity Center does not include disk space for this LPAR in the TOTAL disk space reports.

  Disks that are assigned directly through the physical Fibre Channel adapter to the LPAR in a Virtual SCSI environment are listed as virtual disks in Tivoli Storage Productivity Center.

Login requirements

Use the **padmin** user ID when logging in to a Virtual I/O Server, **padmin** is the main administrator ID for Virtual I/O Servers. Upon login, a password change is required. There is no default password to remember.

The Virtual I/O Server does not support users logging in as root to install or configure Tivoli Storage Productivity Center agents. However, you can debug errors generated by agents as a root user. For example, you can run the `service.sh` script to gather debugging information.

Upon logging in to the Virtual I/O Server, you are placed into a restricted Korn shell. The restricted Korn shell works the same way as a regular Korn shell, except you cannot:

- Change the current working directory.
- Set the value of the SHELL, ENV, or PATH variable.
- Specify the path name of the command that contains a redirect output of a command with a >, >|, <>, or >.

As a result of these restrictions, you cannot run commands that are in locations not defined in your PATH variable. These restrictions prevent you from directly sending the output of the command to a file, instead requiring you to pipe the output to the **tee** command.

Installation requirements

You must use the **padmin** user ID to perform the following actions:

- Install the Storage Resource agent. See "Installing the Storage Resource agent on the Virtual I/O Server" on page 235.
- Run the commands to configure, start, and stop the agent. Go to the information center. Search for **managing the daemon**.
- For information about disk space requirements for the Storage Resource agent, see "Hardware requirements" on page 2 (Hardware requirements for the Tivoli Storage Productivity Center agents).

The installation requirements are:

- Tivoli Storage Productivity Center supports Virtual I/O Server version 1.5.2 or later.
- The Virtual I/O Server for POWER5 processor-based systems, must have a minimum of 512 MB of memory.
- The Virtual I/O Server for POWER6 processor-based systems, must have a minimum of 768 MB of memory.

General procedure to install and configure agents

The general steps to install and configure the agents on the Virtual I/O Server are:

1. Install Tivoli Storage Productivity Center.
2. Install and configure the agents on the Virtual I/O Server.
3. When you install the agents, they are automatically started
4. Run a discovery job for the agents on the Virtual I/O Server.
5. Collect information about a Virtual I/O Server by running probes, scans, and ping jobs.
6. View the storage information gathered by the monitoring jobs through the topology viewer and reports that you can generate through Fabric Manager and Data Manager.

See the *IBM Tivoli Storage Productivity Center User's Guide* for more information about how to use the Tivoli Storage Productivity Center user interface to collect

and view information gathered by agents on the Virtual I/O Servers.

Planning for SQL access to IBM Tivoli Storage Productivity Center's views

Use this topic for planning to retrieve storage information that is stored in the Tivoli Storage Productivity Center database repository.

Tivoli Storage Productivity Center stores the information collected by data collection jobs in its DB2 database repository. The information in the database repository is organized into a set of tables. A view is a way of describing data that exists in one or more of these tables.

You can use the Structured Query Language (SQL) to retrieve the information from these views and then use that information in reports and applications outside of Tivoli Storage Productivity Center, such Business Intelligence and Reporting Tools (BIRT), Microsoft Excel, and Crystal Reports. Other applications can also use these views to gather and import information that is collected by Tivoli Storage Productivity Center.

The views are created within a DB2 schema when you install the Tivoli Storage Productivity Center server. A schema defines the views and the fields within each view. The name of the schema containing the views for Tivoli Storage Productivity Center is TPCREPORT. To view this schema using the DB2 Control Center, complete these steps:

1. Start the DB2 Control Center.
2. Expand **All Databases** > **TPCDB** in the left pane.
3. Click **Schemas**.
4. Click **TPCREPORT** in the content pane.

Before you can work with the data stored in views, keep in mind the following planning considerations:

- 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, and the data in a view is determined by a SELECT statement that is run on the base tables.
- The views are available in Tivoli Storage Productivity Center version 4.1 or higher. If you want to retrieve data from the database repositories of previous versions of the product, you must first upgrade to version 4.1 or higher. See Chapter 4, “Upgrading and migrating the Tivoli Storage Productivity Center family,” on page 379 for information about how to upgrade Tivoli Storage Productivity Center.
- You must run data collection jobs to collect data about the storage resources in your environment before you can retrieve data from views. The data collected by probes, scans, pings, switch performance monitors, and subsystem performance monitors populates those views.
- You must have proper authority to access and retrieve data from the views. The views can be accessed by the following users:
  - DB2 administrator
  - The DB2 user that Tivoli Storage Productivity Center uses to store and retrieve information in the database repository
  - A view-only user that uses the operating system credentials associated to DB2 (for example, db2admin/password).

- For information about performing the following tasks, go to the information center. Search for *Accessing views in the database repository*.
  - Use the Business Intelligence and Reporting Tools (BIRT) report designer to create a report based on the information available in Tivoli Storage Productivity Center's exposed views. See the following Web site for more information about BIRT: http://www.eclipse.org/birt/phoenix/.
  - View the report you created in BIRT through the Tivoli Common Reporting v1.2 feature of Tivoli Enterprise Portal. See the following web site for more information about Tivoli Common Reporting: http://www.ibm.com/developerworks/spaces/tcr.
- See the following page on Tivoli Storage Productivity Center's support web site for detailed information about the views provided with Tivoli Storage Productivity Center: http://www.ibm.com/systems/support/storage/software/tpc. Search for **Views Documentation**.

Setting up a view-only user on Windows

This topic provides information about how to set up a user with view-only privileges for the exposed views in the TPCREPORT schema. This procedure is for Windows only.

To set up a user with view-only privileges on Windows, complete these steps:

1. **Create a new user on Windows.** For example, create a user named TPCRPT.
   a. Open **Administrative Tools** > **Computer Management** in the Control Panel.
   b. Click **New User** on the **Action** menu.
   c. Enter information about the new user and click **Create**.
   d. Click **Close**.
2. **Add the new user to the DB2USERS group.** When you add the new user to the DB2USERS group it automatically provides that user with SQL connect authority only.
   a. Open **User Accounts** in the Control Panel.
   b. On the **Users** tab under **Users for this computer**, select the user account name and click **Properties**.
   c. On the **Group Membership** tab, select the **DB2USERS** group and then **OK**.
3. **Use the DB2 Control Center to assign SELECT only authority to the new user**.
   a. Start the DB2 Control Panel.
   b. Expand **TPCDB** > **User and Group Objects** in the left pane.
   c. Select the **DB Users** folder. A list of users appears in the content pane.
   d. Select the new user that you created in step 1.
   e. Select **Change...** from the **Selected** menu while the new user is highlighted.
   f. Click the **View** tab. A list of views appears in the content pane. A green check mark in the columns next to each view indicates that the user can perform the actions represented by those columns. A red circle with a line through it indicates that the user cannot perform the actions represented by those columns.
   g. Do not change the privileges for the views in the TPCREPORT schema. The default settings for the views in the TPCREPORT schema indicate that the new user cannot perform the following SQL commands against those views: INSERT, UPDATE, DELETE, CONTROL. Accept the default settings that restrict the new user from performing actions against the views.

h. Click the **Schema** tab. A list of schemas and the users level of authority for each appears in the content pane.
i. Make sure that the user does not have **CREATEIN**, **DROPIN**, or **ALTERIN** privileges. This check ensures that the user has read-only access to the exposed Tivoli Storage Productivity Center views.

Monitoring your DB2 environment

It is a good practice to monitor your DB2 environment to better understand what is happening inside your DB2 data server. DB2 9.7 includes a new monitoring infrastructure that you can access through new event monitors.

This infrastructure is a superior alternative to the existing system monitor, event monitors, snapshot commands, and snapshot SQL interfaces.

For more information about monitoring your DB2 environment, see http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.wn.doc/doc/c0055019.html.

Planning for the Discovery Library Adapter

Use this topic to plan for using the Discovery Library Adapter.

IBM Tivoli Storage Productivity Center provides a Discovery Library Adapter which can be used in a multiple IBM Tivoli products environment to exchange data with other Tivoli products. The data gathered by IBM Tivoli Storage Productivity Center is put into a common data model format and written to a file using the Identification Markup Language. This file can be loaded into the change and configuration management database (CCMDB) so that other Tivoli products such as Tivoli Application Dependency Discovery Manager or Tivoli Business Service Manager can use that data.

The Discovery Library Adapter also enables Tivoli Application Dependency Discovery Manager to be able to perform a launch in context operation to the IBM Tivoli Storage Productivity Center server. The following IBM Tivoli Storage Productivity Center panels are accessible:

- Topology view for computers
- Topology view for fabrics
- Topology view for switches
- Topology view for storage systems

Tivoli Storage Productivity Center Monitoring Agent

The 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 will be installing the Tivoli Storage Productivity Center Monitoring Agent in the Tivoli Storage Productivity Center environment. You will also be installing the Tivoli Storage Productivity Center Monitoring Agent support files in the IBM Tivoli Monitoring server environment.

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_6.2.2/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_6.2.2fp2/itm622fp2_qsg_en.htm .

This agent is on the installation media disk 1, part 4.

Hardware and software requirements

You must install Tivoli Agent Management Services version 6.2 or later on Windows, AIX or Linux. The following Tivoli Monitoring components are required:

Tivoli Enterprise Monitoring Agent Framework - TEMA

- Tivoli Enterprise Monitoring Agent Framework
- Warehouse Proxy
- Summarization and Pruning Agent
- Tivoli Enterprise Services User Interface Extensions

Tivoli Enterprise Monitoring Server - TEMS

- Tivoli Enterprise Monitoring Server
- Warehouse Proxy
- Summarization and Pruning Agent
- Monitoring Agent support for the corresponding operating system on which the agent is installed

Tivoli Enterprise Portal Server - TEPS

- Tivoli Enterprise Portal Server
- Warehouse proxy
- Monitoring Agent support for the operating system on which the agent is installed
- Summarization and Pruning Agent

Tivoli Enterprise Desktop Client - TEPD

- Tivoli Enterprise Desktop Client
- Warehouse Proxy
- Monitoring Agent support for the operating system on which the agent is installed
- Summarization and Pruning Agent

Tivoli Enterprise Portal Browser Client

IBM Eclipse Help Server

Note: IBM Tivoli Monitor does not support the Tivoli Enterprise Portal Desktop Client (TEPDC) on AIX.

The following disk space is required:

Memory
Approximately 100 MB of RAM per agent instance.

Disk space
Approximately 300 MB of disk space for the IBM Tivoli Storage Productivity Center Monitoring Agent. For information about disk space required for historical data, see the *IBM Tivoli Monitoring Administrator's Guide* at http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.2/itm_admin.htm.

To install the IBM Tivoli Storage Productivity Center Monitoring Agent, you must have Administrator authority on Windows and root authority on AIX and Linux.

For information about installing the IBM Tivoli Enterprise Monitoring Server, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.2/itm_install86.htm.

Tivoli Agent Management Services

The Tivoli Agent Management Services extension of IBM Tivoli Monitoring monitors the health and availability of Storage Resource agents that are run as daemon processes.

The purpose of Agent Management Services is to provide customers with a central view of various Tivoli agents, and to perform the following functions:

- Monitors the health of Storage Resource agents
- Starts, stops, and manages agents

Through these functions, Tivoli Agent Management Services provides the following benefits for agent management:

- Adds a dimension of fault tolerance to agents by keeping a deployed Storage Resource agent up and running.
- Reduces the mean time to recovery of failed agents to seconds, limiting the number of situations that need to be configured to those reported as unrecoverable faults.
- Reduces the number of events flowing to your event managers to those reported by Tivoli Agent Management Services.

For information about Tivoli Monitoring, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.2fp2/itm622fp2_qsg_en.htm .

A Tivoli Monitoring Agent for OS must be installed on every computer that you want to monitor with a Storage Resource agent. The Tivoli Monitoring Agent for OS monitors all agents that have CAP files located in the CAP file directories. The Tivoli Agent Management Services can stop, start, and check the status of agents listed in the CAP file. The CAP file is created when a daemon Storage Resource agent is installed on a system.

For more information about the Tivoli Agent Management Services, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.2fp2/agentmgmtsvcs_intro_tema.htm .

For information about Tivoli Monitoring supported platforms, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.2fp2/ch2.4.htm#ch2.4.

On Windows, the status, start, and stop functions are handled through the Windows services.

The CAP file is installed in the following directories:

Windows

```
%ALLUSERSPROFILE%\Application Data\IBM\CAP
```

For example:

```
Windows 2008: C:\Program Data\IBM\CAP
Windows 2003: C:\Documents and Settings\All Users\Application Data\IBM\CAP
```

Non-Windows

```
/opt/IBM/CAP
```

CAP files are not configurable. You cannot change the location of these directories.

If the creation of the CAP file fails, you receive an error message in the installation log file and trace log file. If the CAP is not created, the lack of a CAP file does not affect the function of the Storage Resource agent. You can reinstall the agent with the **-force** option to create the CAP file.

For more information about the CAP file or Tivoli Agent Management Services, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.2fp2/agentmgmtsvcs_intro_tema.htm .

Tivoli Storage Productivity Center installation

Use this information to learn about the steps required to install IBM Tivoli Storage Productivity Center as a new installation.

Overview

When you install Tivoli Storage Productivity Center, you are also installing Tivoli Integrated Portal and Tivoli Storage Productivity Center for Replication.

External non-WebSphere applications authenticate single sign-on users using the authentication service provided by Tivoli Integrated Portal.

The Tivoli Storage Productivity Center user interface can be launched from Tivoli Integrated Portal using a web browser and Java Web Start. When you launch the Tivoli Storage Productivity Center user interface by using Java Web Start, the user interface is downloaded from the Java Network Launching Protocol (JNLP) servlet within the Tivoli Storage Productivity Center Device server. Once downloaded, the user interface opens a new window.

When you install Tivoli Storage Productivity Center, you can specify a location to install a new instance of Tivoli Integrated Portal or you can reuse an existing Tivoli Integrated Portal instance, as long as its authentication mechanism is the operating system or an LDAP-compliant repository.

When you install Tivoli Storage Productivity Center and Tivoli Integrated Portal, you can elect to use the operating system (OS) or an LDAP-compliant directory as the centralized authentication repository. If you select local OS authentication, then no other Tivoli Integrated Portal-based product is supported in this instance of Tivoli Integrated Portal.

For more information about Tivoli Integrated Portal, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.tip.doc_1.1/welcome_tip_ic.htm.

When you install the Tivoli Storage Productivity Center family, both the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication components are installed. Canceling installation at any time during the installation process results in the uninstallation of all components. For example, if you click **Cancel** during the installation of the Tivoli Storage Productivity Center for Replication component, this action results in the uninstallation of Tivoli Storage Productivity Center, Tivoli Integrated Portal, and Tivoli Storage Productivity Center for Replication.

Fully qualified host names

Tivoli Storage Productivity Center requires fully qualified host names. Some systems might be configured to return a short host name, such as `ddunham` instead of a fully qualified host name, such as `ddunham.myorg.mycompany.com`. Install Tivoli Storage Productivity Center on a computer that has a fully qualified host name. If you install Tivoli Storage Productivity Center on a system without a fully qualified host name, the installation might appear to be successful but the single sign-on feature for Tivoli Integrated Portal, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication might fail to operate.

For information about fully qualified host names, go to the Tivoli Storage Productivity Center for Replication information center. Search for **Checking for a fully qualified host name**.

Note: The Tivoli Storage Productivity Center installation wizard enforces having a fully qualified host name during installation. If you enter a short host name, the installation wizard displays an error message and you are not allowed to continue with the installation.

Static IP address

The Tivoli Storage Productivity Center server requires the use of a static IP address.

Tivoli Storage Productivity Center installation package

The installation images shipped with Tivoli Storage Productivity Center include the Tivoli Monitoring server and DB2.

The Tivoli Monitoring server shipped with Tivoli Storage Productivity Center has a limited license that allows you to use the product to manage or monitor your Tivoli Storage Productivity Center environment. See the product license for specific entitlements.

The Tivoli Monitoring server installation images shipped with Tivoli Storage Productivity Center are:

- ITM 6.2.2.2 Server for Windows
- ITM 6.2.2.2 Server for Linux
- ITM 6.2.2.2 Server for UNIX
- ITM 6.2.2.2 Agent (this is the Tivoli Monitoring Agent for OS)
- ITM 6.2.2.2 National Language Support

The Tivoli Storage Productivity Center documentation does not describe how to install the Tivoli Monitoring server or Tivoli Monitoring Agent for OS. For information about installing these components, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.2fp2/installation.htm .

For information about the supported platforms for Tivoli Monitoring, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.2fp2/ch2.4.htm#ch2.4.

The DB2 installation image shipped with Tivoli Storage Productivity Center are:

- DB2 9.7 for Windows (32-bit)
- DB2 9.7 for Windows (64-bit)
- DB2 9.7 for Linux (64-bit)
- DB2 9.7 for AIX (64-bit)

The Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication installation images are:

- Tivoli Storage Productivity Center for AIX
- Tivoli Storage Productivity Center for Linux
- Tivoli Storage Productivity Center for Windows
- Tivoli Storage Productivity Center Storage Resource Agent
- Tivoli Storage Productivity Center National Language Support

Note: The Tivoli Storage Productivity Center installation image has four parts:

- Disk 1, Part 1 -- Base IBM Tivoli Storage Productivity Center installation program
- Disk 1, Part 2 -- Base IBM Tivoli Integrated Portal and IBM Tivoli Storage Productivity Center for Replication installation programs
- Disk 1, Part 3 -- IBM Tivoli Integrated Portal PTF installation program
- Disk 1, Part 4 -- IBM Tivoli Storage Productivity Center Monitoring Agent installation program for IBM Tivoli Monitoring. This is an optional program and only required if you are using Tivoli Monitoring.

The **disk 1** images part 1, part 2, part 3, and part 4 must be downloaded to the same directory.

The Tivoli Storage Productivity Center for Replication installation images are:

- Tivoli Storage Productivity Center for Replication Two Site Business Continuity License Key for Windows
- Tivoli Storage Productivity Center for Replication Two Site Business Continuity License Key for Linux
- Tivoli Storage Productivity Center for Replication Two Site Business Continuity License Key for AIX
- Tivoli Storage Productivity Center for Replication Three Site Business Continuity License Key for Windows
- Tivoli Storage Productivity Center for Replication Three Site Business Continuity License Key for Linux
- Tivoli Storage Productivity Center for Replication Three Site Business Continuity License Key for AIX

Procedure

If you are installing Tivoli Storage Productivity Center as a new installation, follow these general steps:

1. Install DB2 V9.7. See "Installing DB2" on page 155.
2. Install Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication. See "Installing the IBM Tivoli Storage Productivity Center family" on page 167.

Upgrading and migrating the Tivoli Storage Productivity Center family

You can upgrade previous TotalStorage Productivity Center 3.3.2 or later releases to Tivoli Storage Productivity Center version 4.2.2.1. You can upgrade previous IBM TotalStorage Productivity Center for Replication version 3.4 or later to Tivoli Storage Productivity Center for Replication version 4.2.2.1. This section provides information about upgrading and migrating.

For information about upgrading releases before Tivoli Storage Productivity Center 3.3.2 and Tivoli Storage Productivity Center for Replication 3.4, see the Tivoli Storage Productivity Center 4.2.2 information center.

When you upgrade Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication, if a component fails to upgrade, then just the component is not upgraded.

If Tivoli Storage Productivity Center exists but Tivoli Storage Productivity Center for Replication does not, then this upgrade is a fresh installation of Tivoli Storage Productivity Center for Replication and an upgrade of Tivoli Storage Productivity Center. If a failure occurs, an error message is displayed but there is no rollback.

If Tivoli Storage Productivity Center for Replication exists and Tivoli Storage Productivity Center does not, then this upgrade is a fresh installation of Tivoli Storage Productivity Center and an upgrade of Tivoli Storage Productivity Center for Replication. If there is a failure in the Tivoli Storage Productivity Center installation, a rollback occurs of Tivoli Storage Productivity Center (Tivoli Storage Productivity Center for Replication remains).

Tivoli Storage Productivity Center for Replication is no longer a stand-alone application. Tivoli Storage Productivity Center Version 4.2 or later installs Tivoli Integrated Portal and Tivoli Storage Productivity Center for Replication.

With version 4.2 or later, Tivoli Storage Productivity Center for Replication no longer supports DB2 as the datastore for its operational data. Tivoli Storage Productivity Center for Replication uses an embedded repository for its operational data. The Tivoli Storage Productivity Center for Replication 4.2.2.1 installation program automatically migrates any Tivoli Storage Productivity Center for Replication operational data in an existing Tivoli Storage Productivity Center for Replication DB2 database to the Tivoli Storage Productivity Center for Replication embedded repository as part of upgrading to Tivoli Storage Productivity Center for Replication from an earlier version. New Tivoli Storage Productivity Center for Replication 4.2.2.1 installations use the embedded repository by default.

For more information about upgrading and migrating, see Chapter 4, "Upgrading and migrating the Tivoli Storage Productivity Center family," on page 379.

Upgrading IBM Tivoli Storage Productivity Center from version 3

If you are upgrading Tivoli Storage Productivity Center from version 3.3.2 or later, follow these general steps.

You can upgrade the following TotalStorage Productivity Center releases to Tivoli Storage Productivity Center version 4.2.2.1:

- TotalStorage Productivity Center version 3.3.2 or later to 4.2.2.1.
- TotalStorage Productivity Center for Replication version 3.4 or later to 4.2.2.1. For information about upgrading, see "Upgrading Tivoli Storage Productivity Center for Replication from V3.4 or later to V4.2.2.1 or later" on page 383.

Follow these general steps to upgrade:

1. If you are using DB2 version 8, you must migrate to DB2 version 9.5. See "Upgrading DB2 8.2 to DB2 9.7" on page 386.
2. If you have Agent Manager 1.2, you can optionally upgrade to Agent Manager 1.3.2. For information about upgrading, see the Tivoli Storage Productivity Center 4.1.1 Information Center.

   Note: If your Agent Manager 1.2 is running fine, do not upgrade.
3. Upgrade Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication to version 4.2.2.1. See Chapter 4, "Upgrading and migrating the Tivoli Storage Productivity Center family," on page 379.
4. Migrate the Tivoli Storage Productivity Center database.

Upgrading Tivoli Storage Productivity Center scenarios

This topic describes the various scenarios that are valid when upgrading Tivoli Storage Productivity Center.

Upgrade scenarios for Tivoli Storage Productivity Center V4.1.1, upgrade server to V4.2 or later

When you have a Tivoli Storage Productivity Center 4.1.1 (or earlier) server installed, these are the valid upgrade scenarios.

Table 27. Upgrade scenarios for Tivoli Storage Productivity Center V4.1.1 (or earlier)

Tivoli Storage Productivity Center server V4.1.1 (or earlier) installed	Upgrade Tivoli Storage Productivity Center server to V4.2 or later
Data agent or Fabric agent or both installed on the server (version 4.1.1 or earlier)	You have a choice: • Leave the Data agent or Fabric agent at the down level version • Migrate the Data agent or Fabric agent to Storage Resource agent
A daemon or RXA Storage Resource agent is installed	You have a choice to upgrade or not upgrade the Storage Resource agent to 4.2.
No agent installed	The default Storage Resource agent is installed.

Note:

- You cannot use the Tivoli Storage Productivity Center V4.1.1 (or earlier) installation program on a Tivoli Storage Productivity Center V4.2 or later system.
- You can use the Tivoli Storage Productivity Center V4.2 installation program to install a local Storage Resource agent on a system that does not have the Tivoli Storage Productivity Center server installed. You can also use the Tivoli Storage Productivity Center GUI to deploy the Storage Resource agents (from the server system).

Upgrade scenarios for Tivoli Storage Productivity Center V4.2 or later, install Storage Resource agent

When you have a Tivoli Storage Productivity Center V4.2 or later server installed, and are installing a Storage Resource agent, these are the valid upgrade scenarios.

Table 28. Upgrade scenario for Tivoli Storage Productivity Center V4.2 or later

Tivoli Storage Productivity Center server V4.2 or later installed	Use Tivoli Storage Productivity Center V4.2 installation program on server or to install a Storage Resource agent on a local machine
Data agent or Fabric agent or both installed (version 4.1.1 or earlier)	You can elect to migrate the Data agent or Fabric agent to a Storage Resource agent
A daemon or RXA Storage Resource agent is installed	The Storage Resource agent is upgraded to V4.2
No agent installed	The default Storage Resource agent is installed.

Upgrade scenarios for Tivoli Storage Productivity Center agents using the GUI (V4.2 or later)

When you are upgrading a Tivoli Storage Productivity Center agent using the GUI, these are the valid upgrade scenarios.

Table 29. Upgrade scenario for Tivoli Storage Productivity Center agents using V4.2 GUI

Tivoli Storage Productivity Center agent	Upgrade agent using Tivoli Storage Productivity Center V4.2 GUI
Data agent or Fabric agent or both on local computer	Not supported. You can migrate a Data agent or Fabric agent to a Storage Resource agent.
Storage Resource agent V4.1 on local computer	The Storage Resource agent is upgraded to the latest V4.2 level.

Table 29. Upgrade scenario for Tivoli Storage Productivity Center agents using V4.2 GUI (continued)

Tivoli Storage Productivity Center agent	Upgrade agent using Tivoli Storage Productivity Center V4.2 GUI
Storage Resource agent V4.2 on local computer	The Storage Resource agent is upgraded to the latest V4.2 level (use the **force** option).

Upgrade scenarios for Tivoli Storage Productivity Center agents using the Storage Resource agent command (V4.2 or later)

When you are upgrading a Tivoli Storage Productivity Center agent using the command line, these are the valid upgrade scenarios. For information about the Storage Resource agent command, see "Upgrading Storage Resource agents using a command" on page 406.

Table 30. Upgrade scenario for Tivoli Storage Productivity Center agents using Tivoli Storage Productivity Center V4.2 command line

Tivoli Storage Productivity Center agent	Upgrade agent using Tivoli Storage Productivity Center V4.2 or later command line
Data agent or Fabric agent or both on local computer	Not supported. Migrate the Data agent or Fabric agent to a Storage Resource agent. See "Migrating Data agents and Fabric agents to Storage Resource agents" on page 407.
Storage Resource agent V4.1 on local computer	The Storage Resource agent is upgraded to the latest V4.2 or later level (cannot change **commtype**).
Storage Resource agent V4.2 on local computer	The Storage Resource agent is upgraded to the latest V4.2 or later level (must use **force** option).

Platform support for Tivoli Storage Productivity Center agents with V4.2

Table 31. Platform support for Tivoli Storage Productivity Center agents

Platform	Earliest level of Data agent or Fabric agent	Latest level of Data agent or Fabric agent	New Data agent or Fabric agent install	Upgrade Data agent or Fabric agent	Earliest Storage Resource agent	Latest Storage Resource agent
Linux	3.3.2	4.1.1	Not supported	Not supported	4.1	4.2
Windows with NetWare	3.3.2	4.1.1	Local Data agent install	Local Data agent upgrade	Not supported	Not supported
HP-UX PA-RISC (11iv2)	3.3.2	4.1.1	Local Data agent or Fabric agent install	Local Data agent or Fabric agent upgrade	Not supported	Not supported
HP-UX Itanium (11iv2)	3.3.2	4.1.1	Remote Data agent install	Fabric agent not supported	Not supported	Not supported
HP-UX Itanium (11iv3)	Not supported	Not supported	Not supported	Not supported	4.2	4.2
AIX	3.3.2	4.1.1	Not supported	Not supported	4.1	4.2

Table 31. Platform support for Tivoli Storage Productivity Center agents (continued)

Platform	Earliest level of Data agent or Fabric agent	Latest level of Data agent or Fabric agent	New Data agent or Fabric agent install	Upgrade Data agent or Fabric agent	Earliest Storage Resource agent	Latest Storage Resource agent
AIX with Virtual I/O Server	3.3.2	4.1.1	Local Data agent and Fabric agent install	Local Data agent and Fabric agent upgrade	Not supported	Not supported
Oracle Solaris	3.3.2	4.1.1	Not supported	Not supported	4.2	4.2

Upgrade scenarios for Tivoli Storage Productivity Center V4.1.1 agents with V4.1.1 installation program

When you have a Tivoli Storage Productivity Center 4.1.1 (or earlier) agent installed, these are the valid upgrade scenarios.

Table 32. Upgrade scenario for Tivoli Storage Productivity Center V4.1.1. agents with V4.1.1 installation program

Tivoli Storage Productivity Center agent V4.1.1 (or earlier) installed	Use Tivoli Storage Productivity Center V4.1.1 installation program on non-TPC server for local install
Data agent or Fabric agent or both installed (version 4.1.1 or earlier) on local machine	If the Data agent or Fabric agent is down level, the agent will be upgraded to the latest V4.1.1 level. If the Data agent or Fabric agent is at the latest V4.1.1 level, you see a message that the agent is already installed.
A daemon or RXA Storage Resource agent is installed on local machine	If the Storage Resource agent is at the latest V4.1.1 level, the Storage Resource agent is left as is. If the Storage Resource agent is not at the latest V4.1.1 level, the agent is migrated to a Data agent or Fabric agent.
No agent installed	The Data agent or Fabric agent is installed.

Chapter 2. Installing the Tivoli Storage Productivity Center family

This section describes how to install DB2, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication. This section also provides information about the preparation work required before installing the Tivoli Storage Productivity Center family.

Installing Tivoli Storage Productivity Center overview

This section provides information about installing Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication.

Overview

You can install all the Tivoli Storage Productivity Center components using typical installation or custom installation. Custom installation allows you to see what components are being installed and where the components are being installed as well as set different passwords for user IDs.

Tivoli Storage Productivity Center for Replication is no longer a stand-alone application. Tivoli Storage Productivity Center Version 4.2 or later installs Tivoli Integrated Portal and Tivoli Storage Productivity Center for Replication.

When you install Tivoli Storage Productivity Center, you have these installable components:

- Database repository
- Data Server and Device server
- Graphical user interface (GUI)
- Command Line Interface (CLI)
- Storage Resource agent

You need to install most of these components in order to get Tivoli Storage Productivity Center to work. The CLI is considered optional. The GUI is installed wherever a user might want to control Tivoli Storage Productivity Center.

If you are using a remote database for Tivoli Storage Productivity Center, you must install the database schema on that remote database after you have installed DB2 on that remote machine.

After Tivoli Storage Productivity Center is installed, the installation program will start the Tivoli Storage Productivity Center for Replication installation wizard.

Tivoli Storage Productivity Center also supports a Java Web-based GUI which gives you the option of installing the GUI wherever you want to use Tivoli Storage Productivity Center. For information about the Java Web-based GUI, see “Starting the Tivoli Storage Productivity Center GUI using a web browser” on page 347.

If you are using IPv6, see “Planning for Internet Protocol Version 6” on page 51.

When you install the Tivoli Storage Productivity Center family, both the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication

© Copyright IBM Corp. 2001, 2012

components are installed. Canceling installation at any time during the installation process results in the uninstallation of all components. For example, if you click **Cancel** during the installation of the Tivoli Storage Productivity Center for Replication component, this action results in the uninstallation of Tivoli Storage Productivity Center, Tivoli Integrated Portal, and Tivoli Storage Productivity Center for Replication.

DB2 is a separate installation program from Tivoli Storage Productivity Center. The database repository, Data Server, Device server, Storage Resource agents, GUI, and CLI are all components of Tivoli Storage Productivity Center. They can be installed all at once or at different times.

If you install the Tivoli Storage Productivity Center components at different times, you must install the components in this order:

1. DB2 (DB2 is required for Tivoli Storage Productivity Center)
2. Tivoli Storage Productivity Center database repository on the DB2 server (installed using the Tivoli Storage Productivity Center installation program)
3. Data server and Device server (installed using the Tivoli Storage Productivity Center installation program)
4. Storage Resource agent, GUI, or CLI (you can install one, all, or a combination of these components using the Tivoli Storage Productivity Center installation program)

Static IP address

The Tivoli Storage Productivity Center server requires the use of a static IP address.

Fully qualified host names

Tivoli Storage Productivity Center requires fully qualified host names. Some systems might be configured to return a short host name, such as `ddunham` instead of a fully qualified host name, such as `ddunham.myorg.mycompany.com`. Tivoli Storage Productivity Center must be installed on a computer that has a fully qualified host name. If you install Tivoli Storage Productivity Center on a system without a fully qualified host name, the installation might appear to be successful but the single sign-on feature for Tivoli Integrated Portal, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication might fail to operate. For information about fully qualified host names, go to the Tivoli Storage Productivity Center information center. Search for **Checking for a fully qualified host name**.

Note: The Tivoli Storage Productivity Center installation wizard enforces having a fully qualified host name during installation. If you enter a short host name, the installation wizard displays an error message and you are not allowed to continue with the installation.

Accessing Tivoli Storage Productivity Center for Replication

If you install Tivoli Storage Productivity Center using operating system authentication, the Tivoli Storage Productivity Center user group is not automatically added to the Tivoli Storage Productivity Center for Replication administration role. If you want to start Tivoli Storage Productivity Center for Replication from Tivoli Storage Productivity Center using the same user ID, you must explicitly add the Tivoli Storage Productivity Center user group (that has that

user ID) to Tivoli Storage Productivity Center for Replication. For an example of adding a user group to Tivoli Storage Productivity Center for Replication, go to the Tivoli Storage Productivity Center information center. Search for **Adding a Tivoli Storage Productivity Center user group name to Tivoli Storage Productivity Center for Replication**.

InstallShield limitations

These are some limitations you might encounter when installing Tivoli Storage Productivity Center:

- When running the Tivoli Storage Productivity Center installation program on Oracle Solaris systems, some of the graphical elements in the installation panels might not display correctly. For example, in the Select the type of installation you want to run panel, the **TPC Installation Location** button might appear truncated. This behavior is due to an InstallShield limitation on Solaris systems.
- When using the Tivoli Storage Productivity Center installation program to install IBM Tivoli Integrated Portal on AIX systems, the progress bar incorrectly indicates that IBM Tivoli Integrated Portal installation is 100% complete even though it is not yet complete. You must continue to wait until installation is complete. This behavior is due to an InstallShield limitation on AIX systems that prevents the progress bar from correctly reflecting the installation progress.
- When you install Tivoli Storage Productivity Center on AIX, the progress bar can incorrectly display 100% for any component that is installed even though the installation is not complete. This behavior does not affect the installation.

Installing legacy Data agents and Fabric agents

The Tivoli Storage Productivity Center 4.2 installation program does not support installation of the Data agent or Fabric agent. If you want to install the legacy Data agent or Fabric agent, you must have a previous Tivoli Storage Productivity Center installation program that supports installing the Data agent or Fabric agent.

If you are installing local legacy Data agents and Fabric agents, use the local uninstallation program to uninstall these agents.

Installing DB2

This topic describes how to perform a new installation of IBM DB2 Version 9.7 for Linux, UNIX, and Windows. DB2 is required for IBM Tivoli Storage Productivity Center. IBM Tivoli Storage Productivity Center for Replication does not use DB2 as the database repository.

Note:

- If you are installing DB2 9.7, do not use DB2 9.7 with Fix Pack 1 or Fix Pack 2. These fix packs cause issues with Tivoli Storage Productivity Center.
- For information about configuring DB2, AIX, and Linux for IPv6-only environment, see “Configuring DB2, AIX, and Linux for IPv6-only environment” on page 377.

Preparing to install DB2

Before you install DB2, it is important to understand the default user IDs and groups created. The DB2 administrator user ID and password created is the user

ID and password you use to install IBM Tivoli Storage Productivity Center. You can use the default user ID and password or provide your own user ID and password.

Note: Windows platforms do not support systems configured for IPv6–only (Windows is enabled for both IPv6 and IPv4.)

Preparing to install DB2 on Windows

If you have an administrator user ID that you want to use for installing DB2 and IBM Tivoli Storage Productivity Center, make sure this user ID is a member of DB2ADMNS group and the Administrator group.

When you install DB2, two groups are created:

- DB2ADMNS - this group has all the required user rights assigned for administrative authority.
- DB2USERS - this group has user rights assigned for DB2 users.

Assign users that need administrative authority to the DB2ADMNS group and DB2 users to the DB2USERS group. When you assign a user ID to either group, that user ID has all the user rights required for that group; you do not have to assign individual user rights to each user ID.

Important:

- The Windows user ID "Administrator" does not have the authority to create the Tivoli Storage Productivity Center database TPCDB. The DB2 administrator user ID that you use when you install IBM Tivoli Storage Productivity Center must have DB2 administrative rights. The user ID used to install IBM Tivoli Storage Productivity Center must be a member of the DB2ADMNS and Administrators groups. The DB2 user ID db2admin created when you install DB2 has the proper authority for Tivoli Storage Productivity Center. For valid DB2 user IDs and passwords, see http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.admin.dbobj.doc/doc/c0007248.html.

  Note: When you assign the Windows user ID Administrator to the DB2ADMNS group, this authority becomes active the next time that the Administrator user ID logs in.
- If you are installing Tivoli Storage Productivity Center, the user ID that you use must be a member of the DB2ADMNS group and Windows Administrators group. If you do not use the default user ID `db2admin` to install Tivoli Storage Productivity Center, then the alternate user ID must also be a member of these groups.

Preparing to install DB2 on UNIX or Linux

Three users and groups are required to operate DB2 on UNIX or Linux: the instance owner, the fenced user, and the DB2 administration server user. These users and groups are automatically created during the installation of DB2.

The user ID you use to install IBM Tivoli Storage Productivity Center must be the instance owner of the instance that you want to use. This user ID and password are created when you install DB2.

Table 33. DB2 Setup wizard default user and group names

Required user	Default user name	Default group name	Description
instance owner	db2inst1	db2iadm1	Created in the instance owner home directory. This user ID controls all DB2 processes and owns all file systems and devices used by the databases contained within the instance.
fenced user	db2fenc1	db2fadm1	Used to run user-defined functions (UDFs) and stored procedures outside of the address space used by the DB2 database.
DB2 administration server user	dasusr1	dasadm1	Used to run the DB2 administration server on your system. Used by the DB2 GUI tools to perform administration tasks. This user does not contain any databases and there is only one administration server per machine.

Note: For valid DB2 user IDs and passwords, see http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.admin.dbobj.doc/doc/c0007248.html.

Installing DB2 on Windows

Use this information to install DB2 on Windows. DB2 must be installed before you install IBM Tivoli Storage Productivity Center. You must also register the DB2 license key.

To install DB2, follow these steps:

1. Log on with a user ID with Administrator authority on Windows.
2. Insert the DB2 product CD into the CD-ROM drive or use the extracted DB2 installation image. Windows Autorun starts the Launchpad.

   The DB2 installation program starts within 15 - 30 seconds if you have Windows Autorun mode set on your system. If the installation program does not start, do one of the following steps:

   - Use a command prompt to change to the root of the CD-ROM and enter:

     ```
     setup.exe
     ```

   - Use Windows Explorer, open the root of the CD-ROM, and double-click the `setup.exe` file.
3. On the DB2 Setup Launchpad panel, click **Install a Product**. In the content pane, click **Install New** for **DB2 Enterprise Server Edition Version 9.7**.

4. On the Welcome to the DB2 Setup wizard for DB2 Enterprise Server Edition, Version 9.7 panel, click **Next**.
5. On the Software License Agreement panel, read the license agreement and select **I accept the terms in the license agreement** if you agree to the terms. Click **Next**. If you click **Cancel**, the installation program ends without installing any programs.
6. In the Select the installation type panel, select **Typical**. Click **Next**.
7. The Select installation, response file creation, or both panel is displayed. Select **Install DB2 Enterprise Server Edition on this computer and save my settings in a response file**. Enter a response file name or accept the default. Click **Next**.
8. In the Select the installation folder panel, enter a directory or use the default. (For ease of maintenance, use the default directory.) Click **Next**.
9. The Set user information for the DB2 Administration Server panel is displayed. Enter the following information for User information:
   a. Domain (leave this field blank).
   b. User name (Type a user name of the DB2 user account that you want to use. DB2 adds this user ID to the DB2ADMNS group.)
   c. Password and confirm password

   Select the check box for **Use the same user name and password for the remaining DB2 services**. The user ID and password are used to install Tivoli Storage Productivity Center. Click **Next**.
10. The Configure DB2 instances panel is displayed. Click **Next** to accept the default instance of DB2.

    Note: This panel appears if you are installing DB2 on a computer that belongs to an accessible domain.
11. The Prepare the DB2 tools catalog panel is displayed. Accept the defaults and click **Next**. (The default is to not prepare the DB2 tools catalog.)
12. The Set up notifications panel is displayed. Clear the check box for **Set up your DB2 server to send notifications**. Click **Next**.
13. The Enable operating system security for DB2 objects panel is displayed. Accept the defaults for this panel and click **Next**. (The default is to enable operating system security.)

    Note: If you have installed DB2 before on this system, the DB2ADMS group might still be on the machine. If so, you see a warning dialog when you click **Next** in the security panel. The message is:

    ```
    Warning
    The group; name "DB2ADMS" already exists
    in the system and will be granted complete
    access to the DB2 folders and objects through
    the operating system.
    Click OK to continue installation, or click
    Cancel to input another group name.
    ```
14. The Start copying files and create response file panel is displayed. After you have reviewed the summary information, click **Finish**.
15. You see a progress window as DB2 is installed. Wait for installation to complete.
16. The Setup is complete panel is displayed. Read the information about the panel and click **Next**.

17. The Install additional products panel is displayed. Do not install additional products at this time. Click **Finish**.
18. The setup wizard closes and you are back on the DB2 Welcome screen which now has "Welcome to First Steps" in the left pane. This panel includes a **Product Updates** button in the left pane. Click the button to see if there are any product updates. Click **Exit** when you have finished checking for updates.
19. From the Tivoli Storage Productivity Center download image or DVD package, find the file for the DB2 Enterprise Server Edition, Version 9.7 license key: `db2ese_o.lic`. You must register the license key for DB2. Ensure that you understand and agree to the license conditions for DB2. The license terms are located in the `/db2/license/Windows` directory.

    You can register the license key by using one of the following methods:

    From the command line:

    To register the license key, enter the **db2licm** command from the command line:

    ```
    db2licm -a x:\db2\license\license_filename.lic
    ```

    where *x* represents the DVD drive that contains the DB2 license key file and *license_filename.lic* is the license key file: `db2ese_o.lic`.

    From the License Center:

    To register the license key, follow these steps:

    a. Start the DB2 Control Center and click **Tools** > **License Center**.
    b. Select the system and installed product for DB2, and then click **License** > **Add**.
    c. In the Add License window, select the directory containing the license file.
    d. Select the license file, and click **OK**.
20. Restart the system.

Installing DB2 on UNIX or Linux - GUI install

This topic describes how to install DB2 on UNIX or Linux using a GUI installation program.

Note: You must have the X11 graphic capability installed before installing DB2 using the GUI.

To install DB2, complete the following steps:

1. Log in as a user ID with root authority.
2. If you are not using a CD to install DB2, you can skip to step 5 and run **db2setup** from the directory where you have the DB2 source installation program. If you are installing DB2 using the CD, create a mount point or choose an existing mount point. To create a mount point called `/cdrom`, enter the following command: **mkdir /cdrom**.
3. Insert the DB2 CD into the CD-ROM drive. Mount the CD-ROM file system at the wanted mount point.

    On AIX, you can use the **crfs** command to add an entry to `/etc/filesystems` for the CD-ROM mount point. Run the following commands:

    ```
    /usr/sbin/crfs -v cdrfs -p ro -d'cd0' -m'/cdrom' -A'no'
    mount /cdrom
    ```

The **crfs** command only has to be run once for a given mount point, and after that you can use **mount** and **umount** for each CD or DVD you put in the drive, for example, **mount /cdrom** and **umount /cdrom**.

On Linux, run these commands:

```
mkdir /cdrom
mount -o ro /dev/cdrom /cdrom
```

4. Change to the directory where the CD-ROM is mounted by entering the following command: **cd /cdrom**.
5. Enter the **./db2setup** command to start the DB2 Setup wizard.
6. On the DB2 Setup Launchpad panel, click **Install a Product**. In the content pane, click **Install New** for **DB2 Enterprise Server Edition Version 9.7**.
7. In the Welcome to the DB2 Setup wizard panel, click **Next**.
8. On the Software License Agreement panel, read the license agreement. Select **Accept** if you agree with the terms of the license agreement. Click **Next**. If you click **Cancel**, the installation program ends without installing any programs.
9. In the Select the installation type panel, select **Typical**. Click **Next**.
10. The Select installation, response file creation, or both panel is displayed. Select **Install DB2 Enterprise Server Edition on this computer and save my settings in a response file**. Enter a response file name or accept the default. Click **Next**.
11. In the Select the installation directory panel, enter a directory or use the default. (For ease of maintenance, use the default directory.) Click **Next**.
12. The Install the IBM Tivoli System Automation for Multiplatforms (SA MP) panel is displayed. Select **Do not install SA MP**. Click **Next**.
13. The Set user information for the DB2 Administration Server panel is displayed. Select **New user**. Enter the following information for a new user:
    a. User name (Type a user name and password of the DB2 user account that you want to use. DB2 adds this user ID to the DB2ADMNS group. If this user ID does not exist, DB2 creates it.)
    b. Group name
    c. Password and confirm password
    d. Home directory

    Select the options **Use default UID** and **Use default GID**. The DAS user is used to administer the DB2 database.

    Click **Next**.
14. The Set up a DB2 instance panel is displayed. Select **Create a DB2 instance**. Click **Next**.
15. The Set up partitioning options for the DB2 instance panel is displayed. Select **Single partition instance**. Click **Next**.
16. The Set user information for the DB2 instance owner panel is displayed. The DB2 instance owner user is the user you enter when installing IBM Tivoli Storage Productivity Center. Select **New user**. Enter the following information:
    - User name
    - Group name
    - Password and confirm password
    - Home directory

    Select the options **Use default UID** and **Use default GID**. The DAS user is used to administer the DB2 database.

    Click **Next**.

17. The Set user information for the fenced user panel is displayed. Select **New user**. Enter the following information:
    - User name
    - Group name
    - Password and confirm password
    - Home directory

    Select the options **Use default UID** and **Use default GID**. The DAS user is used to administer the DB2 database.

    Click **Next**.
18. The Prepare the DB2 tools catalog panel is displayed. Select **Do not prepare the DB2 tools catalog**. Click **Next**.
19. The Set up notifications panel is displayed. Select **Do not set up your DB2 server to send notifications at this time**. Click **Next**.
20. The Start copying files and create response file panel is displayed. After you have reviewed the summary information, click **Finish**.
21. You see a progress window as DB2 is installed. Wait for installation to complete.
22. The Setup has completed successfully panel is displayed. Read the information about the panel and click **Finish**.
23. The setup wizard closes.
24. From the Tivoli Storage Productivity Center download image or DVD package, find the file for the DB2 Enterprise Server Edition, Version 9.7 license key: `db2ese_o.lic`. You must register the license key for DB2. Ensure that you understand and agree to the license conditions for DB2. The license terms are located in the `/db2/license/Windows` directory.

    You can register the license key by using one of the following methods:

    From the command line:

    To register the license key, enter the **db2licm** command from the command line:

    ```
    db2licm -a x:\db2\license\license_filename.lic
    ```

    where *x* represents the DVD drive that contains the DB2 license key file and *license_filename.lic* is the license key file: `db2ese_o.lic`.

    From the License Center:

    To register the license key, follow these steps:

    a. Start the DB2 Control Center and click **Tools** > **License Center**.
    b. Select the system and installed product for DB2, and then click **License** > **Add**.
    c. In the Add License window, select the directory containing the license file.
    d. Select the license file, and click **OK**.

Note: After you have completed installing DB2, edit the file `/etc/group` and add **root** to the `db2iadm1` group. The `db2iadm1` line in `/etc/group` looks like the following line:

```
db2iadm1:x:102:root
```

Installing DB2 on AIX using the command line

This topic describes how to install DB2 on AIX using a command line.

To install DB2, complete the following steps:

1. Log in as a user ID with root authority.
2. If you are not using a CD to install DB2, you can skip to step 5 and work from the directory where you have the DB2 source installation program. If you are installing DB2 using the CD, create a mount point or choose an existing mount point. To create a mount point called `/cdrom`, entering the following command: **mkdir /cdrom**.
3. Insert the DB2 CD into the CD-ROM drive. Mount the CD-ROM file system at the wanted mount point.

   On AIX, you can use the **crfs** command to add an entry to `/etc/filesystems` for the CD-ROM mount point. Run the following commands:

   ```
   /usr/sbin/crfs -v cdrfs -p ro -d'cd0' -m'/cdrom' -A'no'
   mount /cdrom
   ```

   The **crfs** command only has to be run once for a given mount point, and after that you can use **mount** and **umount** for each CD or DVD you put in the drive, for example, **mount /cdrom** and **umount /cdrom**.
4. Change to the directory where the CD-ROM is mounted by entering the following command: **cd /cdrom**.
5. Install the DB2 Manager.
   a. Create a temporary directory (for example, `db2temp`) to hold the DB2 installer tar file and untarred files. These files require 2-3 GB of hard disk drive space.
   b. Copy or download the DB2 installer into `db2temp`. The name of the file varies depending on the location from where the file is downloaded or copied and the language to which it is associated.
   c. (optional) Use the appropriate extracting tool to extract the DB2 installer file if it is compressed. For example, if the name of the file is `v9fp2_aix_ese.tar.gz`, type **gunzip v9fp2_aix_ese.tar.gz**. In this example, `v9fp2_aix_ese.tar` appears in the `db2temp` directory.
   d. Make an installation directory. For example, enter **mkdir install**.
   e. Change to the new installation directory. Enter **cd install**.
   f. Extract the DB2 installation file. For example, if the name of the installation file is `v9fp2_aix_ese.tar`, enter **tar xvf ../v9fp2_aix_ese.tar**.
   g. Go to the directory that contains the DB2 installation program. Enter **cd /db2temp/installer/ese/disk1**.
   h. Run the command-line installation program. Enter **./db2_install**.
   i. Select the product to install: **DB2.ESE**. A number of messages appear in the command window. The `Installation Summary` eventually appears and indicates a successful installation. DB2 is installed in `/opt/IBM/db2`.
6. Create users and groups for use with DB2.
   a. Enter **mkgroup id=999 db2iadm1**.
   b. Enter **mkgroup id=998 db2fadm1**.
   c. Enter **mkgroup id=997 dasadm1**.
   d. Enter **mkuser id=1004 pgrp=db2iadm1 groups=db2iadm1 home=/home/db2inst1 db2inst1**.
   e. Enter **mkuser id=1003 pgrp=db2fadm1 groups=db2fadm1 home=/home/db2fenc1 db2fenc1**.

f. Enter **mkuser id=1002 pgrp=dasadm1 groups=dasadm1 home=/home/dasusr1 dasusr1**.

g. Verify the owner of the directories. To do this step, enter **ls -ld /home/db2inst1**. The directory ownership appears as follows:

- `/home/db2inst1` shows group `db2iadm1` and user `db2inst1`
- `/home/db2fenc1` shows group `db2fadm1` and user `db2fenc1`
- `/home/dasusr1` shows group `dasadm1` and user `dasusr1`

If the directory ownership is not correct, run the following commands as appropriate:

```
chown -R db2inst1:db2iadm1 /home/db2inst1
chown -R db2fenc1:db2fadm1 /home/db2fenc1
chown -R dasusr1:dasadm1 /home/dasusr1
```

h. Enter **passwd db2inst1** *password*, where *password* represents the password you want to use for the DB2 instance.

i. Enter **pwdadm -f NOCHECK db2inst1**.

j. Enter **passwd db2fenc1** *password*, where *password* represents the password you want to use for the fenced user.

k. Enter **pwdadm -f NOCHECK db2fenc1**.

l. Enter **passwd dasusr1** *password*, where *password* represents the password you want to use for the DB2 administration server (DAS) user.

m. Enter **pwdadm -f NOCHECK dasusr1**.

n. Enter **chgroup users=db2inst1,root db2iadm1**.

7. Create a DB2 Administrative Server (DAS). Enter **/opt/db2/V9.5/instance/dascrt -u dasusr1**.

8. Create a DB2 instance:

a. Enter **/opt/db2/V9.5/instance/db2icrt -a server -u db2fenc1 db2inst1**.

b. Enter **. /home/db2inst1/sqllib/db2profile**.

9. Change the default location for database repositories. By default, this location is `/home/db2inst1`. `/home` is typically not large enough for database repositories. Choose a file system with enough free space (30 GB or higher) to contain the IBM Tivoli Storage Productivity Center repository. To change the default location, complete the following steps:

a. Enter **db2 update dbm cfg using DFTDBPATH** *new_repository_path* **IMMEDIATE**, where *new_repository_path* represents the new location for the repository.

b. Enter **chown -R db2inst1:db2iadm1** *new_repository_path* to assign ownership to `db2inst1` and permission to anyone in `db2iadm1`(same as the ownership for `/home/db2inst1`).

10. Configure DB2 communication:

a. Edit `/etc/services` and verify or add the following line at the end of the file: `db2c_db2inst1 50000/tcp`

b. Enter **db2 update dbm cfg using svcename db2c_db2inst1**.

c. Enter **db2set**. An example of the output is: DB2COMM=tcpip.

11. From the Tivoli Storage Productivity Center download image or DVD package, find the file for the DB2 Enterprise Server Edition, Version 9.7 license key: `db2ese_o.lic`. You must register the license key for DB2. Ensure that you understand and agree to the license conditions for DB2. The license terms are located in the `/db2/license/Windows` directory.

To register the license key, enter the **db2licm** command from the command line:

```
db2licm -a x:\db2\license\license_filename.lic
```

where *x* represents the DVD drive that contains the DB2 license key file and *license_filename.lic* is the license key file: `db2ese_o.lic`.

12. Enter the following commands to restart DB2:

```
db2stop force
db2 terminate
db2start
```

13. Enter the following commands to test your DB2 instance:
    a. Source the environment: **. /home/db2inst1/sqllib/db2profile**.
    b. Get information about the instance and DB2: **db2level**.
    c. Create a test database: **db2 create db test**.
    d. List information about all the databases created for this instance: **db2 list db directory**
    e. Connect to the test database: **db2 connect to test user db2inst1 using** *password* where *password* represents the password you defined in step 6 for the DB2 instance.
    f. Disconnect from the test database: **db2 disconnect test**.
    g. Drop the test database: **db2 drop db test**

Installing DB2 on AIX 6.1 POWER7

This topic describes how to install DB2 on AIX 6.1 POWER7.

To install DB2, complete the following steps. For more information about DB2 installation, see http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.qb.server.doc/doc/t0006833.html .

1. Log on as a user ID with root authority.
2. If you are not using a CD to install DB2, you can skip to step 5 and run **db2setup** from the directory where you have the DB2 source installation program. If you are installing DB2 using the CD, create a mount point or select an existing mount point. To create a mount point called `/cdrom`, enter the following command: **mkdir /cdrom**.
3. Insert the DB2 CD into the CD drive. Mount the CD file system at the wanted mount point.

   On AIX, you can use the **crfs** command to add an entry to `/etc/filesystems` for the CD mount point. Run the following commands:

```
/usr/sbin/crfs -v cdrfs -p ro -d'cd0' -m'/cdrom' -A'no'
mount /cdrom
```

   The **crfs** command only must be run once for a given mount point, and after that you can use **mount** and **umount** for each CD or DVD you put in the drive, for example, **mount /cdrom** and **umount /cdrom**.
4. Change to the directory where the CD is mounted by entering the following command: **cd /cdrom**.
5. Before installing DB2, run this command from a command prompt window:

```
export JAVA_COMPILER=none
```

   This command is necessary as a result of APAR IC63450.
6. Run the following command from the installation image:

```
./db2_install
```

   Provide any input required.

   Note: If you see the following message:

```
WARNING: A minor error occurred while installing
```

Check the installation log for errors. If you do not see any errors, ignore this warning message.

7. Create the DB2 groups, users, and instances by running the following commands.
   a. For the DB2 groups, run the following commands:
   ```
   mkgroup id=999 db2iadm1
    mkgroup id=998 db2fadm1
    mkgroup id=997 dasadm1
   ```
   b. For the DB2 users, run the following commands:
   ```
   mkuser id=1004 pgrp=db2iadm1 groups=db2iadm1
     home=/home/db2inst1 db2inst1
    mkuser id=1003 pgrp=db2fadm1 groups=db2fadm1
     home=/home/db2fenc1 db2fenc1
    mkuser id=1002 pgrp=dasadm1 groups=dasadm1
     home=/home/dasusr1 dasusr1
   ```
   c. Set the password for `db2inst1`, `db2fenc1`, and `dasusr1` by running the following commands:
   ```
   passwd db2inst1
    passwd db2fenc1
    passwd dasusr1
   ```
   d. Setting the new passwords might require that you log in once to change the initial password.
   ```
   login db2inst1
   ```
   Change the initial password and replace with a new password. Repeat this step for `db2fenc1` and `dasusr1` users.
   e. Create the `db2inst1` instance by running the following command:
   ```
   /opt/IBM/db2/V9.7/instance/db2icrt -a server -u db2fenc1 db2inst1
   ```
   f. Create the `dasusr1` instance by running the following command:
   ```
   /opt/IBM/db2/V9.7/instance/dascrt dasusr1
   ```
8. Assign a communication port for the DB2 instance in the `/etc/services` file.
   a. Select a name that is related to the instance name.
   b. Select a free port, for example, 50000. Add the following line to the `/etc/services` file:
   ```
   db2inst1c 50000/tcp
   ```
9. Configure the communication port in the database manager configuration for the instance:
   ```
   db2 update dbm cfg using SVCENAME db2inst1c
   ```
10. Enable the communication protocol TCP/IP for the instance:
    ```
    db2set DB2COMM=TCPIP
    ```
11. From the Tivoli Storage Productivity Center download image or DVD package, find the file for the DB2 Enterprise Server Edition, Version 9.7 license key: `db2ese_o.lic`. You must register the license key for DB2. Ensure that you understand and agree to the license conditions for DB2. The license terms are located in the `/db2/license/Windows` directory. To register the license key, enter the **db2licm** command from the command line:
    ```
    db2licm -a x/db2/license/license_filename.lic
    ```
    where *x* represents the DVD drive that contains the DB2 license key file and *license_filename.lic* is the license key file: `db2ese_o.lic`.
12. Restart the database instance:

```
su - db2inst1
db2stop
db2start
```

Note: After you have finished installing DB2, edit the file /etc/group and add **root** to the db2iadm1 group. The db2iadm1 line in /etc/group looks like the following line:

```
db2iadm1:x:102:root
```

Verifying that DB2 is installed correctly

You can verify that DB2 has been installed properly using the command-line processor (CLP) or the First Steps GUI.

The general steps to verify that DB2 has been installed properly is as follows:

1. Create the SAMPLE database.
2. Connect to the SAMPLE database.
3. Run a query against the SAMPLE database.
4. Drop the SAMPLE database.

Verifying DB2 installation using the command-line processor (CLP)

You can verify that DB2 has been successfully installed using the command-line processor (CLP).

To verify that DB2 has been installed successfully, complete the following steps:

1. Log on to the system as a user with SYSADM authority.
2. Start the database manager by entering the **db2start** command.

   Note: For UNIX and Linux, you must source the db2profile before you run the **db2start** command. for information about how to source the profile, go to the Tivoli Storage Productivity Center information center. Search for **Using the command line on UNIX and Linux**.
3. Enter the **db2sampl** command to create the SAMPLE database. This command might take a few minutes to process. There is no completion message. When the command prompt returns, the process is complete. The SAMPLE database is automatically cataloged with the database alias SAMPLE when it is created.
4. Enter the following DB2 commands from a DB2 command window to connect to the SAMPLE database, retrieve a list of all the employees that work in Department 20, and reset the database connection:

   ```
   db2 connect to sample
   db2 "select * from staff where dept = 20"
   db2 connect reset
   ```
5. After you have verified the installation, you can remove the SAMPLE database to free up disk space. Enter the **db2 drop database sample** command to drop the SAMPLE database.

   For more information about verifying DB2 installation, refer to the DB2 documentation for your operating system.

Verifying DB2 installation using First Steps

You can use First Steps to verify that DB2 has been installed successfully.

You must have the Control Center and the First Steps component installed to perform this task. First Steps are part of the getting started component grouping in

the DB2 Setup wizard. It is installed as part of a typical installation or might be selected when performing a Custom installation.

To verify the DB2 installation, complete the following steps:

1. Log on to the system with the user account that you want to use to verify the installation. Ensure that the domain user account you use to create the sample database has SYSADM or SYSCTRL authority.
2. Start First Steps.
   - On UNIX or Linux, type **db2fs**.
   - On Windows, type **db2fs.bat**.
3. Select **Create Sample Databases** in the First Steps launchpad to open the Create Sample Databases window.
4. In the Create Sample Databases window, select the databases you want to create. The DB2 UDB sample database is used to verify the installation.
5. Click **OK**. By default, the SAMPLE database is created on the computer where DB2 is installed. The command might take a few minutes to process. When the SAMPLE database has been created, you receive a completion message. Click **OK**.
6. Once the database is created, open the DB2 Control Center. You can perform administration tasks on different instance and database objects through the Control Center. In the left pane of the Control Center screen, expand the object tree to view the SAMPLE database and SAMPLE database objects. Select the Tables object to view the SAMPLE database tables in the content pane of the Control Center screen.
7. After you have verified the installation, you can remove the SAMPLE database to free up disk space. Enter the **db2 drop database sample** command to drop the SAMPLE database. For more information about verifying DB2 installation, refer to the DB2 documentation for your operating system.

Installing the IBM Tivoli Storage Productivity Center family

This topic describes how to prepare for and install IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication.

Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication provides an installation program that helps guide you through the installation process. You can use typical installation or custom installation to install Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication.

Preparing for installation

This topic provides information about what you must do to prepare for installing IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication.

Before you begin the installation, complete the following steps:

1. Make sure that you have the correct version of DB2 installed on your system. Also make sure that DB2 is running.
2. Ensure that you do not have any port conflicts. See “TCP/IP ports used by the IBM Tivoli Storage Productivity Center family” on page 33.
3. Make sure that your computers are using fully qualified host names. Make sure that you change the HOSTS file. For information about the HOSTS file, see “Changing the HOSTS file” on page 346.

4. Review the IBM Tivoli Storage Productivity Center readme file for Version 4.2.2 at https://www.ibm.com/support/docview.wss?uid=swg21509235. The readme file contains last-minute information that could not be included in the documentation.
5. If you are installing the server on a UNIX computer, you must set up your shell environment to point to the instance where the repository is installed. To source the environment, run the **db2profile** script for the wanted instance. For example, if the DB2 instance is `db2inst1`, perform the following steps:
   a. Source the file by entering `. /home/db2inst1/sqllib/db2profile`.
   b. Start the installation program.
6. Plan for the user IDs and passwords you expect to use. See "User IDs and passwords for installing and using Tivoli Storage Productivity Center" on page 37.
7. If you install Tivoli Storage Productivity Center and select to authenticate users against an LDAP-compliant directory, then make sure that you have all the necessary LDAP information:
   - Fully qualified host name of the server running the LDAP-compliant directory
   - Port on which the LDAP-compliant directory is listening
   - The Bind Distinguished Name and password for connecting to the LDAP-compliant directory
   - The Relative Distinguished Names for the Tivoli Storage Productivity Center users and groups in the LDAP-compliant directory
   - The attributes used to "name" the Tivoli Storage Productivity Center users and groups in the LDAP-compliant directory
   - The LDAP TPC Administrator user name and password; the user name and password must be an existing entry in the LDAP-compliant directory in the "branch" specified by the "Relative Distinguished Name for usernames" value
   - The LDAP TPC Administrator group; this group must be an existing entry in the LDAP-compliant directory in the "branch" specified by the "Relative Distinguished Name for groups" value; the LDAP TPC Administrator user must be a member of this group

Note:

- On Windows, you cannot run the Tivoli Storage Productivity Center installation program from a network share (for example, Universal Naming Convention format). You must map your network share to your local drive and run the installation program from that drive.
- For custom installation: in UNIX, the Tivoli Storage Productivity Center installation program validates the WebSphere user ID and then gives an error if that user ID does not exist. In Windows, the Tivoli Storage Productivity Center installation program creates the WebSphere user ID if that user ID does not exist.
- If you are using an IPv6 system, see "Planning for Internet Protocol Version 6" on page 51.
- When canceling the installation wizard, click **Cancel** to clean up any residue after you cancel the installation. Clicking **X** in the window does not clean up the canceled installation. For example, when you try to reinstall the program, you might see the message that the TSRMsvr1 user ID exists after a canceled installation.
- If you get an error message with a return code when installing Storage Resource agents, see the Tivoli Storage Productivity Center Information Center. Search for **Return codes used by Storage Resource agent**.

Starting the installation program

This topic describes how to start the IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication installation program.

Note: Note: If you are installing Tivoli Storage Productivity Center on Windows, do not login with a domain user ID for installation. You must use a login with a local administrative user ID to install Tivoli Storage Productivity Center.

If you are installing Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication using the electronic images, there are two installation images:

disk 1 Disk 1 has four parts:

- Disk 1, Part 1 -- Base IBM Tivoli Storage Productivity Center installation program
- Disk 1, Part 2 -- Base IBM Tivoli Integrated Portal and IBM Tivoli Storage Productivity Center for Replication installation programs
- Disk 1, Part 3 -- IBM Tivoli Integrated Portal PTF installation program
- Disk 1, Part 4 -- IBM Tivoli Storage Productivity Center Monitoring Agent installation program for IBM Tivoli Monitoring. This is an optional program and only required if you are using Tivoli Monitoring.

The **disk 1** images part 1, part 2, part 3, and part 4 (if you are using Tivoli Monitoring) must be downloaded to the same directory.

Storage Resource Agent
Contains the files to perform local Storage Resource agent installations. Download the file for the platform you want the agent to reside on.

The **Storage Resource Agent** location and operating system file name is the same as the **disk 1** image.

For Storage Resource agents, the image is located in the following location:
DVD/data/sra/*operating_system_name*

The operating systems supported are listed in the following table.

Table 34. Operating system for Storage Resource agents

Operating system	Operating system name
AIX	aix_power
HP-UX	hp-ux_itanium
Linux x86	linux_ix86
Linux for Power Systems Servers	linux_power
Linux s390	linux_s390
Oracle Solaris	solaris_spark
Windows	windows

Note:

- When you install Tivoli Storage Productivity Center on AIX or Linux, you must source the db2profile before running the installation program. You must also source the db2profile when you install the database and license.
- If you are using Tivoli Storage Productivity Center electronic installation images that are downloaded to a Windows directory, ensure that the directory does not

have spaces or unusual special characters in the name. This action causes the Tivoli Storage Productivity Center installation to fail.

For example, if you have a directory name:

```
C:\tpc 41 standard edition\disk1
```

The Tivoli Storage Productivity Center installation fails.

Change the directory name to one that does not have spaces, for example:

```
C:\tpc41se\disk1
```

This action ensures a successful Tivoli Storage Productivity Center installation.

- If you install the Storage Resource agent using the disk 1 image, you must uninstall the agent using the uninstallation wizard (InstallShield). If you uninstall the agent using the Tivoli Storage Productivity Center user interface, InstallShield does not recognize the uninstallation. Therefore, you must uninstall the agent that is installed from the disk 1 image using InstallShield.
- You cannot install the Storage Resource agent from the disk 1 image on machines that do not have a Data server installed. If you are installing only a Storage Resource agent, use the Tivoli Storage Productivity Center user interface. You can install the Storage Resource agent if you are also installing the Data server using the disk 1 image. You can, however, install the Storage Resource agent on any machine using the Storage Resource Agent installation image or DVD.
- If you are deploying Storage Resource agent on UNIX or Linux systems, you must use root as the user ID.

To start the Tivoli Storage Productivity Center installation program, complete the following steps:

1. For Windows DVD (**disk 1**): Open Windows Explorer and go to the Tivoli Storage Productivity Center DVD drive or directory. Go to:

   ```
   cd DVD_drive\disk1
   ```

   Double-click **setup.exe**.
2. For Windows electronic image: Download the electronic image into a directory, extract the file, and run **setup.exe** from the source installation directory.
3. For Windows DVD (**Storage Resource Agent**): If Windows Autorun is enabled, the installation program starts automatically. If it does not, open Windows Explorer and go to the Tivoli Storage Productivity Center CD-ROM drive. Double-click **setup.exe**.
4. For Linux DVD (**disk 1**):
   a. Create a mount point called `/cdrom`:

      ```
      mkdir /cdrom
      ```

   b. Insert the DVD into the DVD drive. Mount the DVD file system at the wanted mount point. For example, to mount a DVD into a DVD drive known as `/dev/cd0` at mount point `/cdrom`, enter the following command:

      ```
      mount -o ro /dev/cdrom /cdrom
      ```

   c. Change to the directory where the DVD is mounted by entering the following command:

      ```
      cd /cdrom
      ```

   d. Set up your shell environment to point to the instance where the database repository is installed. To source the environment, run the **db2profile** script for the wanted instance. For example, if the DB2 instance is **db2inst1**, perform the following steps:

1) Source the file by typing:

```
. /home/db2inst1/sqllib/db2profile
```

Note: Please note that there is a space between **.** and **/home**.

2) Start the installation program from the **disk 1** directory (for example, *TPC_source_install_dir***/disk1**). Run **./setup.sh**.

Note: Sometimes **setup.sh** exits prematurely because the media (DVD) is not ready yet. In this case, run **./setup.sh** again until the installation wizard is displayed.

5. For AIX DVD (**disk 1**):
   a. Insert the DVD into the DVD drive. Mount the DVD file system at the wanted mount point.

      On AIX, you can use the **crfs** command to add an entry to `/etc/filesystems` for the CD-ROM mount point. Run the following commands:

      ```
      /usr/sbin/crfs -v cdrfs -p ro -d'cd0' -m'/cdrom' -A'no'
      mount /cdrom
      ```

      The **crfs** command only has to be run once for a given mount point, and after that you can use **mount** and **umount** for each CD or DVD you put in the drive, for example, **mount /cdrom** and **umount /cdrom**.

   b. Change to the directory where the DVD is mounted by entering the following command: **cd /cdrom**.

   c. Set up your shell environment to point to the instance where the database repository is installed. To source the environment, run the **db2profile** script for the wanted instance. For example, if the DB2 instance is **db2inst1**, perform the following steps:

      1) Source the file by entering:

      ```
      . /home/db2inst1/sqllib/db2profile
      ```

      Note: Please note that there is a space between **.** and **/home**.

      2) Start the installation program from the **disk 1** directory (for example, *TPC_source_install_dir***/disk1**). Run **./setup.sh**.

      Note: Sometimes **setup.sh** exits prematurely because the media (DVD) is not ready yet. In this case, invoke **./setup.sh** again until the installation wizard is displayed.

6. For Linux or AIX electronic image (**disk 1**):
   a. Create a directory:

      ```
      mkdir /tpcinst
      ```

   b. Change to the directory:

      ```
      cd tpcinst
      ```

   c. FTP the Tivoli Storage Productivity Center installation tar file into a directory:

      ```
      ftp filename.tar
      ```

   d. Run the following command:

      ```
      tar -xvf filename.tar
      ```

   e. Set up your shell environment to point to the instance where the database repository is installed. Source the environment by running the **db2profile** script for the wanted instance.

      For example, if the DB2 instance is **db2inst1**, perform the following steps:

1) Source the file by entering:

```
. /home/db2inst1/sqllib/db2profile
```

Note: Please note that there is a space between **.** and **/home**.

2) Start the installation program. Run **./setup.sh**.

7. For Linux and AIX DVD (**Storage Resource Agent**):
   a. Create a mount point called `/cdrom`:

   ```
   mkdir /cdrom
   ```

   b. Insert the DVD into the CD-ROM drive. Mount the CD-ROM file system at the wanted mount point. For example, to mount a DVD, enter the following command:

   ```
   mount -o ro /dev/cdrom /cdrom
   ```

   c. Change to the directory where the CD-ROM is mounted by entering the following command:

   ```
   cd /cdrom
   ```

   d. Set up your shell environment to point to the instance where the database repository is installed. To source the environment, run the **db2profile** script for the wanted instance.
   e. For example, if the DB2 instance is **db2inst1**, perform the following steps:
      1) Source the file by entering:

      ```
      . /home/db2inst1/sqllib/db2profile
      ```

      Note: Please note that there is a space between **.** and **/home**.

      2) Start the installation program. Run **./setup.sh**.

Using typical installation

This section describes how to install Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication on one computer using typical installation.

The approximate time to install Tivoli Storage Productivity Center and Tivoli Integrated Portal is about 60 minutes. The approximate time to install Tivoli Storage Productivity Center for Replication is about 20 minutes. The time to install Tivoli Integrated Portal on the AIX or Linux operating system is about 15 minutes less than on Windows.

These are things to note when using typical installation:

- If you install the Tivoli Storage Productivity Center servers using typical installation, all user IDs and passwords used by Tivoli Storage Productivity Center for installation are set to share the common user ID and password. The user ID and password entered are entered in the following installation panel: **User ID and password, and server and agent information**. This user ID must have operating system administrator and database administrator authority. If you want to set up different user IDs and passwords, use custom installation.

  The common user ID is used for:

  – DB2 user ID
  – DB2 administrator user ID
  – WebSphere administrator user ID
  – NAS Filer Logon user ID
  – Storage Resource agent user ID

The common user password is used for:

- DB2 user password
- DB2 administrator password
- WebSphere administrator password
- NAS Filer Logon user password
- Storage Resource agent user password
- Host authentication user password

To install Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication using typical installation, follow these steps:

1. Start the installation program. For information about starting the installation program, see "Starting the installation program" on page 169.
2. The Select a language panel is displayed. Select a language from the list box and click **OK**. This language is the language that is used for installing this product.
3. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.

   Note: 8 GB of RAM is required. If you have 4-8 GB of RAM, you can still install Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication. However, you get a warning message. To avoid system load problems, run only Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication because of system load. For information about how to disable Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication after installation, see "Disabling unlicensed or unused Tivoli Storage Productivity Center components" on page 368.
4. The select the type of installation panel is displayed.

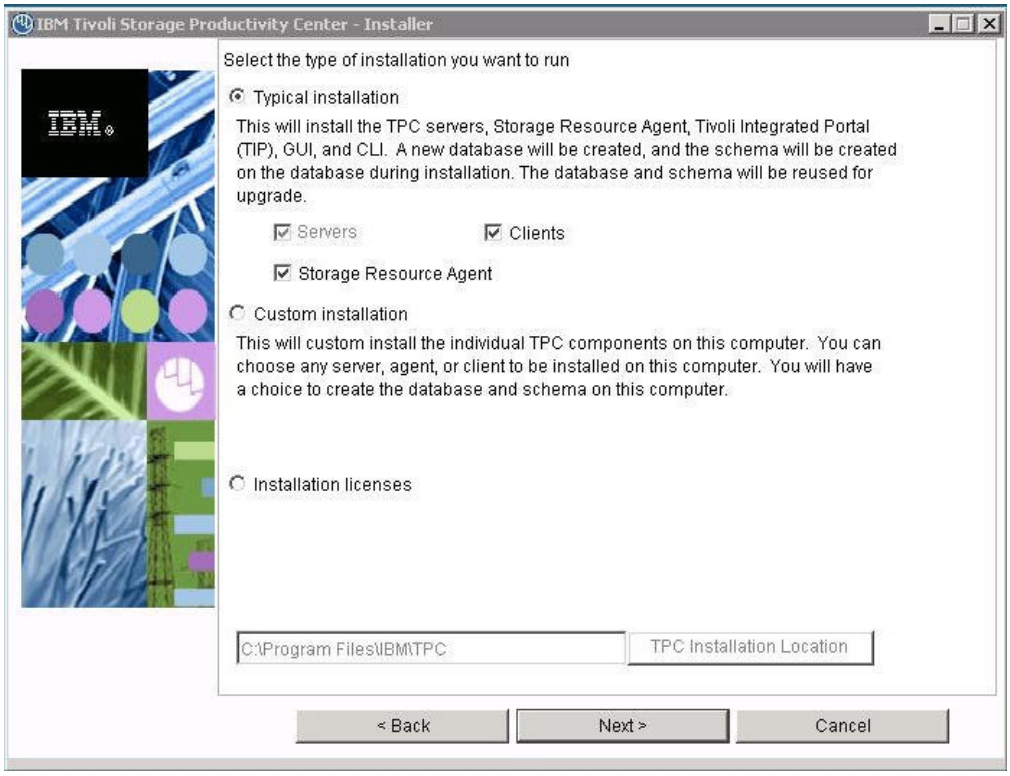


Figure 1. Select the type of installation you want to run panel

Enter the following information and click **Next**.

Typical installation
Allows you to install all the components in a group on the local computer by selecting **Servers**, **Storage Resource Agent**, and **Clients**. This selection requires a minimum amount of input from the user who installs Tivoli Storage Productivity Center.

Typical installation also installs IBM Tivoli Integrated Portal, Tivoli Storage Productivity Center for Replication, the graphical user interface, and CLI. If you want to use an existing Tivoli Integrated Portal, use custom installation.

Because this installation is a typical installation, select **Typical installation**. Under typical installation, you have these options:

Servers
Installs the database schema, the Data Server, and the Device server.

Storage Resource Agent
Installs the Storage Resource agent on the local host.

Clients
Installs the Tivoli Storage Productivity Center GUI and the command-line interface (CLI).

Custom installation
Allows you to install one or more components separately.

Installation licenses
Installs the Tivoli Storage Productivity Center license. The Tivoli Storage Productivity Center license is on the DVD. You only need to

run this option when you upgrade a license to a Tivoli Storage Productivity Center package that has already been installed on your system. For example, if you have installed the Tivoli Storage Productivity Center Basic Edition and you want to upgrade to Tivoli Storage Productivity Center Standard Edition, run the installation program and select **Installation licenses**. This option allows you to upgrade the license from the product DVD.

For information about upgrading the license, see "Adding an installation license" on page 227.

For this scenario, select **Typical installation**, **Servers**, **Storage Resource Agent**, and **Clients**. All the components are installed on the local computer. This option also installs the database schema on the local computer. When you install the database schema, this option creates the database (if the database does not exist), the schema, and the table spaces. At the bottom of the panel, you can enter the directory where you want to install the components. The default location for Windows is `C:\Program Files\IBM\TPC`. The default location for UNIX or Linux is *`usr or opt`*`/IBM/TPC`. Click **Next**.

Note: If the directory path appears as `C:\Program Files (x86)\IBM\TPC` for Windows, remove the `(x86)` from the directory path to avoid installation issues. The directory path should appear as `C:\Program Files\IBM\TPC`. For ease of maintenance, use the default directory path.

5. The User ID and password panel is displayed. Enter the user ID and password that has administrator and database administrator authority.

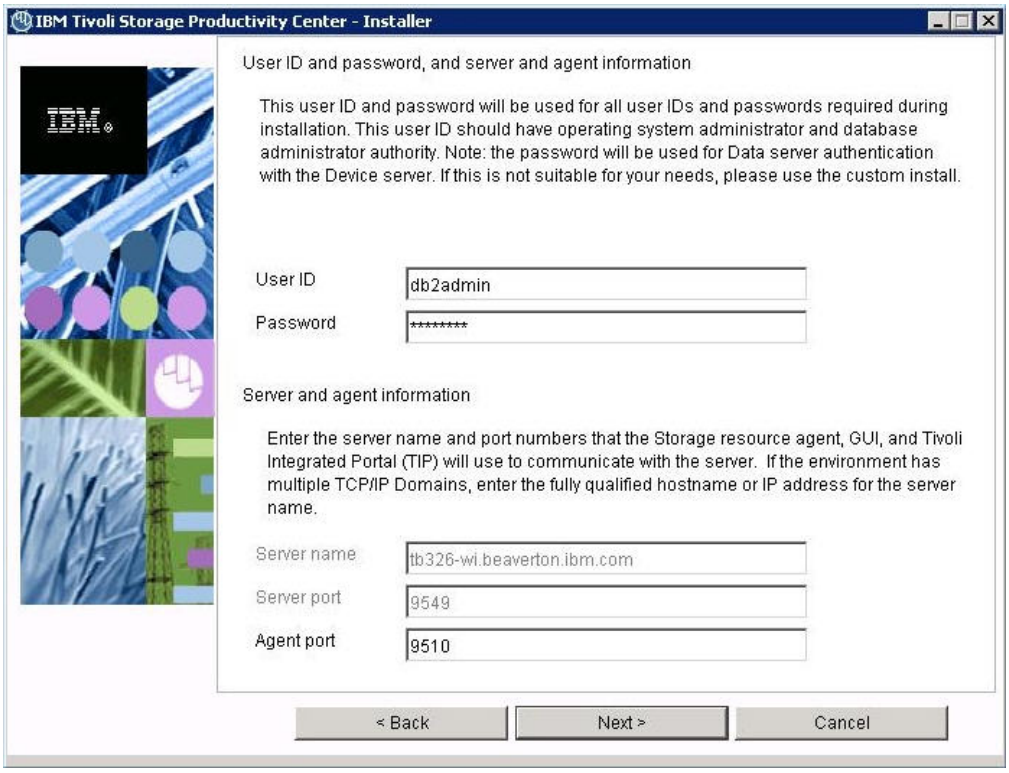


Figure 2. User ID and password panel

Field descriptions:

User ID

Enter the user ID and password that has administrator and database administrator authority. This user ID and password are used for the following items:

- Database administrator user ID and password (for the Data Server or the 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).
- Common agent service logon 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) only when you select "OS" as the authentication mechanism; if you select "LDAP" as the authentication mechanism, then the LDAP TPC Administrator user name and password values you enter during installation are used for the WebSphere administration user ID and password.

For information about these user IDs and passwords, see "Work sheet for user IDs and passwords" on page 40.

Password

Password for the user ID.

Server name

The Tivoli Storage Productivity Center server name. If your environment includes multiple Tivoli Storage Productivity Center servers and IP domains, specify a fully qualified host name. On Windows, the server name might appear as a short host name.

Server port

Port number for the Tivoli Storage Productivity Center server.

The default server port is 9549. This port is assigned to the Data server and port 9550 is assigned to the Device server. If you specify a different server port, for example, 9569, then the the Device server is assigned the next higher port number (9570)

Agent port

Port number for the agent. The default port is 9510. If the default port is not used for the agent port, then the Tivoli Storage Productivity Center agent needs to be installed using custom installation.

Enter the information in the fields or accept the defaults. Click **Next**.

6. The Tivoli Integrated Portal panel is displayed.

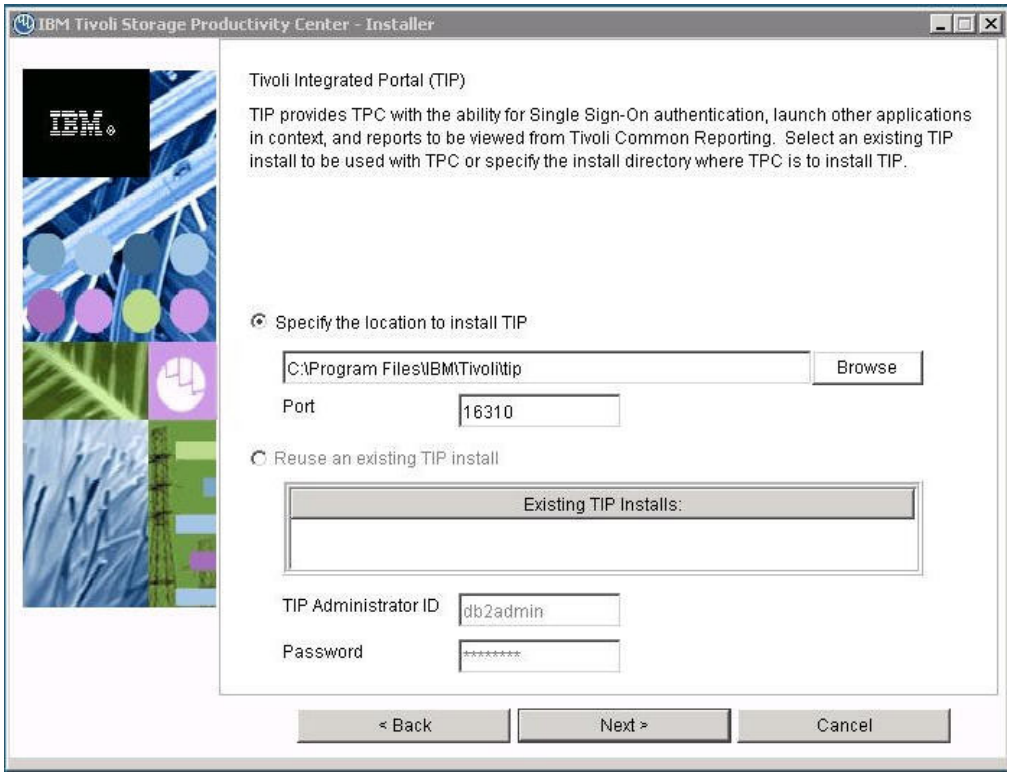


Figure 3. Tivoli Integrated Portal panel

Options and fields:

Specify the location to install TIP
Accept the default or enter a location to install Tivoli Integrated Portal.

Port Enter the port number for Tivoli Integrated Portal or accept the default (16310). Tivoli Integrated Portal uses 10 port numbers. For example, if the default port 16310 is accepted, Tivoli Integrated Portal uses the following ports: 16310, 16311, 16312, 16313, 16315, 16316, 16318, 16320, 16322, and 16323. If you enter a port number other than the default number, ensure that you have the following ports available for Tivoli Integrated Portal:

- *base port*
- *base port* +1
- *base port* +2
- *base port* +3
- *base port* +5
- *base port* +6
- *base port* +8
- *base port* +10
- *base port* +12
- *base port* +13

where *base port* is the port value entered in the **Port** field.

Reuse an existing install TIP install
If you select this option, you see a list of existing Tivoli Integrated Portal installations. Select a Tivoli Integrated Portal installation.

TIP Administrator ID
Enter theTivoli Integrated Portal administrator ID.

Password
Enter the password for the Tivoli Integrated Portal administrator ID.

Click **Next**.

7. The authentication selection panel is displayed.

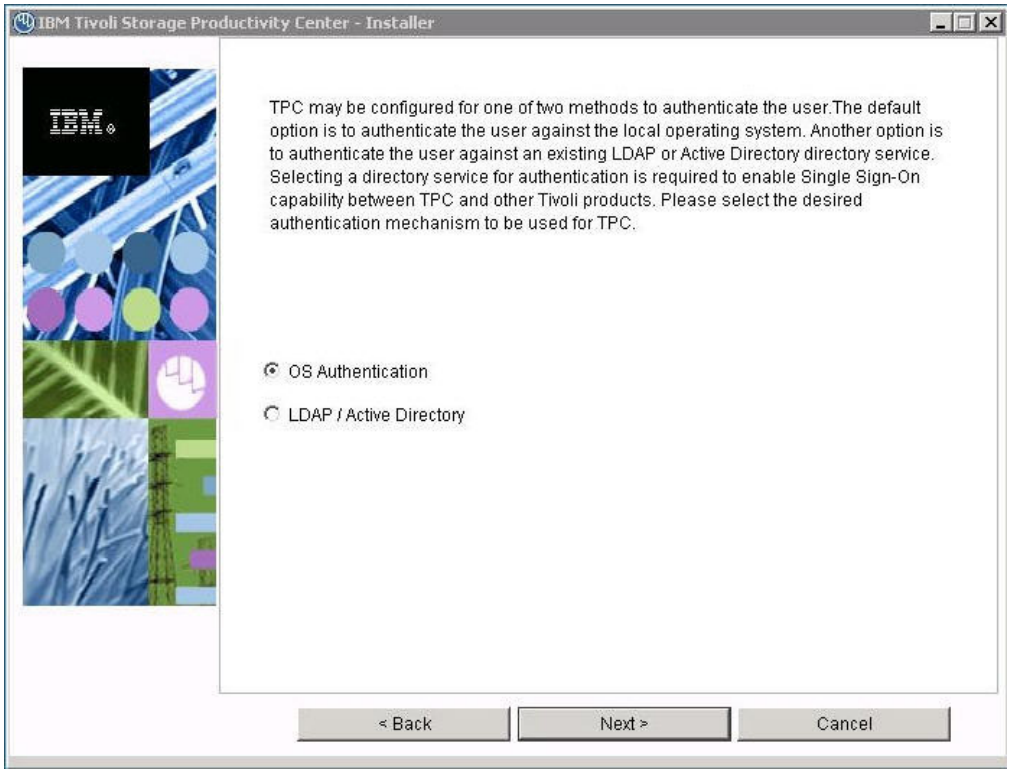


Figure 4. Authentication selection panel

Options:

OS Authentication
Uses the operating system for user authentication. If you select this option, go to step 11 on page 182.

LDAP/Active Directory
If you select LDAP or Microsoft Active Directory for authentication, you must have an LDAP or Active Directory already installed.

If you select this option, go to step 8. Click **Next**.

8. The Lightweight Directory Access Protocol (LDAP server information) panel is displayed. Enter this information and click **Next**.

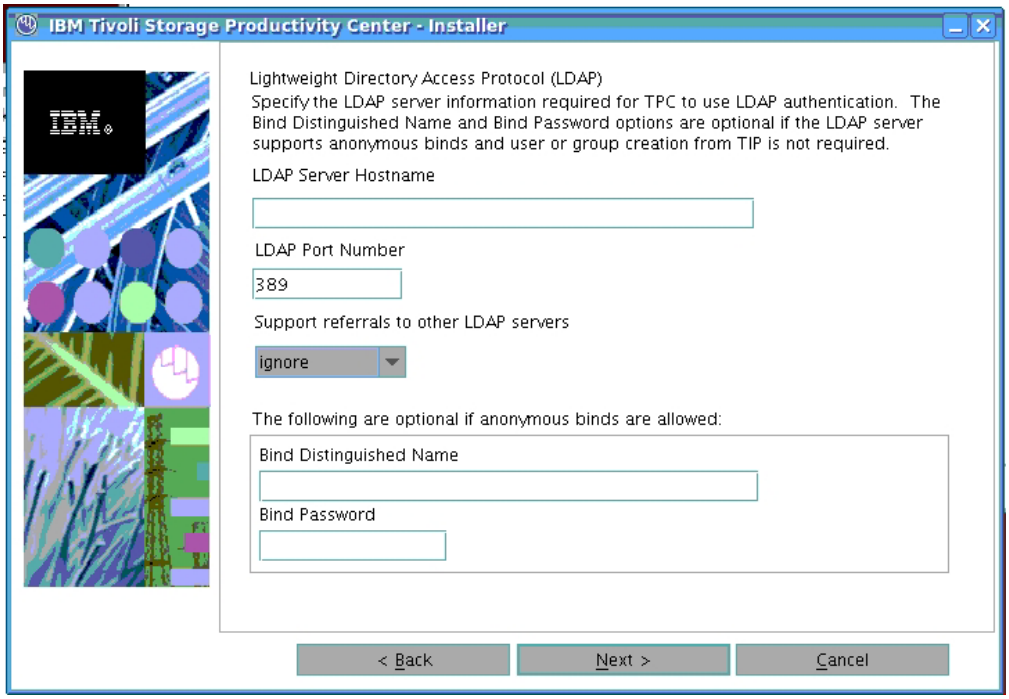


Figure 5. LDAP server information panel

Enter this information and click **Next**.

LDAP Server Hostname
The fully qualified host and domain name of the machine where your LDAP-based directory is running.

LDAP Port Number
The port number on your LDAP server where the LDAP process is listening for communications.

Support referrals to other LDAP servers
The options are **ignore** or **follow** an LDAP referral. The default is **ignore**.

Referral is the process by which an LDAP server, instead of returning a result to a client application, returns a reference (a referral) to another LDAP server which might contain further information related to the client application's request.

Bind Distinguished Name
The distinguished name used to bind to the LDAP-based directory when performing a user or group search. This name is optional because some LDAP-based directories allow anonymous binds and others require you to bind with a specific name and password.

Note: Depending on how the LDAP server is set up, the Bind DN and password might or might not be optional. If the LDAP server allows for anonymous binds, then the Bind DN and password are optional.

If you want the ability to create or modify LDAP users and groups from the Tivoli Integrated Portal admin panel, then the Bind DN and password are required.

Bind Password
The password associated with the Bind Distinguished Name.

9. The Lightweight Directory Access Protocol (LDAP user and group information) panel is displayed.

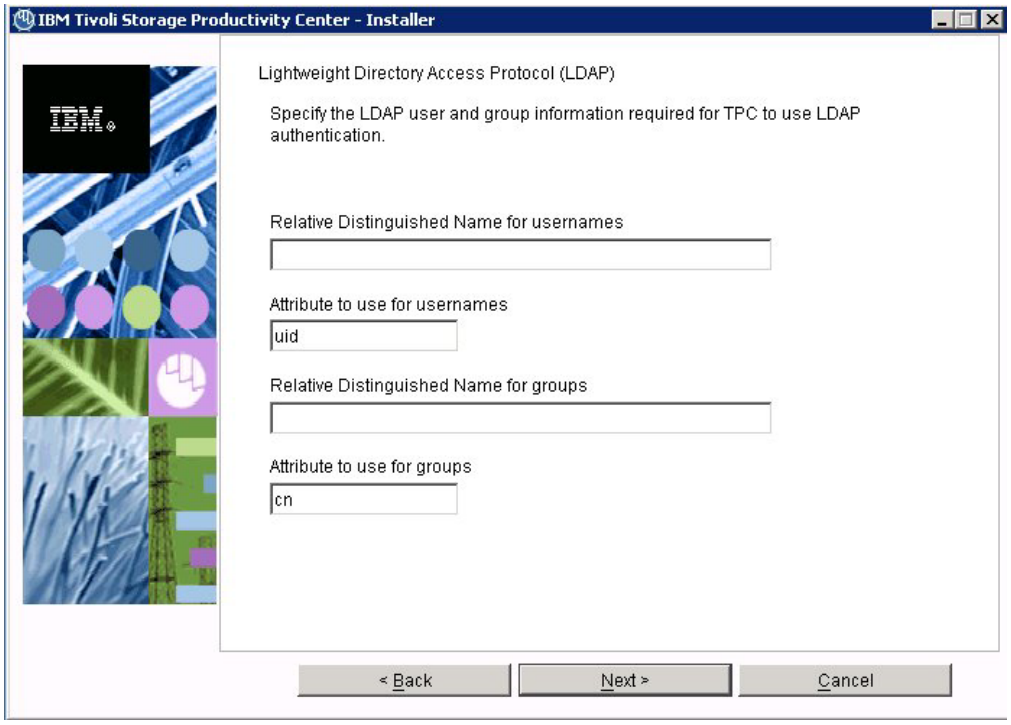


Figure 6. LDAP user and group information panel

Enter this information and click **Next**.

Relative Distinguished Name for usernames
Tells Tivoli Storage Productivity Center where to start a search in the directory when performing user authentication.

Attribute to use for usernames
Tells Tivoli Storage Productivity Center which attribute in a directory entry for the user which contains the name for the user for authentication.

Relative Distinguished Name for groups
Tells Tivoli Storage Productivity Center where in the directory to start a search when performing a group search for authorization.

Attribute to use for groups
Tells Tivoli Storage Productivity Center which attribute in a directory entry for the group which contains the name for the group for authorization.

10. The Lightweight Directory Access Protocol (LDAP TPC Administrator user and group information) panel is displayed.

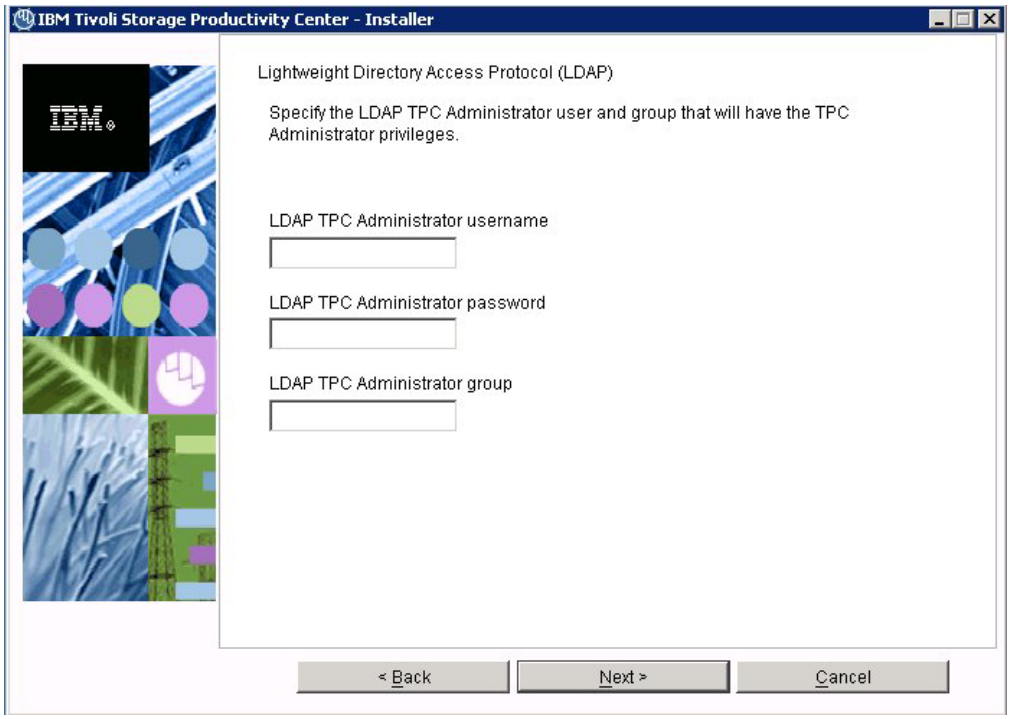


Figure 7. LDAP TPC Administrator user and group information panel

Enter this information and click **Next**.

LDAP TPC Administrator username
The user ID that Tivoli Storage Productivity Center configures as the Tivoli Storage Productivity Center superuser during installation. Ensure that this user name exists in the directory before you start the Tivoli Storage Productivity Center installation. The user name must be a member of your designated LDAP Tivoli Storage Productivity Center Administrator group in the directory.

Note: If you install Tivoli Storage Productivity Center on a Windows system and have Tivoli Storage Productivity Center authenticate users against an LDAP-based repository, then the LDAP Tivoli Storage Productivity Center Administrator user name must **not** contain a space or blank character in it. See the WebSphere Application Server APAR PK77578.

Ensure that you use the correct case when entering the values for the LDAP TPC Administrator user name and LDAP TPC Administrator group. If you use the incorrect case when entering these values, then you cannot log in to Tivoli Storage Productivity Center for Replication with the user name you expected to use.

LDAP TPC Administrator password
Password for the administrator.

LDAP TPC Administrator group
The group that Tivoli Storage Productivity Center maps to the Tivoli Storage Productivity Center superuser role during installation. Ensure that this group exists in the directory before you start the Tivoli

Storage Productivity Center installation. This group must have the designated LDAP Tivoli Storage Productivity Center Administrator as a member.

Note:

Ensure that you use the correct case when entering the values for the LDAP TPC Administrator user name and LDAP TPC Administrator group. If you use the incorrect case when entering these values, then you cannot log in to Tivoli Storage Productivity Center for Replication with the user name you expected to use.

11. The Summary Information panel is displayed. Review the information. Click **Install**.
12. You see the installing panel. When Tivoli Integrated Portal has finished installing, Tivoli Storage Productivity Center starts the Tivoli Storage Productivity Center for Replication installation.
13. Tivoli Storage Productivity Center starts the Tivoli Storage Productivity Center for Replication installation program.
    a. The Welcome panel is displayed. Click **Next**.
    b. The System prerequisites check panel is displayed. The installation wizard checks whether the prerequisites are installed, then confirms whether your operating system is supported and is at the appropriate fix pack or update level. Click **Next**.

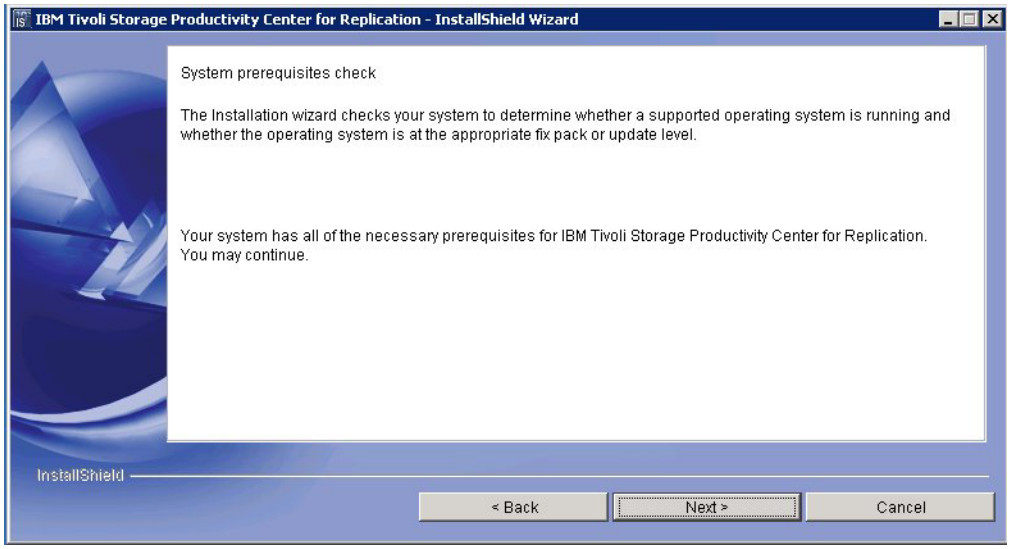


Figure 8. System prerequisites check panel

    c. The License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.
    d. The Directory Name panel is displayed. Accept the default installation directory by clicking **Next**, or specify a different installation directory and click **Next**.

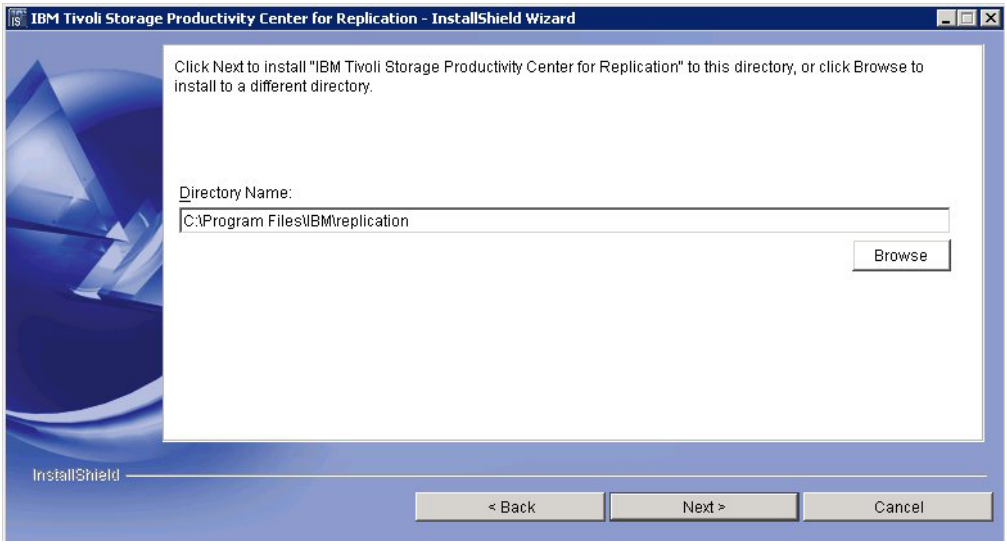


Figure 9. Directory Name panel

e. The Tivoli Storage Productivity Center for Replication Administrator ID and Password panel is displayed. Enter the administrator user ID and password. This ID is typically the operating system administrator user ID. If you use a different ID, create it beforehand and ensure that it has administrator rights.

   Note: There is a limitation on the number of characters for the user name and password on AIX. At the time this document was written, AIX did not properly validate passwords that are longer than eight characters.

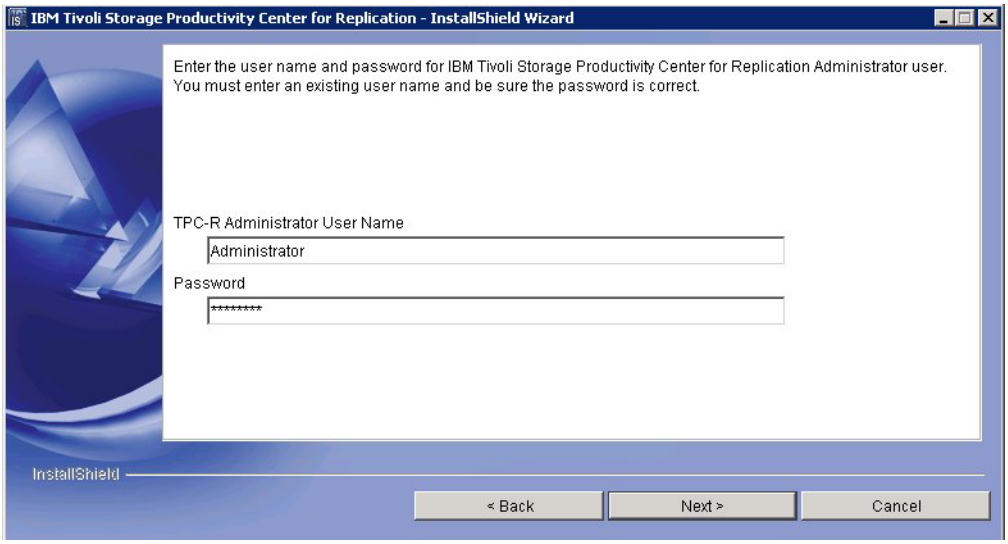


Figure 10. TPC-R Administrator ID and Password panel

f. The Default ports panel is displayed. Accept the defaults. Click **Next**.

   The installation program checks for port conflicts on 5110 and 5120. If there is a conflict, a new port is assigned by increasing the port number, and then retesting for a port conflict until an available port is found. The default port numbers in the fields are set to the available ports when the panel is displayed.

If you assign a port for the Default Client Port or Default Standby Server Port and there is a port conflict, an error message is displayed to let you know of the conflict.

The client port must be set to the same number on both management servers (the active and the standby servers). The standby server port must also be set to the same number on both management servers (the active and the standby servers).

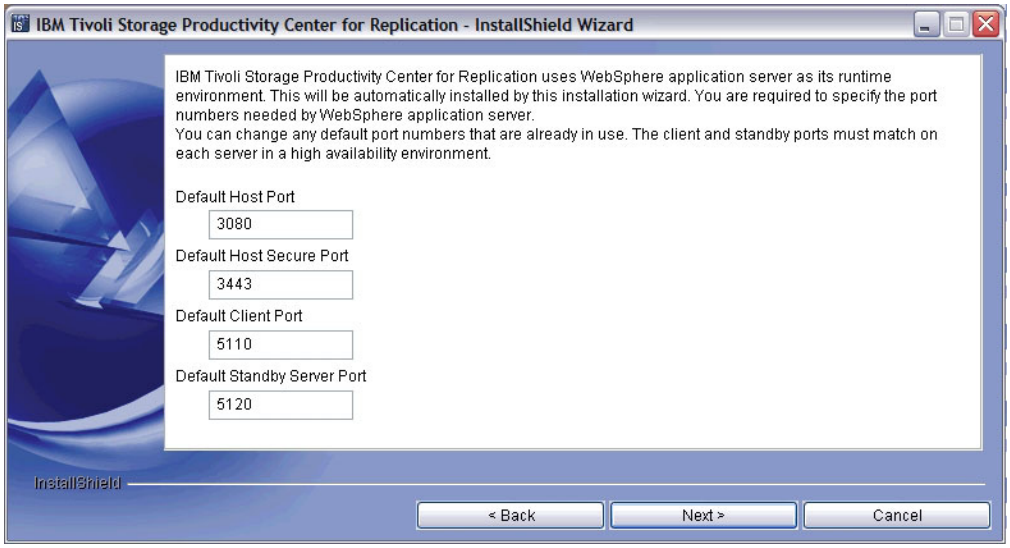


Figure 11. Default ports panel

g. The settings panel is displayed. Review the settings and change them if needed by clicking **Back**. Otherwise, click **Install** to begin the installation.

h. You see the status panel. Wait for installation to complete.

i. You see the summary information panel. Review the information and click **Finish**.

j. Tivoli Storage Productivity Center for Replication is installed with a basic license. If you want the Two Site or Three Site Business Continuity (BC) license, see "Installing the Two Site or Three Site Business Continuity license."

14. The Successfully installed panel is displayed. Click **Finish**.
15. To access the Tivoli Storage Productivity Center user interface, see "Starting the Tivoli Storage Productivity Center GUI using a web browser" on page 347. To access the Tivoli Storage Productivity Center for Replication user interface, go to the Tivoli Storage Productivity Center information center. Search for **Accessing the IBM Tivoli Storage Productivity Center for Replication GUI**.
16. After you have installed Tivoli Storage Productivity Center, you can use the Configure Devices wizard to configure your system. For information about the Configure Devices wizard, see "Configure Devices" on page 266.

Installing the Two Site or Three Site Business Continuity license

This section describes how to install the IBM Tivoli Storage Productivity Center for Replication Two Site or Three Site Business Continuity (BC) license.

To install the license for IBM Tivoli Storage Productivity Center for Replication, you must run as the root user on Linux or AIX, or as the administrator in Windows. Use the Two Site or Three Site CD for the installation of the license.

To install the IBM Tivoli Storage Productivity Center for Replication Two Site or Three Site Business Continuity license, follow these steps:

1. Begin the installation program by double-clicking the setup file:
   - Windows: *installation_source_directory*\setupwin32.exe
   - Linux: *installation_source_directory*/setuplinux.bin
   - AIX: *installation_source_directory*/setupaix.bin
2. On the Welcome page, click **Next**.
3. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.
4. On the Directory name panel, click **Next**.
5. Review the settings and click **Install**.

Using custom installation

This section describes how to install IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication using custom installation.

Note: If you have already installed a DB2 server and IBM Tivoli Storage Productivity Center database schema on a local system, you cannot use a remote DB2 server for the IBM Tivoli Storage Productivity Center installation. The IBM Tivoli Storage Productivity Center installation program selects the local DB2 server by default.

Installing all the components

This procedure provides information about how to install all the Tivoli Storage Productivity Center components and Tivoli Storage Productivity Center for Replication except for the deployment of remote agents.

Note:

- The DB2 database schema name for Tivoli Storage Productivity Center cannot be longer than eight characters.
- For custom installation, all user IDs and passwords used by Tivoli Storage Productivity Center are customized. This means that some or all Tivoli Storage Productivity Center user IDs and passwords are unique and not shared. The user IDs and passwords are:
  - DB2 user ID and password
  - DB2 administrator user ID and password
  - WebSphere administrator user ID and password
  - NAS Filer Logon user ID and password
  - Storage Resource agent user ID and password
  - Host authentication user password

To install all the Tivoli Storage Productivity Center components and Tivoli Storage Productivity Center for Replication, complete the following steps.

1. Start the Tivoli Storage Productivity Center installation program. For information about starting the installation program, see "Starting the installation program" on page 169.

2. The Select a language panel is displayed. Select a language and click **OK**.
3. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.

   Note: 8 GB of RAM is required. If you have 4-8 GB of RAM, you can still install Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication. However, you get a warning message. To avoid system load problems, run only Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication. For information about how to disable Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication after installation, see "Disabling unlicensed or unused Tivoli Storage Productivity Center components" on page 368.
4. The Select the type of installation you want to run panel is displayed.

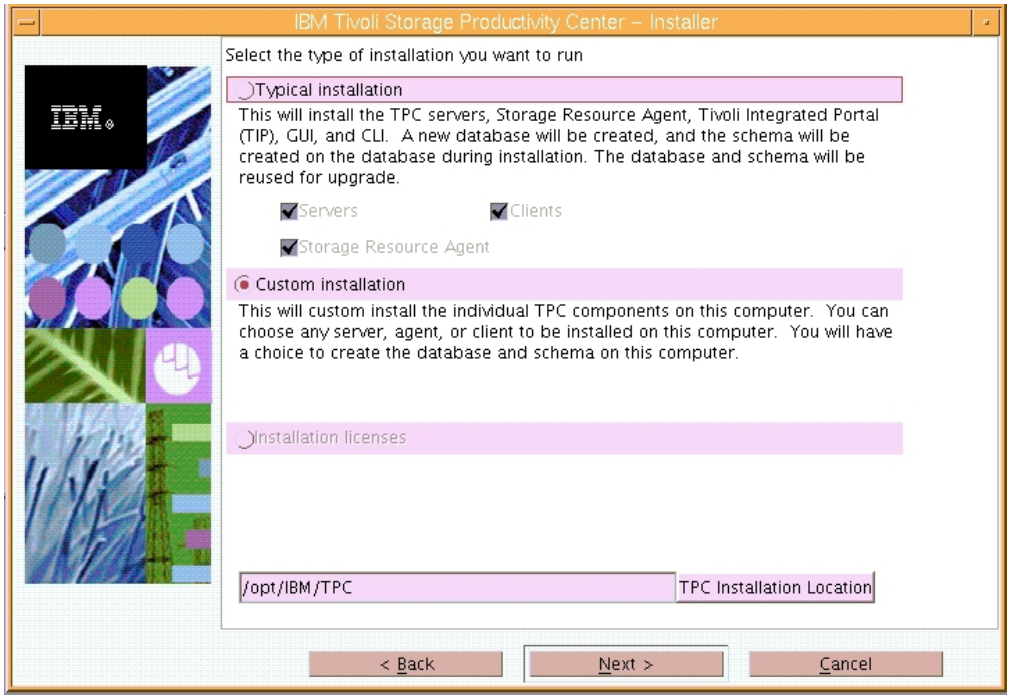


Figure 12. Select the type of installation you want to run panel

You have these options:

Typical installation
For information about typical installation, see "Using typical installation" on page 172.

Custom installation
Allows you to install one or more components separately.

Installation licenses
Installs the Tivoli Storage Productivity Center license. The Tivoli Storage Productivity Center license is on the DVD. You only need to run this option when you upgrade a license to a Tivoli Storage Productivity Center package that has already been installed on your system. For example, if you have installed the Tivoli Storage Productivity Center Basic Edition and you want to upgrade to Tivoli Storage Productivity Center Standard Edition, run the installation

program and select **Installation licenses**. This option allows you to upgrade the license from the product DVD.

For information about upgrading the license, see "Adding an installation license" on page 227.

Select **Custom installation**. Enter the Tivoli Storage Productivity Center Installation Location or accept the default. For Windows, the default location is `C:\Program Files\IBM\TPC`. For UNIX or Linux, the default location is `/opt/IBM/TPC`. Click **Next**.

5. The Select one or more components to installation panel is displayed.

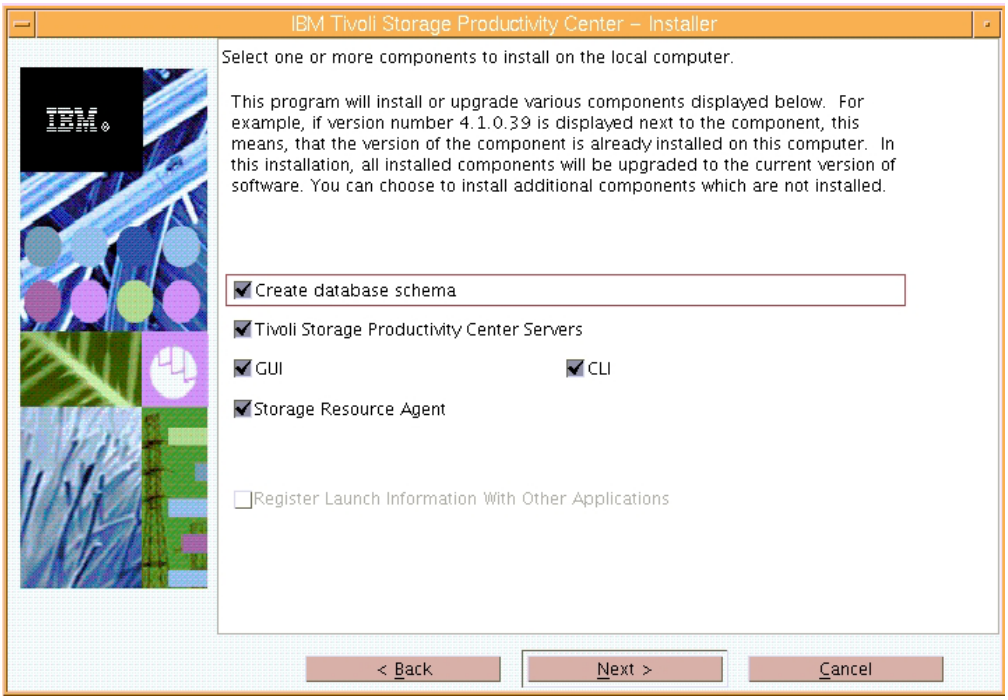


Figure 13. Select one or more components to install panel

Select the following components:

- Create database schema
- Tivoli Storage Productivity Center Servers
- GUI
- CLI
- Storage Resource Agent

Click **Next**.

The **Register Launch Information With Other Applications** option does not install the main Tivoli Storage Productivity Center components (database schema, Data server, Device server, CLI, GUI, Storage Resource agent, Tivoli Integrated Portal, or Tivoli Storage Productivity Center for Replication). Instead, this option only installs and registers the Tivoli Storage Productivity Center Web portal and portlets inside another WebSphere application, the most common being Tivoli Integrated Portal. After you have registered the Tivoli Storage Productivity Center launch information with an existing instance of Tivoli Integrated Portal, then you can launch the Tivoli Storage Productivity Center GUI from within that Tivoli Integrated Portal Console.

Note:

- If you are installing a new Tivoli Storage Productivity Center server using custom installation, install all the components simultaneously rather than individually.

6. The Storage Resource Agent Information panel is displayed.

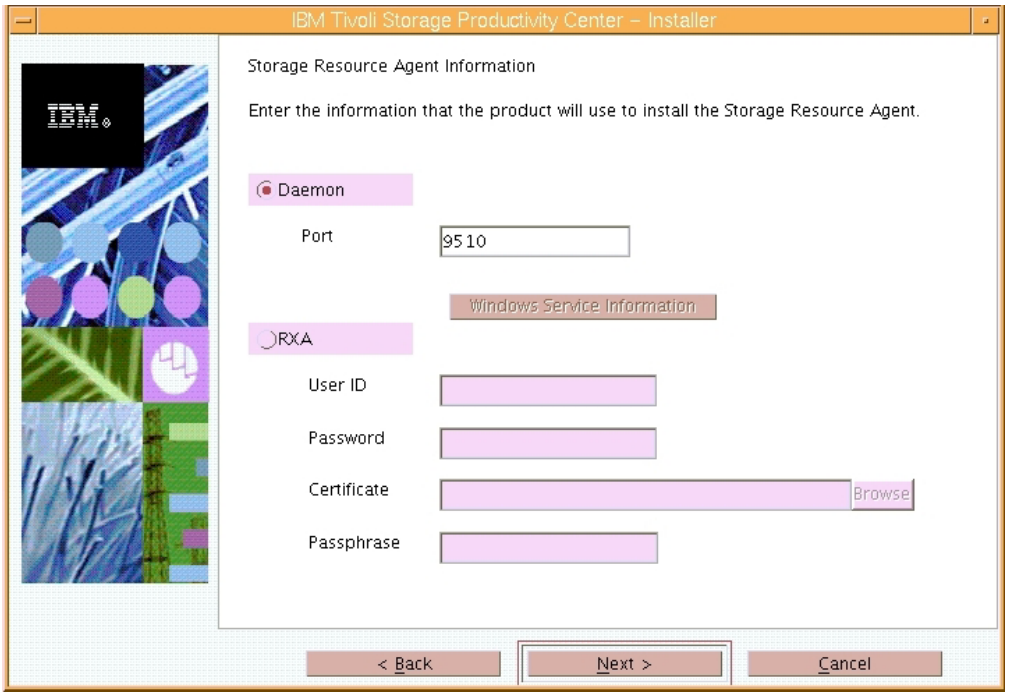


Figure 14. Storage Resource Agent Information panel

Enter the following information:

Daemon

The Storage Resource agent is installed as a daemon service. For more information about the daemon service, see "Storage Resource agent protocol support" on page 50.

Port The Storage Resource agent port. The default port is 9510.

Windows Service Information (Windows only)

Click this button to enter an agent service name, user ID, and password information for the installation program to use when creating a Windows service for the agent.

If you click this button, the Storage Resource Agent service information panel is displayed.

Note: The user ID and password validation is done in two steps. When you first create a new user ID and password, Tivoli Storage Productivity Center checks for a valid user ID and password. When you are back at the Storage Resource Agent service information panel and click **Next**, the password is checked against the local password policy.

For example, this behavior can result in these actions:

- If you specify **Create new local account** and enter an invalid user ID and a valid password, then click **OK** on the

Storage Resource Agent service information panel, you see an error message for the invalid user ID in the pop-up dialog.

- If you enter a valid user ID and invalid password, then click **OK** on the Storage Resource Agent service information panel, then click **Next** on the Storage Resource Agent information panel, you get an error message for the invalid password in the pop-up dialog.

For more information about user IDs and passwords, see "Valid characters for user IDs and passwords" on page 428.

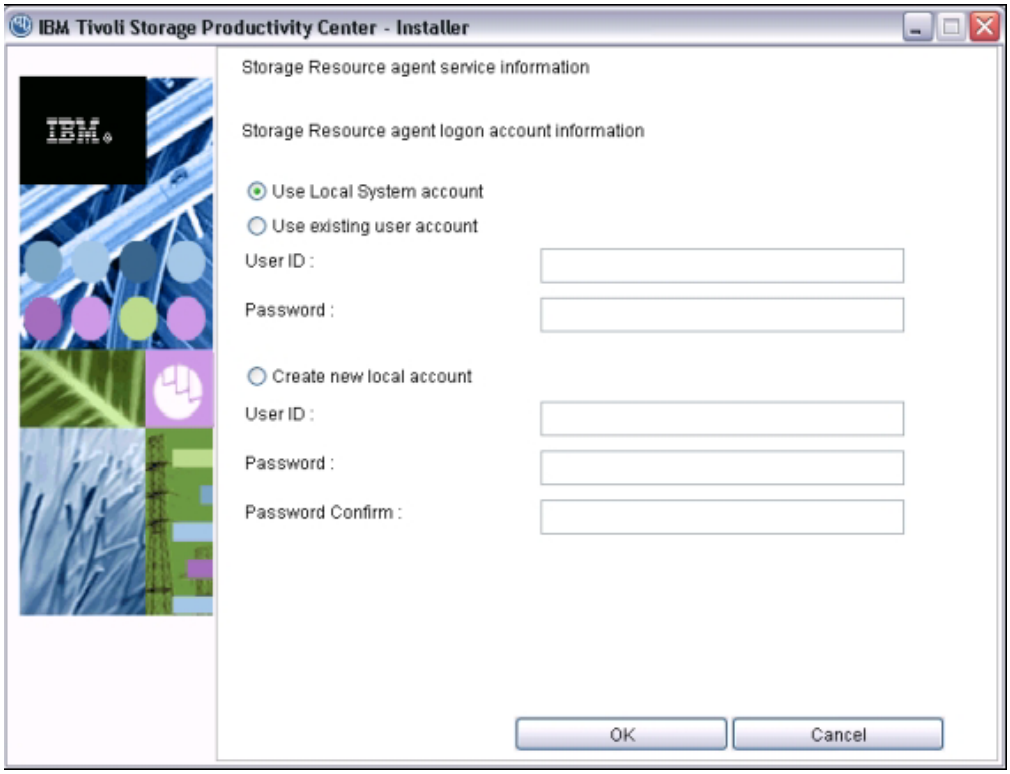


Figure 15. Storage Resource Agent service information panel

Use LocalSystem account
Use the local system account.

Use existing user account
Use an existing service account.

User ID
Enter the user ID for the Windows service.

Password
Enter the password for the Windows service user ID.

Create new local account
Create a new local account.

User ID
Enter the user ID for the Windows service.

Password
Enter the password for the Windows service user ID.

Password confirm
Re-enter the password for the Windows service user ID.

RXA The Storage Resource agent is installed as a non-daemon service. For information about the non-daemon service, see "Storage Resource agent protocol support" on page 50.

Note:

- If you are using the RXA access method to install or deploy the Storage Resource agent on a UNIX or Linux system, you must use the root user ID.
- If you enter a valid root user ID and password, Tivoli Storage Productivity Center does not check the certificate and passphrase.

  If you do not enter a valid password (you enter a valid root user ID only), then Tivoli Storage Productivity Center checks the root user ID, certificate, and passphrase.

User ID
The user ID used to connect to the server.

Password
The password for the user ID to connect to the server.

Certificate
The certificate used to connect to the server. For information about SSH protocol and SSL protocol certificates, see "Creating a certificate for SSH protocol" on page 317 and "Replacing a custom certificate for SSL protocol" on page 322.

Passphrase
The passphrase used to connect to the server.

Click **Next**.

7. The Database administrator information panel is displayed. Enter a database administrator user ID and password. Click **Next**.

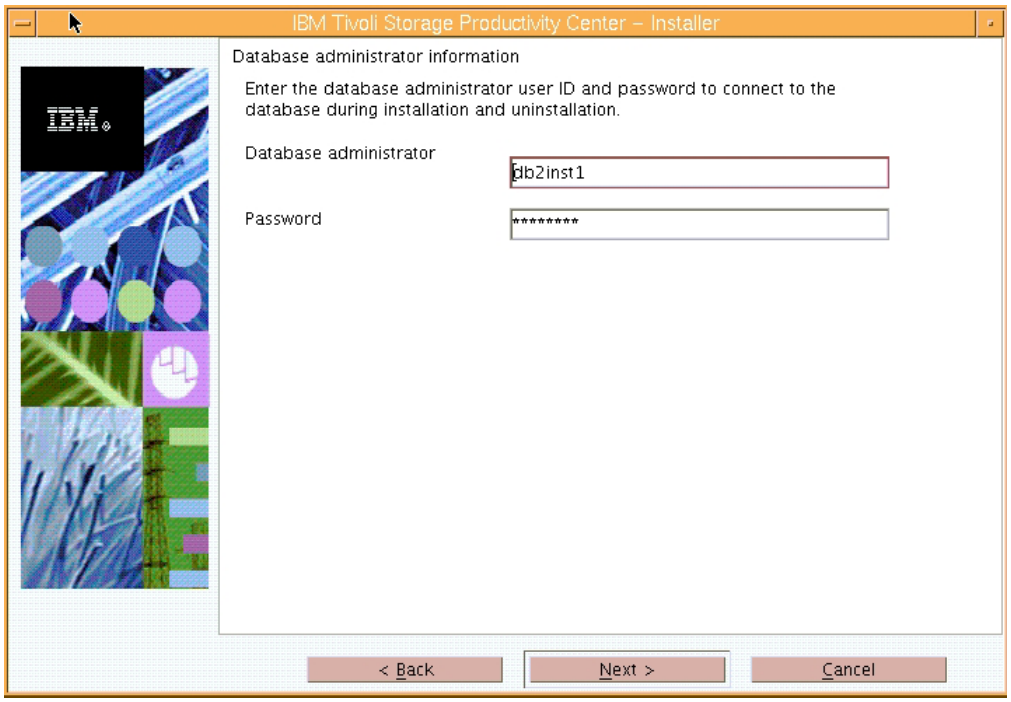


Figure 16. Database administrator Information panel

8. The new database schema information panel is displayed.

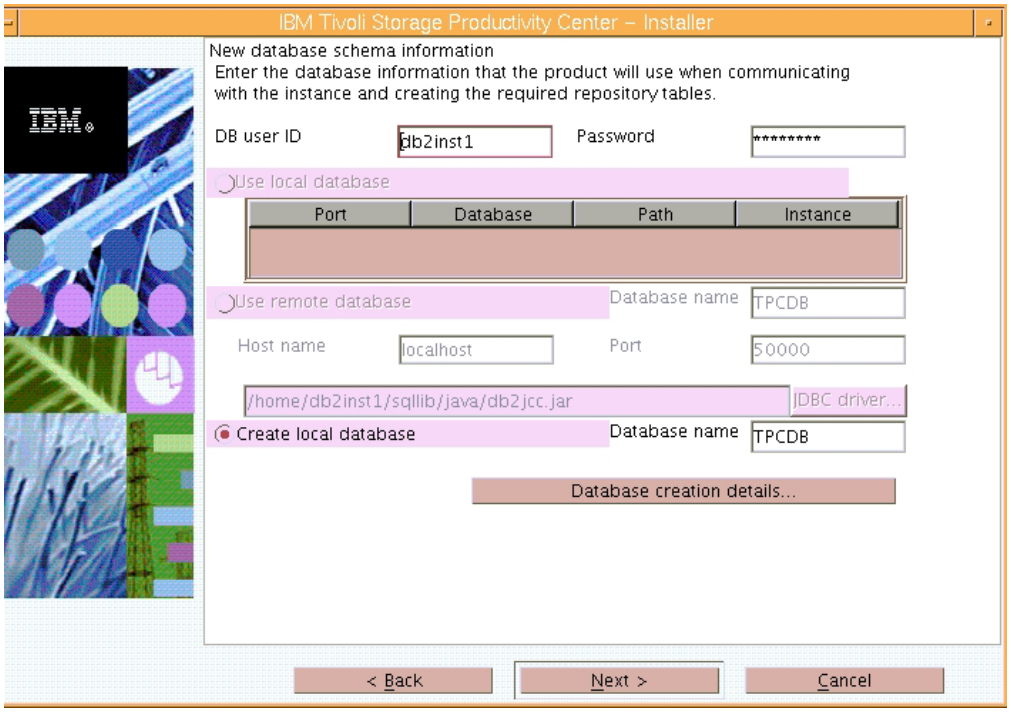


Figure 17. New database schema information panel

If you click **Database creation details**, you see the Database schema creation information panel. Use this panel to change the database schema. Do not change the default values unless you are a knowledgeable DB2 administrator. The Database Management (DMS) and Automatic Storage options are for

advanced users only. For information about table space allocation, see "Estimating tablespace allocation" on page 427. Also refer to your DB2 documentation. If you do change this information, click **OK**.

Note: The DB2 database schema name for Tivoli Storage Productivity Center cannot be longer than eight characters.

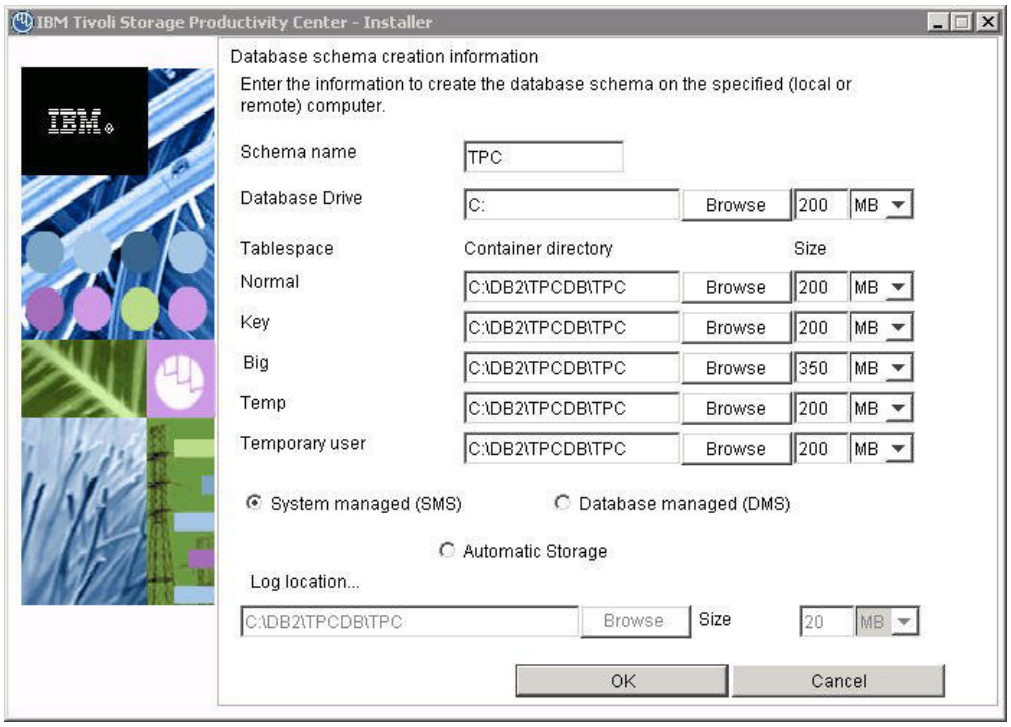


Figure 18. Database schema creation information panel

9. The Data server, Device server, and Storage Resource Agent information panel is displayed.

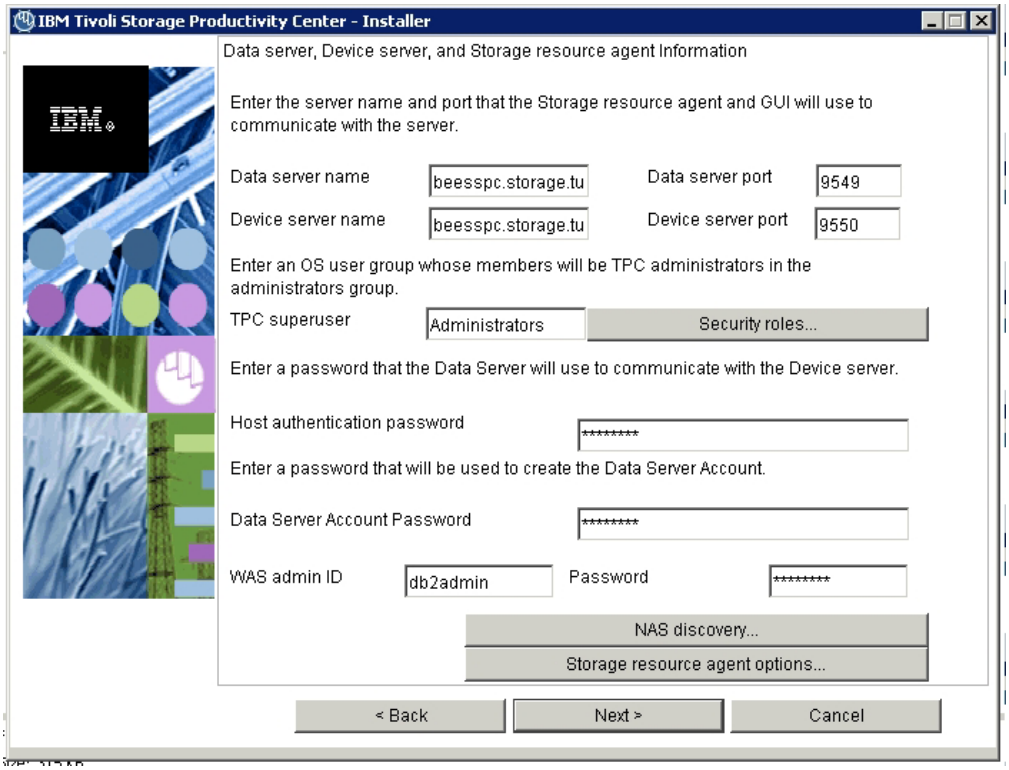


Figure 19. Data server, Device server, and Storage Resource Agent information panel

Enter the following information:

Data Server Name
Enter the fully qualified host name of the Data server.

Data Server Port
Enter the Data server port. The default is 9549.

Device Server Name
Enter the fully qualified host name of the Device server.

Device Server Port
Enter the Device server port. The default is 9550.

TPC Superuser
Enter an operating system group name to associate with the TPC Superuser role. This group must exist in your operating system before you install Tivoli Storage Productivity Center. Membership in this group provides full access to the Tivoli Storage Productivity Center product. You can assign a user ID to this group on your operating system and start the Tivoli Storage Productivity Center GUI using this user ID.

Note: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the value you enter for LDAP TPC Administrator group overrides the value you entered here for the TPC Superuser.

Host authentication password
Password used by the Storage Resource agent to communicate with the Device server. This password must be specified when you install the Device server.

Data Server Account Password

For Windows only. You supply this password. When you supply this password, the Tivoli Storage Productivity Center installation program creates a user ID and the password you supplied to create the Data server account on Windows. This user ID and password are used by the Data server service.

The display name for the Data Server is:

```
IBM Tivoli Storage Productivity Center - Data Server
```

WAS admin ID and Password

The user ID and password required by the Device server to communicate with embedded WebSphere. This is only used at installation time.

Note: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the values you enter for LDAP TPC Administrator user name and password override the value you entered here for WAS Admin ID and password.

Click **Next**.

If you click the Security roles button:

The Advanced security roles mapping panel is displayed. Enter the operating system group for each Tivoli Storage Productivity Center role you want to make an association with. The operating group must exist before you can associate a Tivoli Storage Productivity Center role with a group. Click **OK**. Click **Next**.

Note: You do not have to assign security roles at installation time. You can assign these roles after you have installed Tivoli Storage Productivity Center.

If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the values you enter here for Security Roles Mapping will be deleted and you will have to assign roles to individuals after the installation is complete (use the **Role-to-Group Mapping** node).

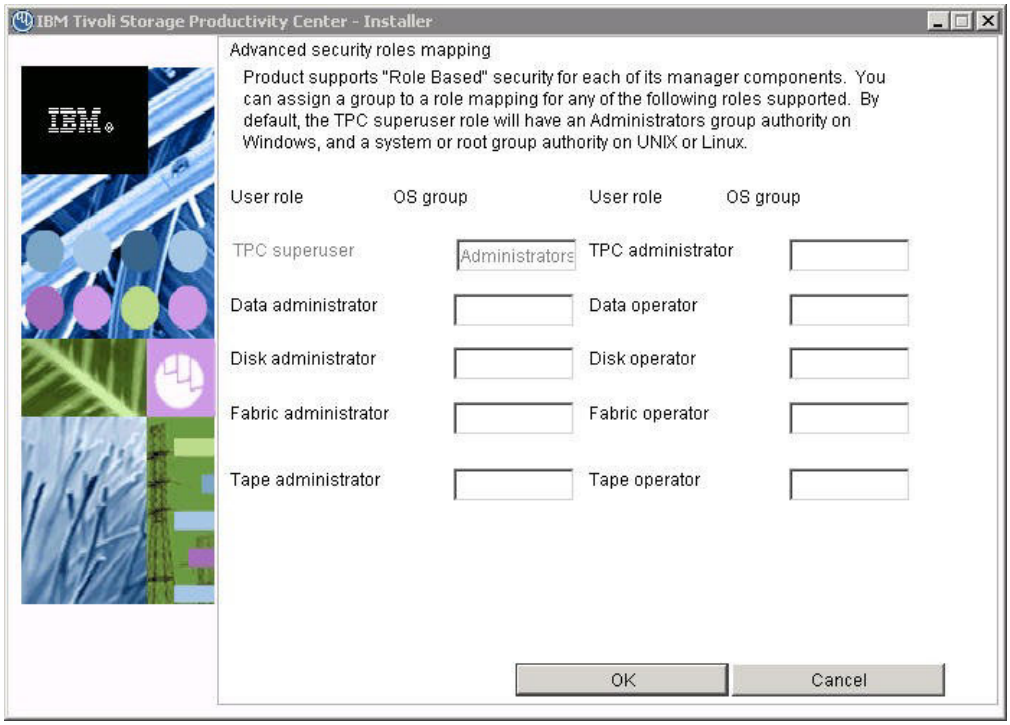


Figure 20. Advanced security roles mapping panel

If you click the NAS discovery button:

The NAS discovery information panel is displayed. Enter the NAS filer login default user name and password. If you want to enter the SNMP communities to be used for NAS discovery, click **Add**, enter the community name, then click **OK**. Click **Next** on the Data server, the Device server, and Agent information panel.

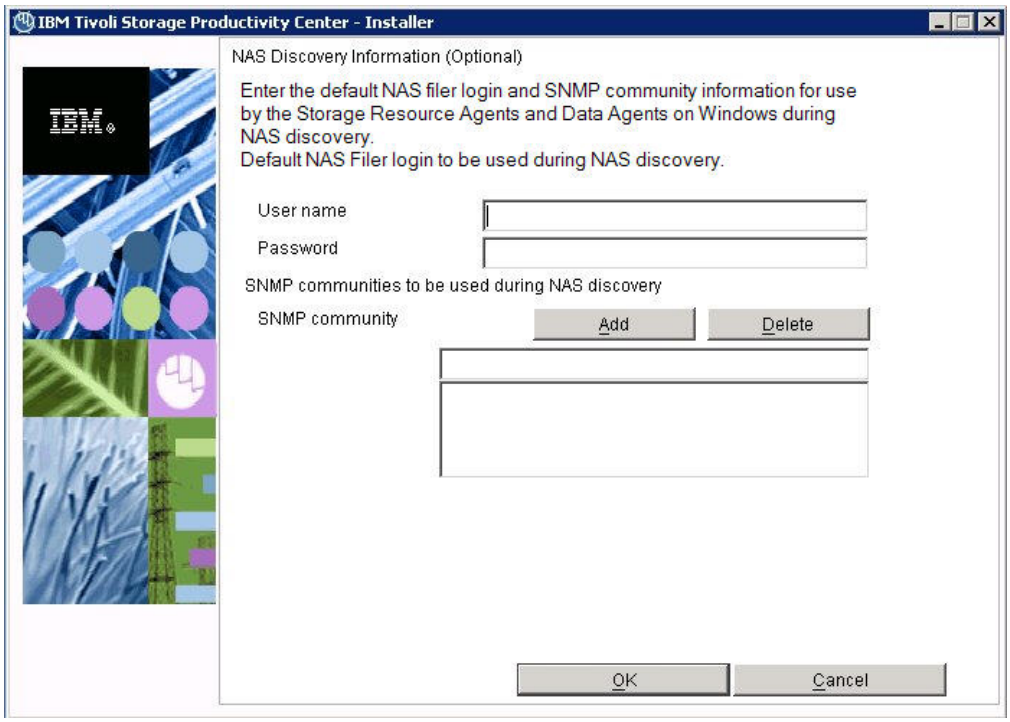


Figure 21. NAS discovery information panel

If you click the Storage Resource Agent options button:

The Storage Resource Agent panel is displayed. You can select whether or not you want the agent to perform an initial scan when first brought up. Make a selection and click **OK**.

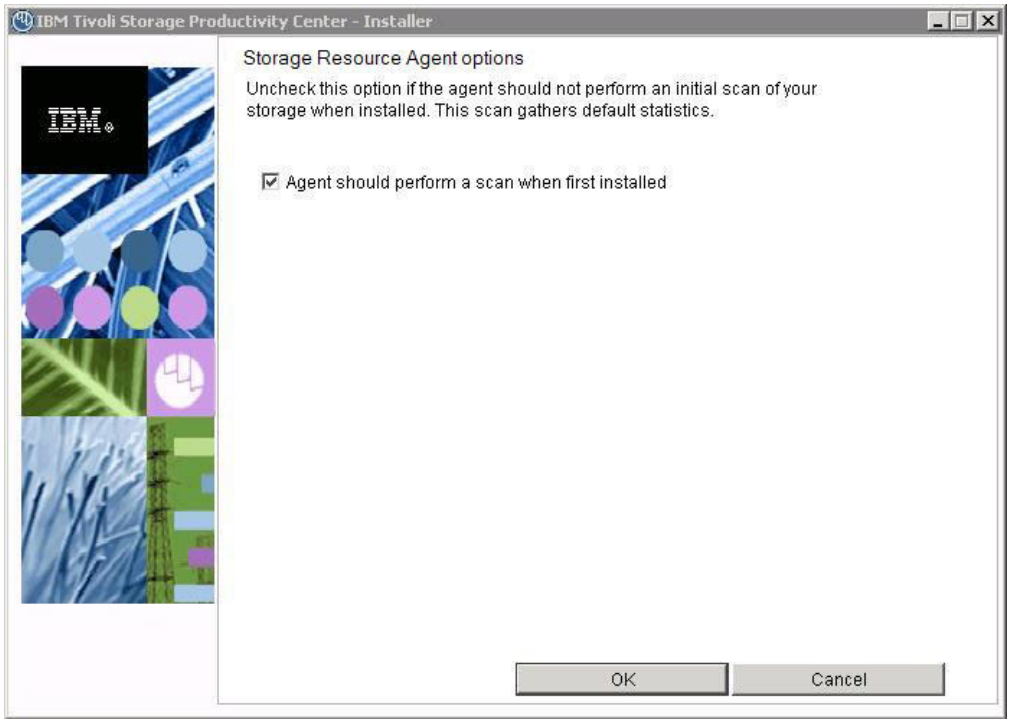


Figure 22. Storage Resource Agent options panel

Agent should perform a scan when first installed
Clear the check box for this option if you do not want Data Manager to perform an initial scan of your storage upon installation. This option is checked by default and gathers default statistics.

10. The Tivoli Integrated Portal installation panel is displayed. This panel lets you select an existing Tivoli Integrated Portal to use or install the Tivoli Integrated Portal program.

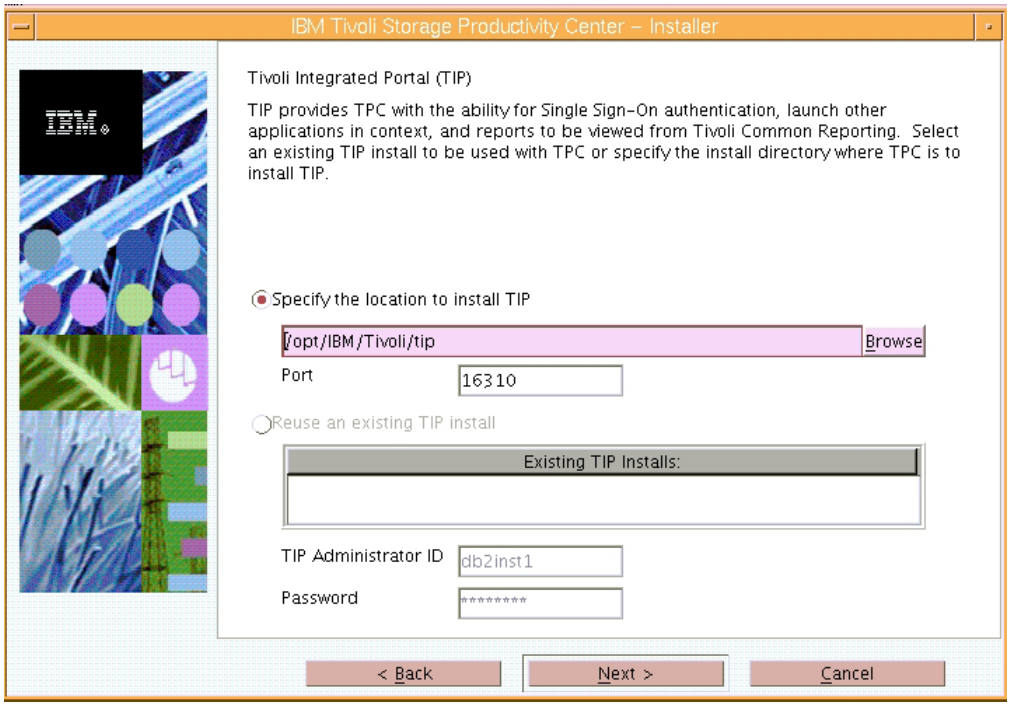


Figure 23. Tivoli Integrated Portal installation panel

Enter the following information:

Specify the location to install TIP
Accept the default or enter a location to install Tivoli Integrated Portal.

Port Enter the port number for Tivoli Integrated Portal or accept the default (16310). Tivoli Integrated Portal uses 10 port numbers. For example, if the default port 16310 is accepted, Tivoli Integrated Portal uses the following ports: 16310, 16311, 16312, 16313, 16315, 16316, 16318, 16320, 16322, and 16323. If you enter a port number other than the default number, ensure that you have the following ports available for Tivoli Integrated Portal:

- *base port*
- *base port* +1
- *base port* +2
- *base port* +3
- *base port* +5
- *base port* +6
- *base port* +8
- *base port* +10
- *base port* +12

- *base port* +13

where *base port* is the port value entered in the **Port** field.

Reuse an existing TIP install
Select this option if you want to use an existing Tivoli Integrated Portal. Under **Existing TIP installs**, select a Tivoli Integrated Portal to use. Enter the following information:

TIP Administrator ID
Enter the Tivoli Integrated Portal administrator ID.

Password
Enter the Tivoli Integrated Portal password of the administrator ID.

Click **Next**.

11. The authentication selection panel is displayed. If a valid IBM Tivoli Integrated Portal instance exists on the system and it uses either the OS-based authentication or LDAP-based authentication, then Tivoli Storage Productivity Center uses that existing authentication method. Otherwise, select the authentication method to use.

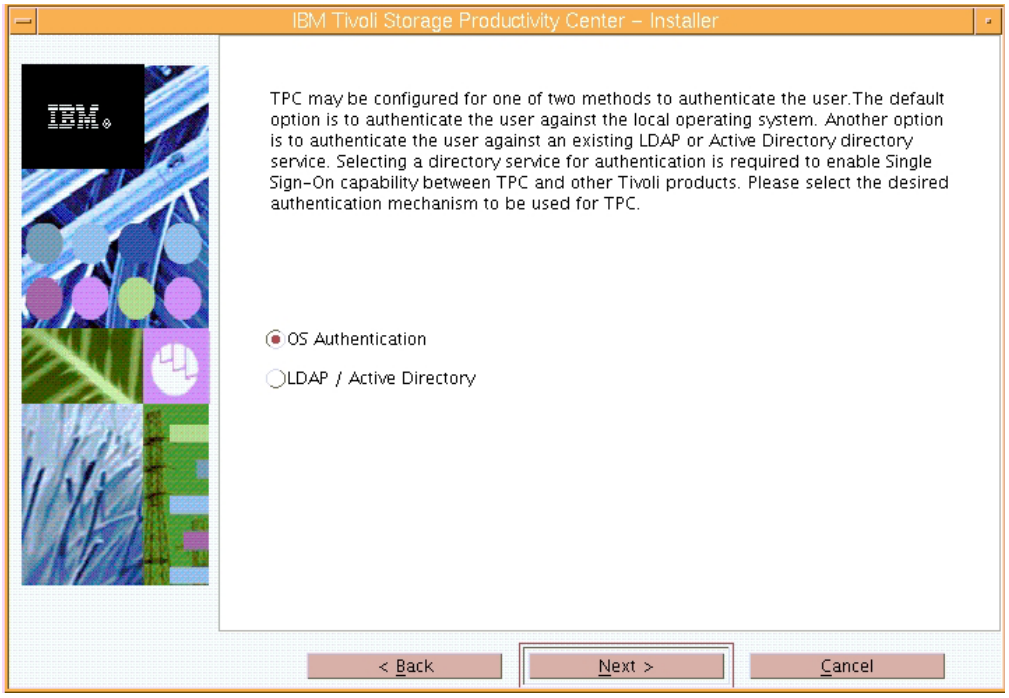


Figure 24. Authentication selection panel

Single sign-on is a method of access control that enables a user to log in once and gain access to the resources of multiple software systems without being prompted to log in again. The single sign-on feature requires a centralized user and group repository, such as an LDAP-compliant directory, that all participating applications can access. The credentials for the user are passed between applications in a secure manner. If you do not want to use LDAP or Active Directory, you can select OS authentication. Select the authentication method to use and click **Next**.The options are:

OS Authentication
Uses the operating system for user authentication. If you select this option, go to step 15 on page 202.

LDAP/Active Directory
If you select LDAP or Microsoft Active Directory for authentication, you must have an LDAP or Active Directory already installed. If you select this option, go to step 12. Click **Next**.

12. If you selected **LDAP/Active Directory**, there is more information that Tivoli Storage Productivity Center needs to be collected.

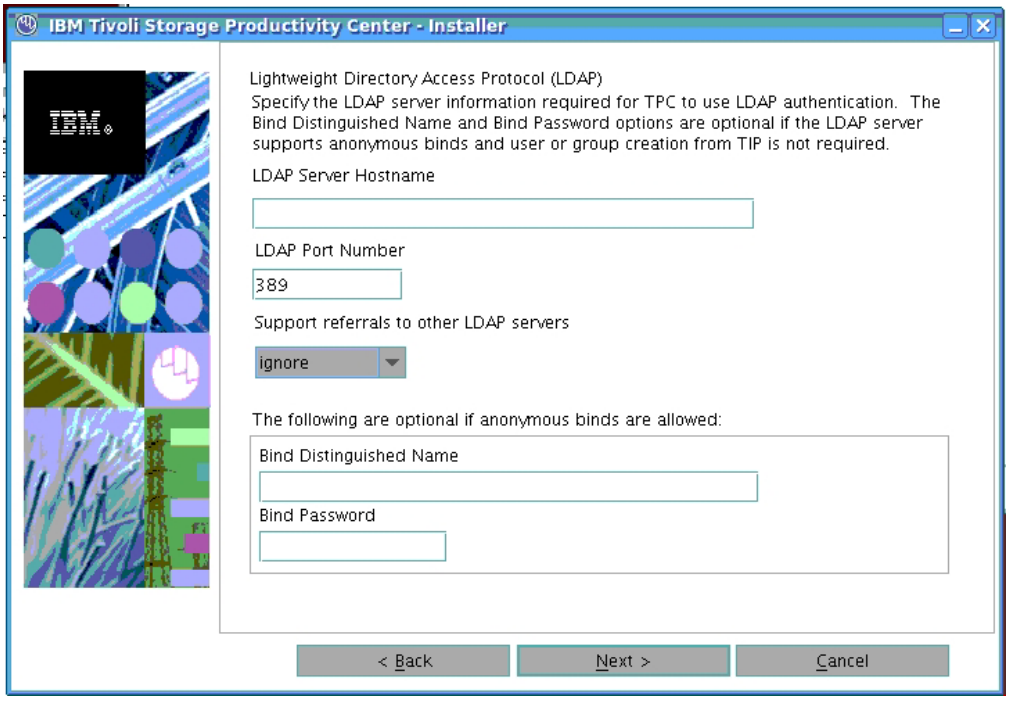


Figure 25. LDAP Server information panel

Enter this information and click **Next**.

LDAP Server Hostname
The fully qualified host and domain name of the machine where your LDAP-based directory is running.

LDAP Port Number
The port number on your LDAP server where the LDAP process is listening for communications.

Support referrals to other LDAP servers
The options are **ignore** or **follow** an LDAP referral. The default is **ignore**.

Referral is the process by which an LDAP server, instead of returning a result to a client application, returns a reference (a referral) to another LDAP server which might contain further information related to the client application's request.

Bind Distinguished Name
The distinguished name used to bind to the LDAP-based directory when performing a user or group search. This name is optional

because some LDAP-based Directories allow anonymous binds and others require you to bind with a specific name and password.

Note: Depending on how the LDAP server is set up, the Bind DN and password might or might not be optional. If the LDAP server allows for anonymous binds, then the Bind DN and password are optional.

If you want the ability to create or modify LDAP users and groups from the Tivoli Integrated Portal admin panel, then the Bind DN and password are required.

Bind Password
The password associated with the Bind Distinguished Name.

Enter this information and click **Next**.

13. The Lightweight Directory Access Protocol (LDAP user and group information) panel is displayed.

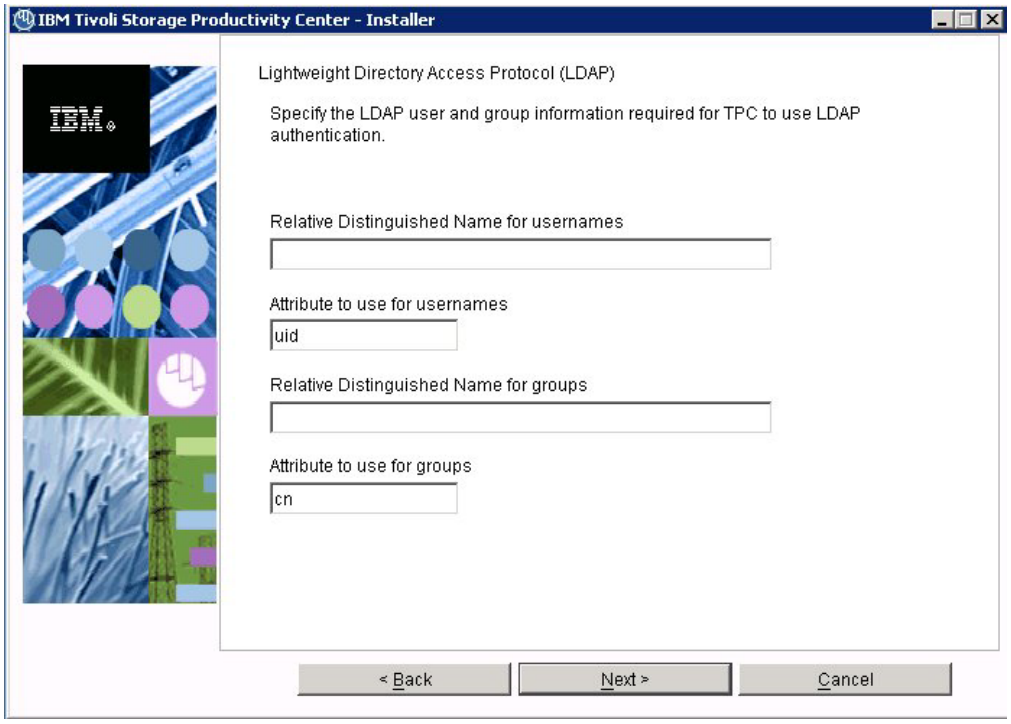


Figure 26. LDAP user and group information panel

Enter this information and click **Next**.

Relative Distinguished Name for usernames
Tells Tivoli Storage Productivity Center where to start a search in the directory when performing user authentication.

Attribute to use for usernames
Tells Tivoli Storage Productivity Center which attribute in a directory entry for the user that contains the name for the user for authentication.

Relative Distinguished Name for groups
Tells Tivoli Storage Productivity Center where in the directory to start a search when performing a group search for authorization.

Attribute to use for groups
Tells Tivoli Storage Productivity Center which attribute in a directory entry for the group that contains the name for the group for authorization.

14. The Lightweight Directory Access Protocol (LDAP TPC Administrator user and group information) panel is displayed.

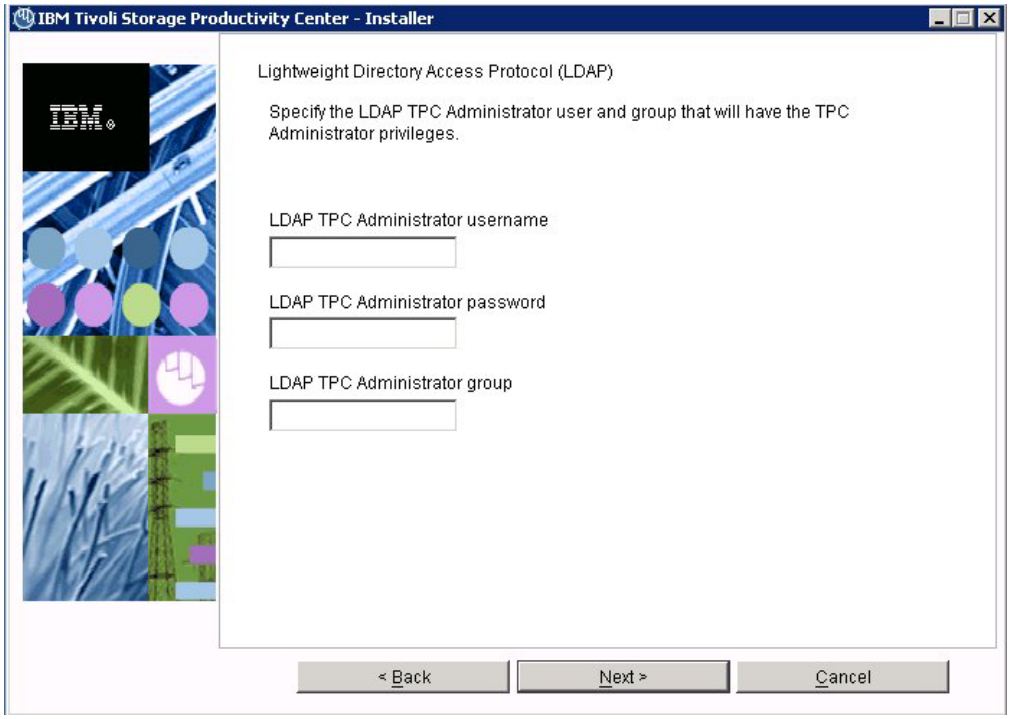


Figure 27. LDAP TPC Administrator user and group information panel

Enter information for these fields:

LDAP TPC Administrator username
The user ID that Tivoli Storage Productivity Center configures as the Tivoli Storage Productivity Center superuser during installation. This user name must exist in the directory before you start the Tivoli Storage Productivity Center installation. The user name must be a member of your designated LDAP Tivoli Storage Productivity Center Administrator group in the directory.

Note: If you install Tivoli Storage Productivity Center on a Windows system and have Tivoli Storage Productivity Center authenticate users against an LDAP-based repository, then the LDAP TPC Administrator user name must **not** contain a space or blank character in it. See the WebSphere Application Server APAR PK77578.

Ensure that you use the correct case when entering the values for the LDAP TPC Administrator username and LDAP TPC Administrator group. If you use the incorrect case when entering these values, then you cannot log in to Tivoli Storage Productivity Center for Replication with the user name you expected to use.

LDAP TPC Administrator password
Password for the administrator.

LDAP TPC Administrator group

The group that Tivoli Storage Productivity Center maps to the Tivoli Storage Productivity Center superuser role during installation. This group must exist in the directory before you start the IBM Tivoli Storage Productivity Center installation. The group must have the designated LDAP Tivoli Storage Productivity Center Administrator as a member.

Note:

Ensure that you use the correct case when entering the values for the LDAP TPC Administrator user name and LDAP TPC Administrator group. If you use the incorrect case when entering these values, then you cannot log in to Tivoli Storage Productivity Center for Replication with the user name you expected to use.

15. The Summary Information panel is displayed. Review the information. Click **Install**.
16. You see the installing panel. When Tivoli Integrated Portal has finished installing, Tivoli Storage Productivity Center starts the Tivoli Storage Productivity Center for Replication installation.
17. Tivoli Storage Productivity Center starts the Tivoli Storage Productivity Center for Replication installation program.
    a. The Welcome panel is displayed. Click **Next**.
    b. The System prerequisites check panel is displayed. The installation wizard checks whether the prerequisites are installed, then confirms whether your operating system is supported and is at the appropriate fix pack or update level. Click **Next**.
    c. The License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.
    d. The Directory Name panel is displayed. Accept the default installation directory by clicking **Next**, or specify a different installation directory and click **Next**.
    e. The Tivoli Storage Productivity Center for Replication Administrator ID and Password panel is displayed. Enter the administrator user ID and password. This ID is typically the operating system administrator user ID. If you use a different ID, create it beforehand and ensure that it has administrator rights.

       Note: There is a limitation on the number of characters for the user name and password on AIX. At the time this document was written, AIX did not properly validate passwords that are longer than eight characters.
    f. The Default ports panel is displayed. Accept the defaults. Click **Next**.

       The installation program checks for port conflicts on 5110 and 5120. If there is a conflict, a new port is assigned by increasing the port number, and then retesting for a port conflict until an available port is found. The default port numbers in the fields are set to the available ports when the panel is displayed.

       If you assign a port for the Default Client Port or Default Standby Server Port and there is a port conflict, an error message is displayed to let you know of the conflict.

The client port must be set to the same number on both management servers (the active and the standby servers). The standby server port must also be set to the same number on both management servers (the active and the standby servers).

g. The settings panel is displayed. Review the settings and change them if needed by clicking **Back**. Otherwise, click **Install** to begin the installation.

h. You see the status panel. Wait for installation to complete.

i. You see the summary information panel. Review the information and click **Finish**.

j. Tivoli Storage Productivity Center for Replication is installed with a basic license. If you want to install the Two Site or Three Site Business Continuity (BC) license, see "Installing the Two Site or Three Site Business Continuity license" on page 184.

18. The successfully installed panel is displayed. Click **Finish**.
19. To access the Tivoli Storage Productivity Center user interface, see "Starting the Tivoli Storage Productivity Center GUI using a web browser" on page 347. To access the Tivoli Storage Productivity Center for Replication user interface, go to the Tivoli Storage Productivity Center information center. Search for **Accessing the IBM Tivoli Storage Productivity Center for Replication GUI**.
20. After you have installed Tivoli Storage Productivity Center, you can use the Configure Devices wizard to configure your system. For information about the Configure Devices wizard, see "Configure Devices" on page 266.

Installing the database schema

This topic provides information about how to install the database schema for Tivoli Storage Productivity Center. You must create a database schema before you can use Tivoli Storage Productivity Center. If you are using a remote database for Tivoli Storage Productivity Center, you must install the database schema on that remote computer after you have installed DB2. This option also allows you to specify the type of table space you want to use to store your data on DB2.

Note:

- The DB2 database schema name for Tivoli Storage Productivity Center cannot be longer than eight characters.

To install the database schema, complete the following steps:

1. If you are installing the database on AIX or Linux, source the `db2profile`.
2. Start the Tivoli Storage Productivity Center installation program. For information about starting the installation program, see "Starting the installation program" on page 169.
3. The Select a language panel is displayed. Select a language and click **OK**.
4. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.
5. The Select the type of installation you want to run panel is displayed. Select **Custom installation**. Enter the Tivoli Storage Productivity Center Installation Location or accept the default. The default location for Windows is `C:\Program Files\IBM\TPC`. The default location for UNIX or Linux is `/opt/IBM/TPC`. Click **Next**.
6. The Select one or more components to install panel is displayed. Select **Create database schema**. Click **Next**.

7. The Database administrator information panel is displayed. Enter the database administrator user ID and password. This user ID and password were specified when you installed DB2.
8. The New database schema information panel is displayed.

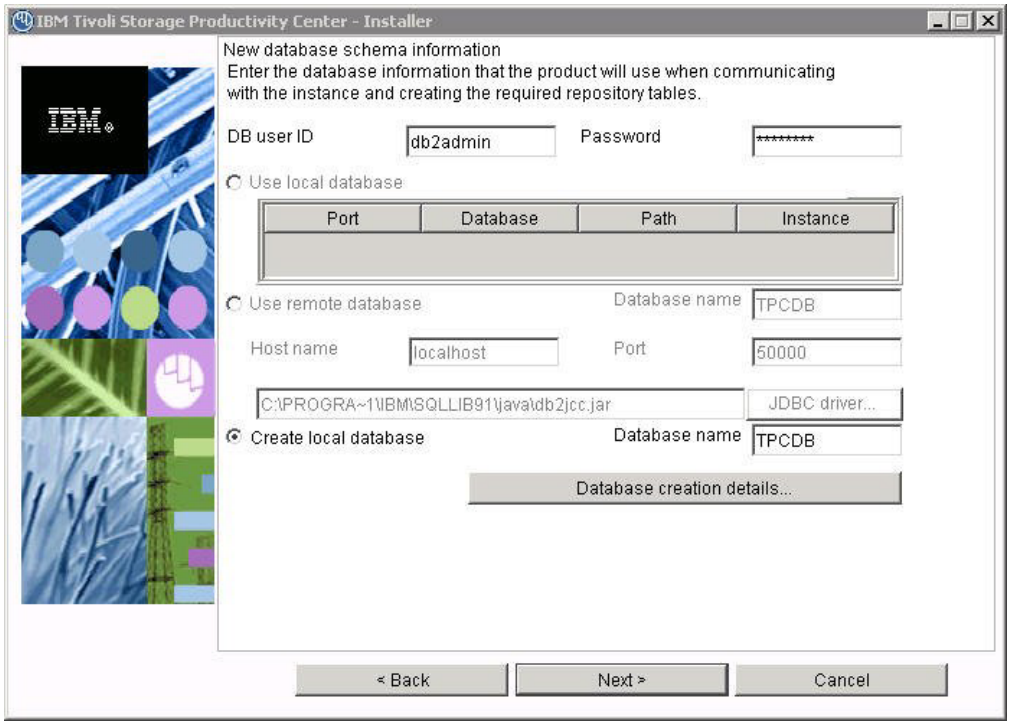


Figure 28. New database schema information panel

Select **Create local database** and enter the database name or accept the default (TPCDB). Click **Next**.

If you click **Database creation details**, you see the Database schema creation information panel. This option allows you to change the database schema.

Do not change the default values unless you are a knowledgeable DB2 administrator. For information about table space allocation, see "Estimating tablespace allocation" on page 427. Also refer to your DB2 documentation.

Make your changes and click **OK**. Click **Next** on the New database schema information panel.

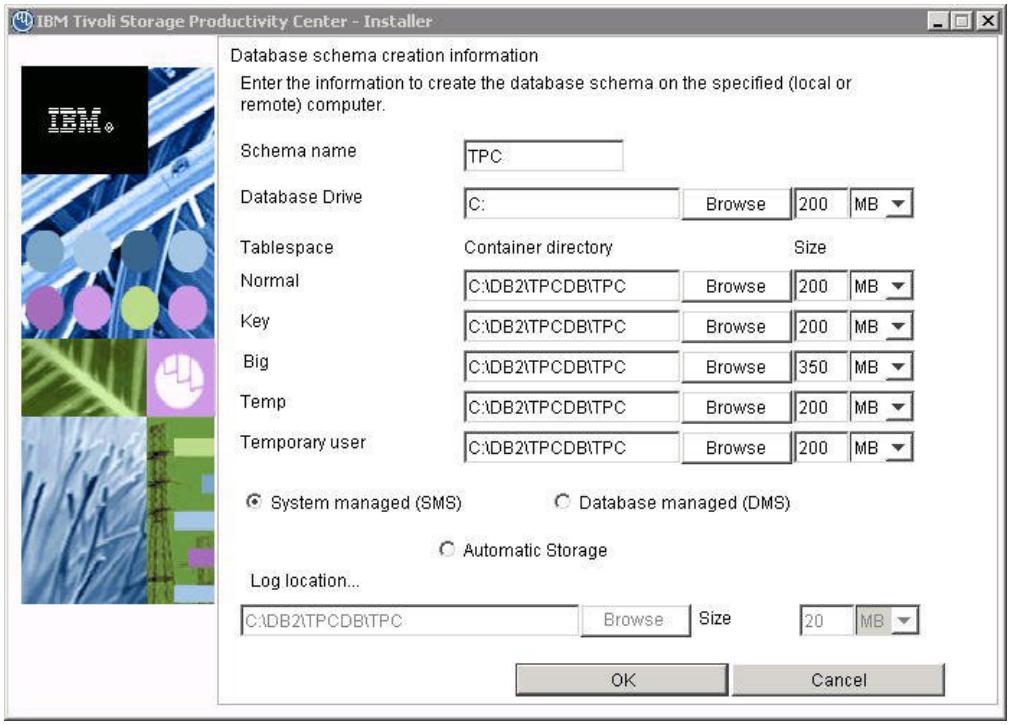


Figure 29. Database schema creation information panel

Note:

- If you are installing on AIX, change the default location of the database schema to a file system with adequate space. The DB2 administrator and the Tivoli Storage Productivity Center installer must have authority to modify the database. The default is for Tivoli Storage Productivity Center to create a database in a directory homed off root.

9. The Summary Information panel is displayed. Review the information. Click **Install**.
10. You see the installing panel. Wait for installation to complete.
11. The successfully installed panel is displayed. Click **Finish**.

Installing the Data server and Device server using a local database

This procedure provides information about how to install the IBM Tivoli Storage Productivity Center servers using a local database.

To install the IBM Tivoli Storage Productivity Center servers, complete the following steps:

1. Start the IBM Tivoli Storage Productivity Center installation program. For information about starting the installation program, see "Starting the installation program" on page 169.
2. The Select a language panel is displayed. Select a language and click **OK**.
3. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.

   Note: 8 GB of RAM is required. If you have 4-8 GB of RAM, you can still install Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication. However, you get a warning message. To avoid system

load problems, run only Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication. For information about how to disable Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication after installation, see "Disabling unlicensed or unused Tivoli Storage Productivity Center components" on page 368.

4. Go to the Select the type of installation you want to run panel. Select **Custom installation**. Click **Next**.
5. On the Select one or more components to install panel, select **Tivoli Storage Productivity Center Servers** and **Create database schema**. Click **Next**.
6. The Database administrator information panel is displayed. Enter a database administrator user ID and password. Click **Next**.
7. The New database schema information panel is displayed. Select **Create local database**. Click **Next**.

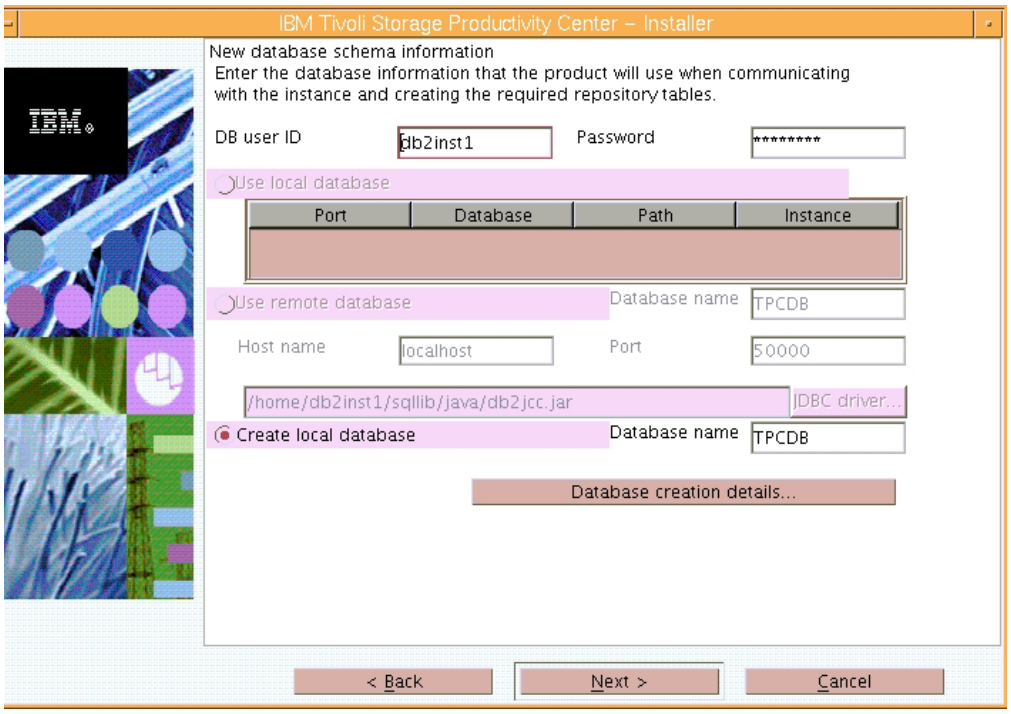


Figure 30. New database schema information panel

If you click **Database creation details**, you see the Database schema creation information panel. This panel allows you to change the default space assigned t the database. You do not need to change these values and you can accept the default values. If you do change this information, click **OK**.

Note: The DB2 database schema name for Tivoli Storage Productivity Center cannot be longer than eight characters.

8. The Data server, Device server, and Storage Resource Agent information panel is displayed. Enter the following information:

Data Server Name
Enter the fully qualified host name of the Data server.

Data Server Port
Enter the Data server port. The default is 9549.

Device server Name
Enter the fully qualified host name of the Device server.

Device server Port
Enter the Device server port. The default is 9550.

TPC Superuser
Enter an operating system group name to associate with the TPC Superuser role. This group must exist in your operating system before you install Tivoli Storage Productivity Center. Membership in this group provides full access to the Tivoli Storage Productivity Center product. You can assign a user ID to this group on your operating system and start the Tivoli Storage Productivity Center GUI using this user ID.

Note: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the value you enter for LDAP TPC Administrator group overrides the value you entered here for the TPC Superuser. For more information about security roles, see "Planning for IBM Tivoli Storage Productivity Center authorization" on page 33.

Host authentication password
The password used by the Storage Resource agent to communicate with the Device server. This password must be specified when you install the Storage Resource agent.

Data Server Account Password
For Windows only. You supply this password. When you supply this password, the Tivoli Storage Productivity Center installation program creates a user ID and the password you supplied to create the Data server account on Windows. This user ID and password are used by the Data server service.

The display name for the Data Server is:

```
IBM Tivoli Storage Productivity Center - Data Server
```

WAS admin ID and Password
The WebSphere administration user ID and password required by the Device server to communicate with embedded WebSphere, if this user ID does not exist.

Note: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the values you enter for LDAP TPC Administrator user name and password override the value you entered here for WAS admin ID and password.

Click **Next**.

If you click the Security roles button:

The Advanced security roles mapping panel is displayed. Enter the operating system group for each Tivoli Storage Productivity Center role you want to make an association with. You do not have to assign security roles at installation time. You can assign these roles after you have installed Tivoli Storage Productivity Center. Click **OK**. Click **Next** on the Data server, the Device server, and Storage Resource agent information panel.

Note: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the values you enter here for Security

Roles Mapping are deleted and you will have to reset the mappings after the installation is complete using the **Role-to-Group Mappings** node.

9. The IBM Tivoli Integrated Portal Information panel is displayed. This panel lets you select an existing IBM Tivoli Integrated Portal to use or install the IBM Tivoli Integrated Portal program. Enter the following information:

   Specify the location to install TIP
   Accept the default or enter a location to install Tivoli Integrated Portal.

   Port Enter the port number for Tivoli Integrated Portal or accept the default (16310). Tivoli Integrated Portal uses 10 port numbers. For example, if the default port 16310 is accepted, Tivoli Integrated Portal uses the following ports: 16310, 16311, 16312, 16313, 16315, 16316, 16318, 16320, 16322, and 16323. If you enter a port number other than the default number, ensure that you have the following ports available for Tivoli Integrated Portal:

   - *base port*
   - *base port* +1
   - *base port* +2
   - *base port* +3
   - *base port* +5
   - *base port* +6
   - *base port* +8
   - *base port* +10
   - *base port* +12
   - *base port* +13

   where *base port* is the port value entered in the **Port** field.

   Reuse an existing TIP install
   Select this option if you want to use an existing IBM Tivoli Integrated Portal. Under **Existing TIP installs**, select a IBM Tivoli Integrated Portal to use. Enter the following information:

   TIP Administrator ID
   Enter the IBM Tivoli Integrated Portal administrator ID.

   Password
   Enter the IBM Tivoli Integrated Portal password of the administrator ID.

   Click **Next**.

10. The authentication selection panel is displayed. If a valid IBM Tivoli Integrated Portal instance exists on the system and it uses either OS-based authentication or LDAP-based authentication, then IBM Tivoli Storage Productivity Center uses that existing authentication method. Otherwise, select the authentication method to use.

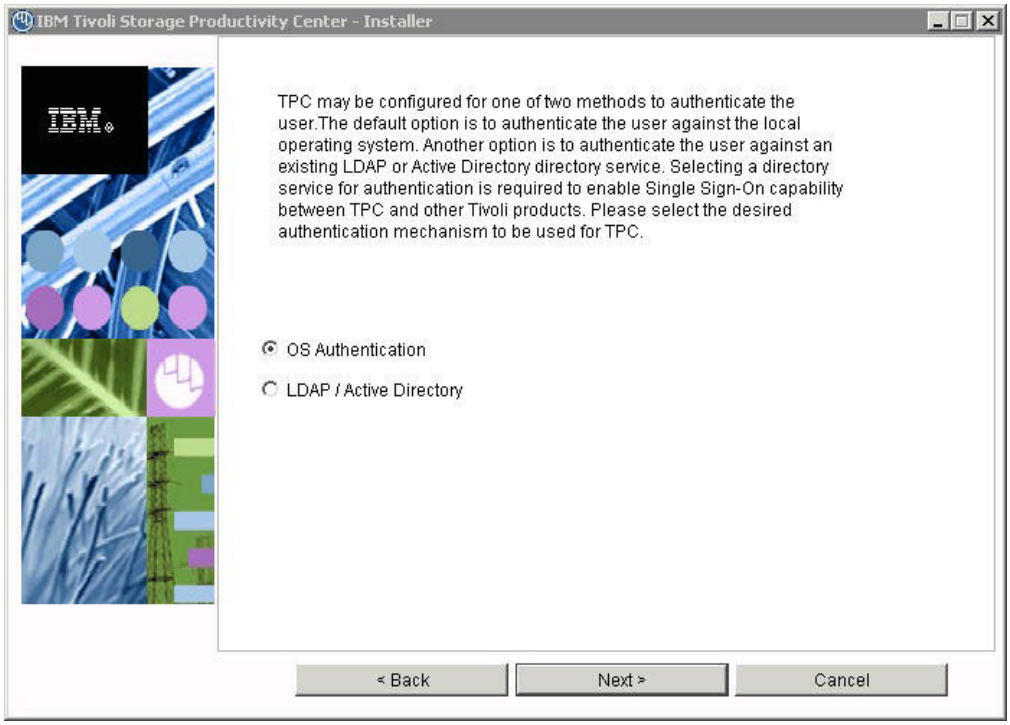


Figure 31. Authentication selection panel

Single sign-on is a method of access control that enables a user to log in once and gain access to the resources of multiple software systems without being prompted to log in again. The single sign-on feature requires a centralized user and group repository, such as an LDAP-compliant directory, that all participating applications can access. The credentials for the user are passed between applications in a secure manner. If you do not want to use LDAP or Active Directory, you can select OS authentication. Select the authentication method to use and click **Next**.The options are:

LDAP/Active Directory
If you select LDAP or Microsoft Active Directory for authentication, you must have an LDAP or Active Directory already installed. If you select this option, go to step 11. Click **Next**.

OS Authentication
Uses the operating system for user authentication. If you select this option, go to step 13 on page 213. Click **Next**.

11. If you selected **LDAP/Active Directory** in step 10, there is more information that Tivoli Storage Productivity Center needs to collect.

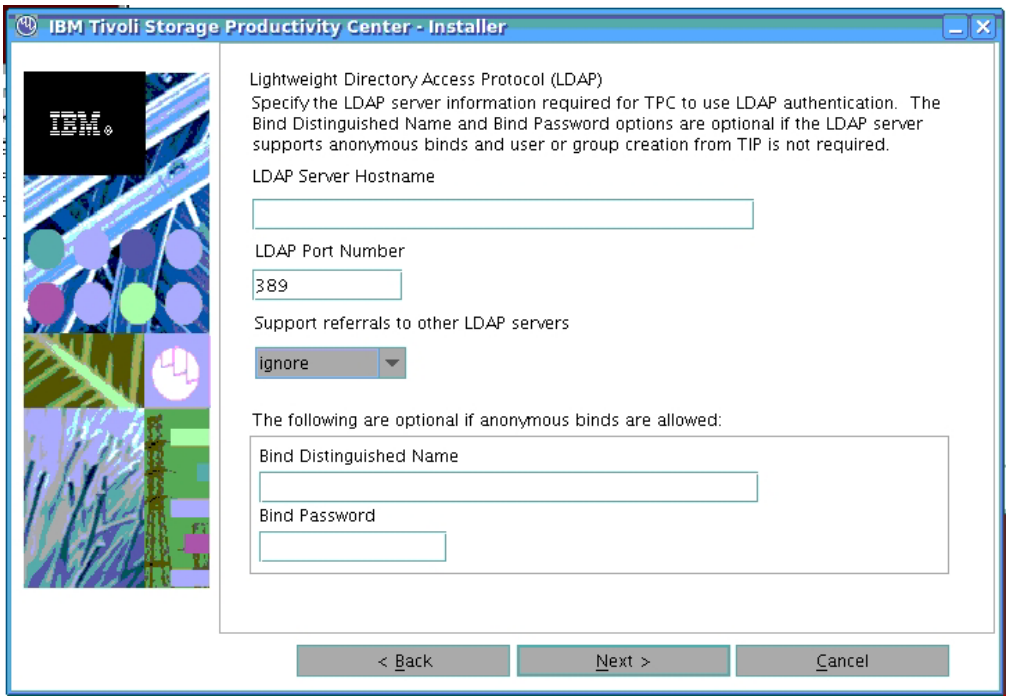


Figure 32. LDAP Server information panel

Enter this information and click **Next**.

LDAP Server Hostname
The fully qualified host and domain name of the machine where your LDAP-based directory is running.

LDAP Port Number
The port number on your LDAP server where the LDAP process is listening for communications.

Support referrals to other LDAP servers
The options are **ignore** or **follow** an LDAP referral. The default is **ignore**.

Referral is the process by which an LDAP server, instead of returning a result to a client application, returns a reference (a referral) to another LDAP server which might contain further information related to the client application's request.

Bind Distinguished Name
The distinguished name used to bind to the LDAP-based directory when performing a user or group search. This name is optional because some LDAP-based Directories allow anonymous binds and others require you to bind with a specific name and password.

Note: Depending on how the LDAP server is set up, the Bind DN and password might or might not be optional. If the LDAP server allows for anonymous binds, then the Bind DN and password are optional. If you want the ability to create or modify LDAP users and groups from the Tivoli Integrated Portal administrative panel, then the Bind DN and password are required.

Bind Password
The password associated with the Bind Distinguished Name.

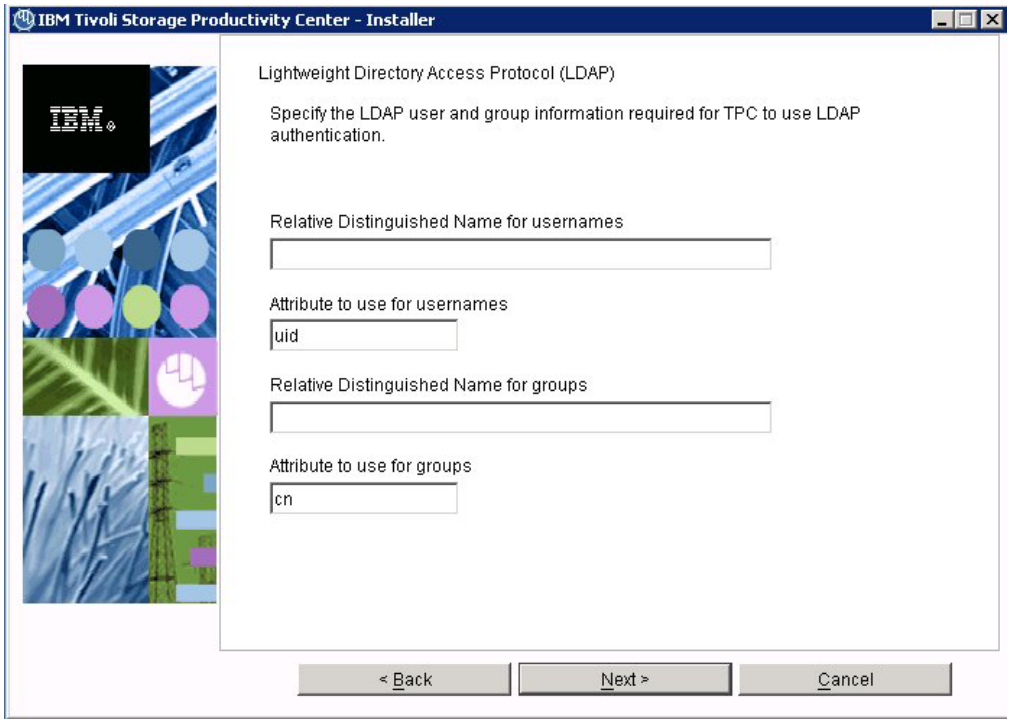


Figure 33. LDAP user and group information panel

Enter this information and click **Next**.

Relative Distinguished Name for usernames
Tells IBM Tivoli Storage Productivity Center where in the directory to start a search when performing user authentication.

Attribute to use for usernames
Tells IBM Tivoli Storage Productivity Center which attribute in a directory entry contains the name for the user for authentication.

Relative Distinguished Name for groups
Tells IBM Tivoli Storage Productivity Center where in the directory to start a search when performing a group search for authorization.

Attribute to use for groups
Tells IBM Tivoli Storage Productivity Center which attribute in a directory entry contains the name for the group for authorization.

12. The LDAP TPC Administrator user and group information panel are displayed.

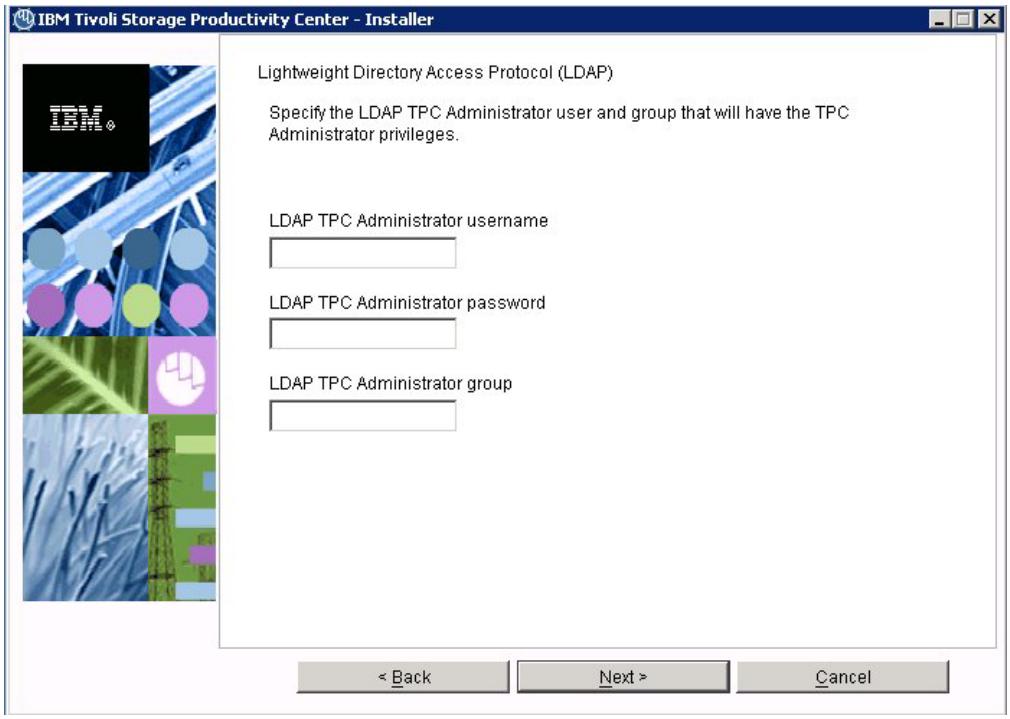


Figure 34. LDAP TPC Administrator user and group information panel

Enter information for these fields:

LDAP TPC Administrator username
The user ID that IBM Tivoli Storage Productivity Center configures as the IBM Tivoli Storage Productivity Center superuser during installation. This user name must exist in the directory before you start the IBM Tivoli Storage Productivity Center installation. The user name must be a member of the LDAP IBM Tivoli Storage Productivity Center Administrator group in the directory.

Note: If you install Tivoli Storage Productivity Center on a Windows system and have Tivoli Storage Productivity Center authenticate users against an LDAP-based repository, then the LDAP Tivoli Storage Productivity Center Administrator user name must **not** contain a space or blank character in it. See the WebSphere Application Server APAR PK77578.

Ensure that you use the correct case when entering the values for the LDAP TPC Administrator user name and LDAP TPC Administrator group. If you use the incorrect case when entering these values, then you cannot log in to Tivoli Storage Productivity Center for Replication with the user name you expected to use.

LDAP TPC Administrator password
Password for the administrator.

LDAP TPC Administrator group
The group that IBM Tivoli Storage Productivity Center maps to the IBM Tivoli Storage Productivity Center superuser role during installation. This group must exist in the directory before you start the

IBM Tivoli Storage Productivity Center installation. The group must have the LDAP IBM Tivoli Storage Productivity Center Administrator as a member.

Note:

Ensure that you use the correct case when entering the values for the LDAP TPC Administrator user name and LDAP TPC Administrator group. If you use the incorrect case when entering these values, then you cannot log in to Tivoli Storage Productivity Center for Replication with the user name you expected to use.

13. The Summary Information panel is displayed. Review the information. Click **Install**.
14. You see the installing panel. When Tivoli Integrated Portal has finished installing, Tivoli Storage Productivity Center starts the Tivoli Storage Productivity Center for Replication installation.
15. Tivoli Storage Productivity Center starts the IBM Tivoli Storage Productivity Center for Replication installation program.
    a. The Welcome panel is displayed. Click **Next**.
    b. The System prerequisites check panel is displayed. The installation wizard checks whether the prerequisites are installed, then confirms whether your operating system is supported and is at the appropriate fix pack or update level. Click **Next**.
    c. The License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.
    d. The Directory Name panel is displayed. Accept the default installation directory by clicking **Next**, or specify a different installation directory and click **Next**.
    e. The Tivoli Storage Productivity Center for Replication Administrator ID and Password panel is displayed. Enter the administrator user ID and password. This ID is typically the operating system administrator user ID. If you use a different ID, create it beforehand and ensure that it has administrator rights.

       Note: There is a limitation on the number of characters for the user name and password on AIX. At the time this document was written, AIX did not properly validate passwords that are longer than eight characters.
    f. The Default ports panel is displayed. Accept the defaults. Click **Next**.

       The installation program checks for port conflicts on 5110 and 5120. If there is a conflict, a new port is assigned by increasing the port number, and then retesting for a port conflict until an available port is found. The default port numbers in the fields are set to the available ports when the panel is displayed.

       If you assign a port for the Default Client Port or Default Standby Server Port and there is a port conflict, an error message is displayed to let you know of the conflict.

       The client port must be set to the same number on both management servers (the active and the standby servers). The standby server port must also be set to the same number on both management servers (the active and the standby servers).
    g. The settings panel is displayed. Review the settings and change them if needed by clicking **Back**. Otherwise, click **Install** to begin the installation.

h. You see the status panel. Wait for installation to complete.

i. You see the summary information panel. Review the information and click **Finish**.

j. Tivoli Storage Productivity Center for Replication is installed with a basic license. If you want to install the Two Site or Three Site Business Continuity (BC) license, see "Installing the Two Site or Three Site Business Continuity license" on page 184.

16. To access the Tivoli Storage Productivity Center user interface, see "Starting the Tivoli Storage Productivity Center GUI using a web browser" on page 347. To access the Tivoli Storage Productivity Center for Replication user interface, go to the Tivoli Storage Productivity Center information center. Search for **Accessing the IBM Tivoli Storage Productivity Center for Replication GUI**.

17. After you have installed Tivoli Storage Productivity Center, you can use the Configure Devices wizard to configure your system. For information about the Configure Devices wizard, see "Configure Devices" on page 266.

Installing Tivoli Storage Productivity Center using a remote database

This section provides information about how to install Tivoli Storage Productivity Center using a remote database.

You will be working with two server computers:

Server A
Install the DB2 server and Tivoli Storage Productivity Center database on this computer. Run the Tivoli Storage Productivity Center installation program to install the database on this computer.

Server B
Install Tivoli Storage Productivity Center servers on this computer. Do **not** install the Tivoli Storage Productivity Center database on this computer.

Preparing to install the Data server and Device server using a remote database:

This topic provides information about the preparation steps required on the DB2 server before installing the Data server and Device server using a remote database.

Note:

- The DB2 server must have at least 6 GB of RAM.
- Use Tivoli Storage Productivity Center version 4.1.1.95 or later for installation with a remote database.

Before installing the Data server and Device server using a remote database, follow these steps on the DB2 server:

1. Install DB2 on the remote computer if you have not already done so.
2. Ensure that the DB2 instance is running on the DB2 host computer before you run the Tivoli Storage Productivity Center installation program.
3. Extract the Tivoli Storage Productivity Center installation package from the download site or use the Tivoli Storage Productivity Center DVD.

   For AIX, use the GNU tar program (version 1.14 or later to extract the package). You can check the version by running the command **tar - -version** from the directory `/opt/freeware/bin/tar`.
4. For AIX or Linux, source the DB2 profile:

   ```
   . /home/DB2_instance_owner/sqllib/db2profile
   ```

where *DB2_instance_owner* is the DB2 instance owner.

5. Run the Tivoli Storage Productivity Center installation program on the DB2 server computer to install the database schema.

   For example, follow these steps:

   a. Start the Tivoli Storage Productivity Center installation program on the DB2 server computer.

   b. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.

   c. On the Select the type of installation you want to run panel, select **Custom Installation**. Click **Next**.

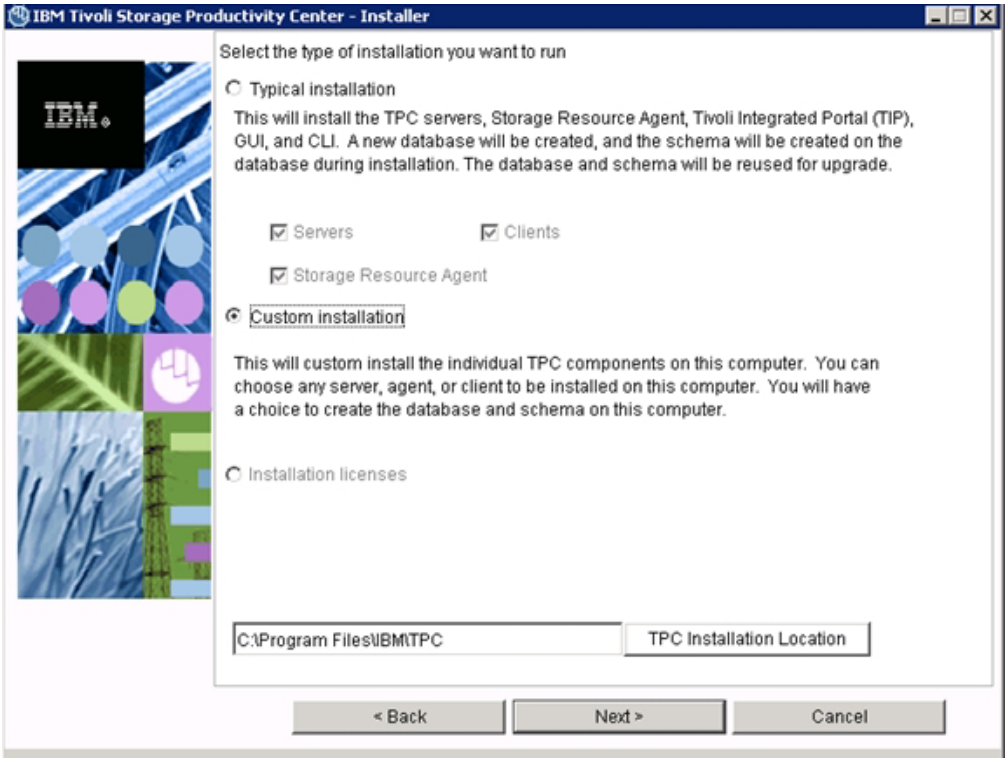


Figure 35. Select the type of installation you want to run panel

   d. On the Select one or more components to install panel, select **Create database schema** and clear all other options. Click **Next**.

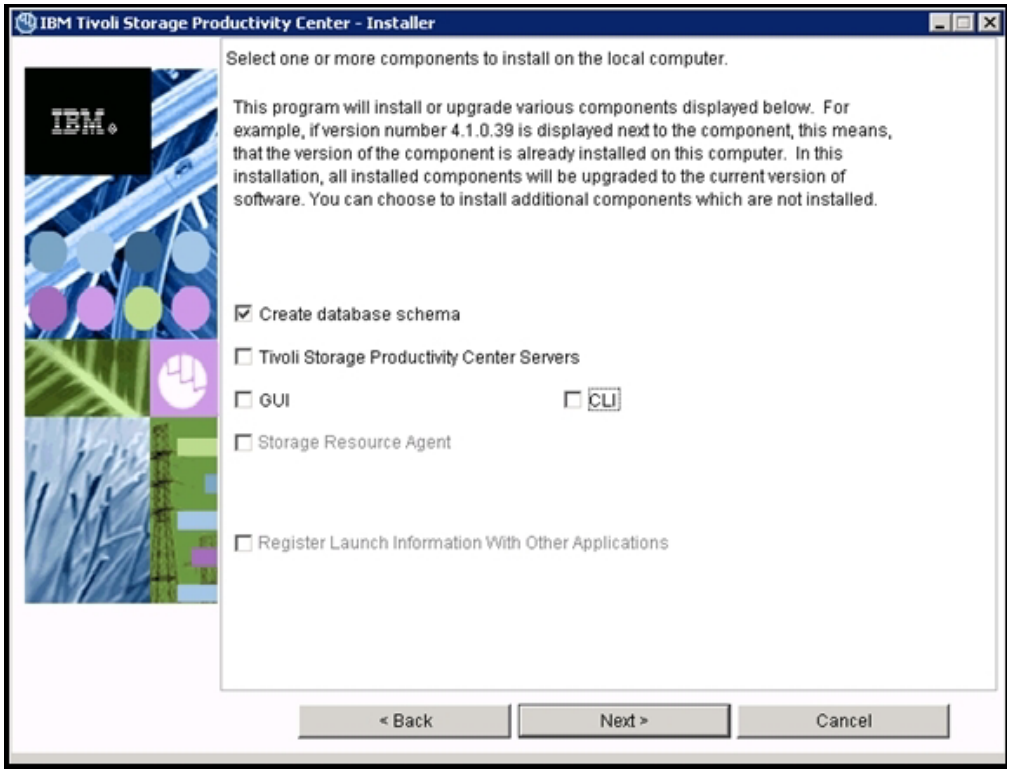


Figure 36. Select one or more components to install panel

e. The Database administrator information panel is displayed. Enter a database administrator user ID and password to create a new database. This user ID must have the authority to access the database and to insert and delete data. Click **Next**.

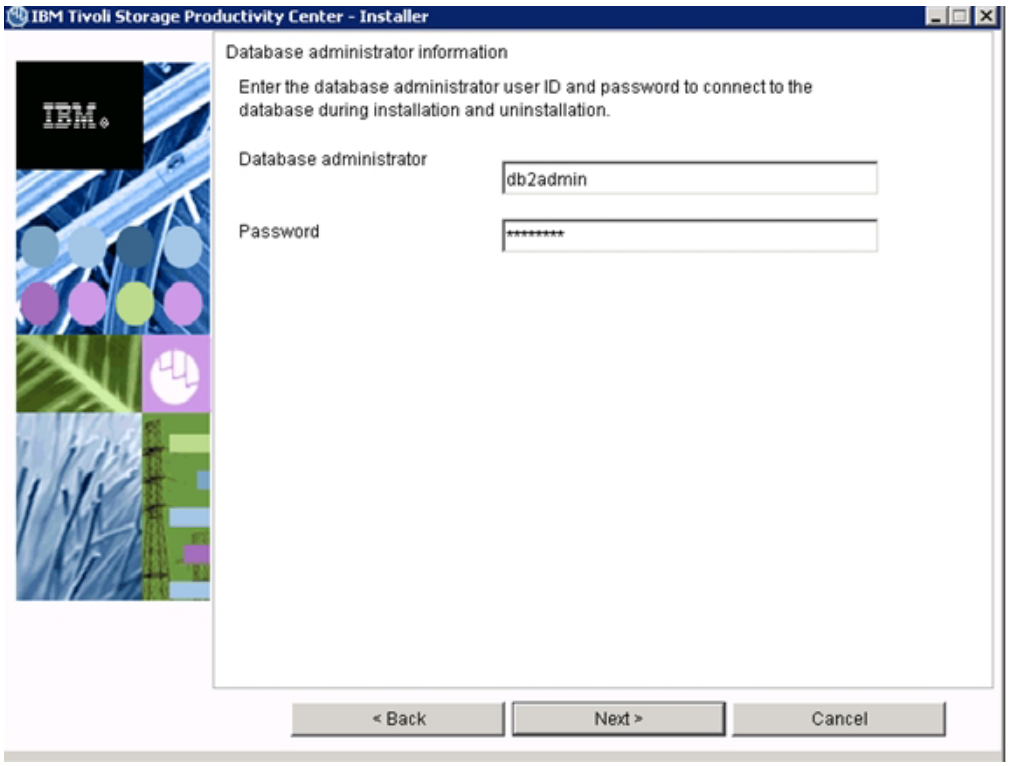


Figure 37. Database administrator Information panel

f. On the New database schema information panel, select **Create local database**. Click **Next**.

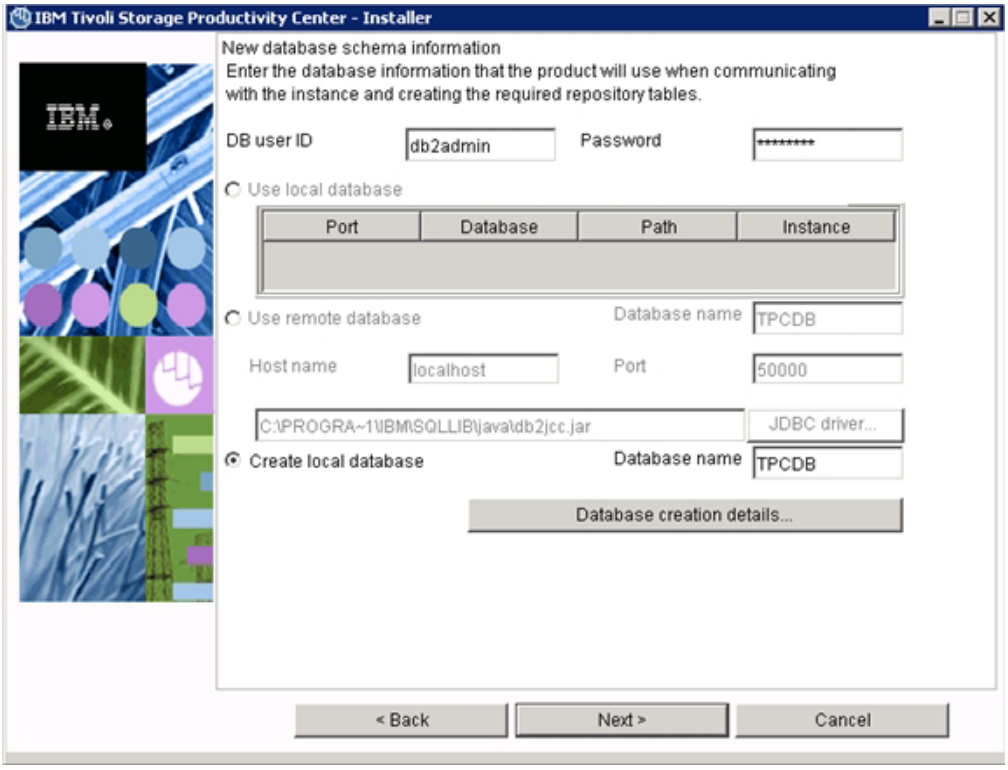


Figure 38. New database schema information panel

If you want to set customized tablespace container locations, click **Database creation details**. The Database schema creation information panel is displayed. Enter the appropriate information and click **OK**.

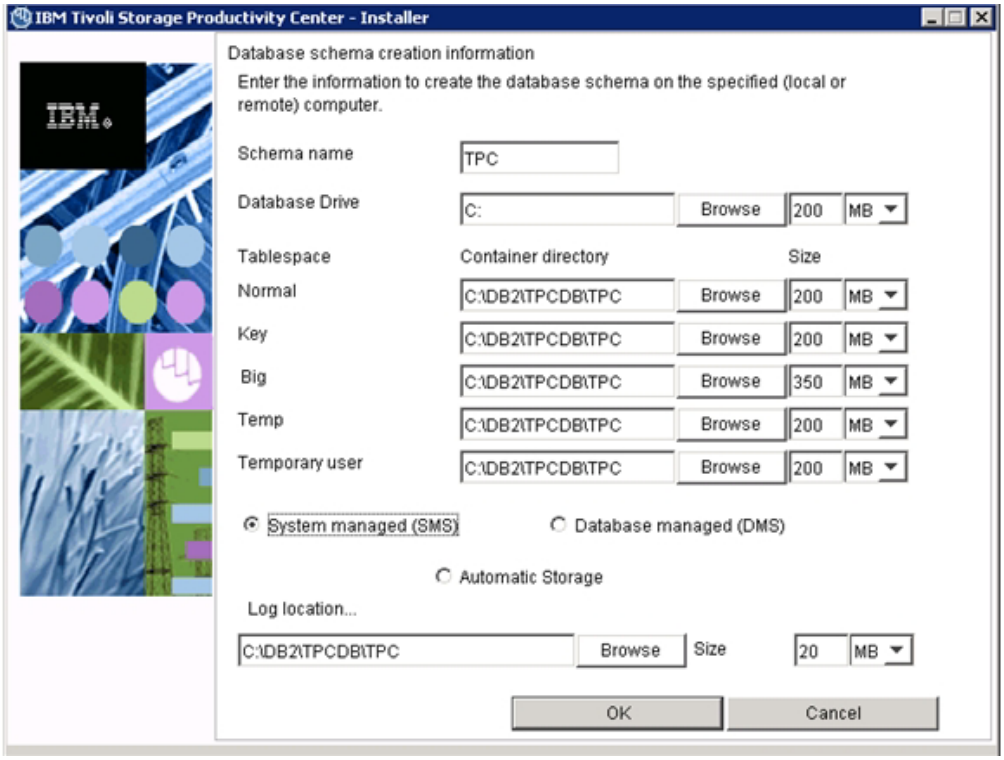


Figure 39. Database schema creation information panel

g. The summary information panel is displayed. Review the information and click **Install**.

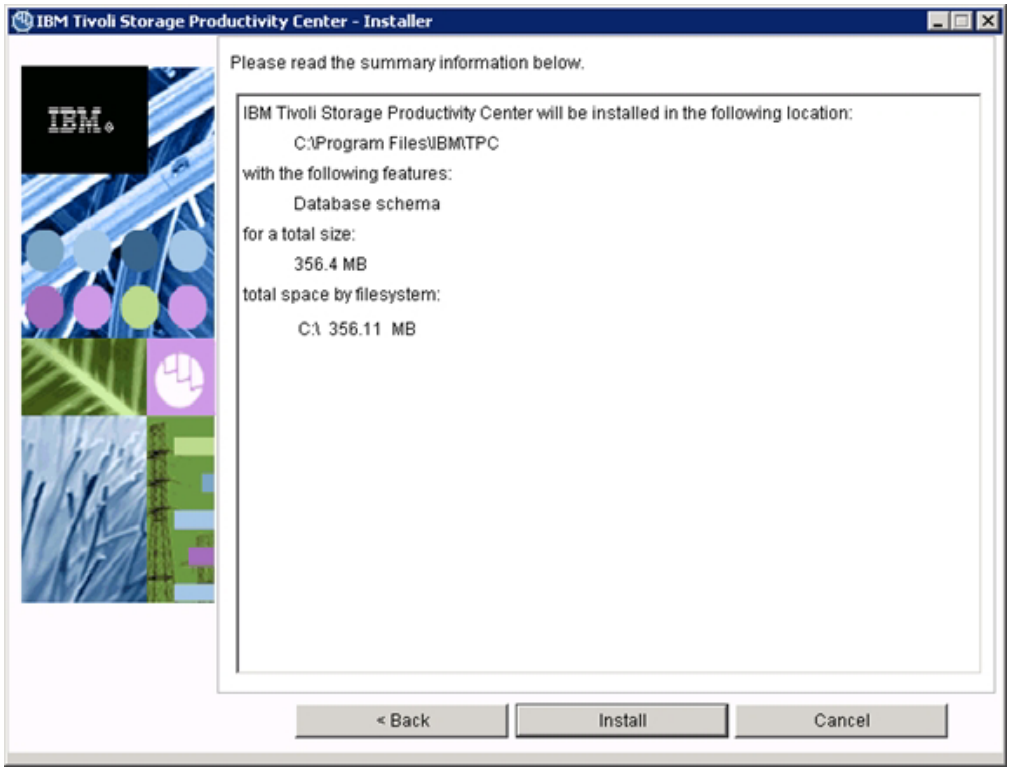


Figure 40. Summary information panel

h. Complete the installation.

Installing Tivoli Storage Productivity Center on Server B:

This procedure provides information about how to install Tivoli Storage Productivity Center on Server B. You will **not** be installing the database schema on this computer.

Note:

- 8 GB of RAM is required.

To install Tivoli Storage Productivity Center using a remote database, follow these steps:

1. Required for AIX: Ensure that the user ID used for installation is part of the AIX system group. By default, any user that is part of the system group can log into the Tivoli Storage Productivity Center GUI.

   Note: If you are creating a user ID for installing Tivoli Storage Productivity Center, exit from root and login with the new user ID. This user ID must be part of the AIX system group.
2. Extract the Tivoli Storage Productivity Center installation package from the download site or use the Tivoli Storage Productivity Center DVD.

   For AIX, use the GNU tar program (version 1.14 or later to extract the package). You can check the version by running the command **tar - -version** from the directory `/opt/freeware/bin/tar`.

3. Ensure that the required JDBC drivers (db2jcc.jar and db2jcc_license_cu.jar) are copied from the remote DB2 server to the local Tivoli Storage Productivity Center computer. The JDBC drivers are located in the Java subdirectory where you installed DB2 on Server A.
4. For AIX or Linux: Do not source the db2profile on Server B.
5. Start the Tivoli Storage Productivity Center installation program on Server B.
6. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.
7. On the Select the type of installation you want to run panel, select **Custom Installation**. Click **Next**.
8. On the Select one or more components to install panel, select **Tivoli Storage Productivity Center Servers** and clear **Create Database Schema**.
9. The Database administrator information panel is displayed. Enter a database administrator user ID and password. Click **Next**.
10. The Existing database schema information panel is displayed.

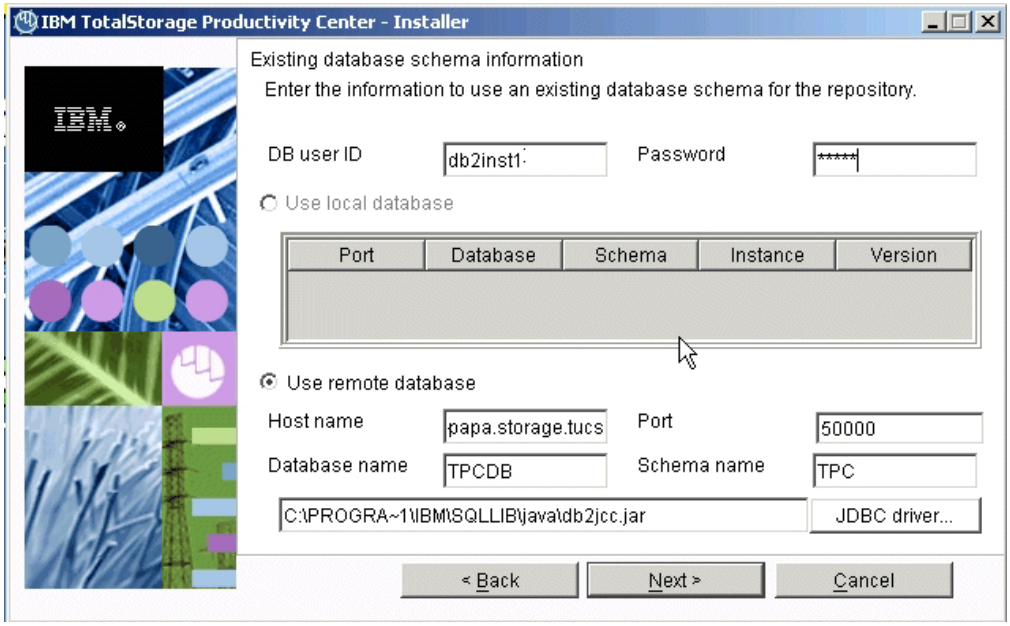


Figure 41. Existing database schema information panel

Select **Use remote database**. Enter the database name, schema name, fully qualified host name or IP address of the remote database server, and its port. Click **JDBC driver** to locate and open the JDBC driver, `db2jcc.jar`, copied from the database host computer. Click **Next**, and the remote database manager message dialog is displayed. Ensure that the remote DB2 instance is running on the database host computer. Click **OK**.

11. The Data server, Device server, and Agent information panel is displayed. Enter the following information:

    Data Server Name
    Enter the fully qualified host name of the Data Server.

    Data Server Port
    Enter the Data Server port. The default is 9549.

Device server Name
Enter the fully qualified host name of the Device server.

Device server Port
Enter the Device server port. The default is 9550.

TPC Superuser
Enter an operating system group name to associate with the TPC Superuser role. This group must exist in your operating system before you install Tivoli Storage Productivity Center. Membership in this group provides full access to the Tivoli Storage Productivity Center product. You can assign a user ID to this group on your operating system and start the Tivoli Storage Productivity Center GUI using this user ID.

Note: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the value you enter for LDAP TPC Administrator group overrides the value you entered here for the TPC Superuser. For more information about security roles, see "Planning for IBM Tivoli Storage Productivity Center authorization" on page 33.

Host authentication password
The password is used by the Storage Resource agent to communicate with the Device server. This password must be specified when you install the Storage Resource agent.

Data Server Account Password
For Windows only. You supply this password. When you supply this password, the Tivoli Storage Productivity Center installation program creates a user ID and the password to create the Data server account on Windows. This user ID and password are used by the Data server service.

The display name for the Data Server is:

```
IBM Tivoli Storage Productivity Center - Data Server
```

WAS admin ID
The WebSphere administration user ID is required by the Device server to communicate with embedded WebSphere. This user ID must be a user with system authority.

Note: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the values you enter for LDAP TPC Administrator user name and password override the value you enter here for WAS Admin ID and password.

Password
Password for the WebSphere administration user ID.

Click **Next**.

If you want to associate operating system groups with Tivoli Storage Productivity Center security roles, click **Security roles**. Enter the operating system group for each Tivoli Storage Productivity Center role you want to make an association with. You do not have to assign security roles at installation time. You can assign these roles after you have installed Tivoli Storage Productivity Center. Click **OK**. Click **Next** on the Data server, Device server, and Storage Resource agent information panel.

Note: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the values you enter here for Security Roles Mapping will be deleted and you must reset the mappings after the installation is complete through the **Role-to-Group Mapping** node.

12. The IBM Tivoli Integrated Portal Information panel is displayed. You can use this panel to select an existing IBM Tivoli Integrated Portal to use or install the IBM Tivoli Integrated Portal program. Enter the following information:

    Specify the location to install TIP
    Accept the default or enter a location to install Tivoli Integrated Portal.

    Port Enter the port number for Tivoli Integrated Portal or accept the default (16310). Tivoli Integrated Portal uses 10 port numbers. For example, if the default port 16310 is accepted, Tivoli Integrated Portal uses the following ports: 16310, 16311, 16312, 16313, 16315, 16316, 16318, 16320, 16322, and 16323. If you enter a port number other than the default number, ensure that you have the following ports available for Tivoli Integrated Portal:

    - *base port*
    - *base port* +1
    - *base port* +2
    - *base port* +3
    - *base port* +5
    - *base port* +6
    - *base port* +8
    - *base port* +10
    - *base port* +12
    - *base port* +13

    where *base port* is the port value entered in the **Port** field.

    Reuse an existing TIP install
    Select this option if you want to use an existing IBM Tivoli Integrated Portal. Under **Existing TIP installs**, select a IBM Tivoli Integrated Portal to use. Enter the following information:

    TIP Administrator ID
    Enter the IBM Tivoli Integrated Portal administrator ID.

    Password
    Enter the IBM Tivoli Integrated Portal password of the administrator ID.

    Click **Next** to continue.

13. The authentication selection panel is displayed. If a valid IBM Tivoli Integrated Portal exists on the system, the authentication method that exists with the program is used. Otherwise, select the authentication method to use.

    Use the single sign-on feature to authenticate a user and gain access to the resources of multiple, trusted applications. To have a single sign-on environment, there must be a centralized authentication repository such as Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory. The LDAP or Active Directory must be accessed by all applications within a single sign-on environment. The user ID and password are passed between applications in an encrypted manner. If you do not want to use LDAP or

Active Directory, you can select OS authentication. Select the authentication method to use and click **Next**.The options are:

LDAP/Active Directory
Use LDAP or Active Directory for the centralized authentication repository. If you select **LDAP/Active Directory**, go to step 11.

OS Authentication
Use operating system authentication. If you select **OS Authentication,** go to step 17 on page 22514.

14. If you selected **LDAP/Active Directory** in step 10, there is more information that Tivoli Storage Productivity Center must collect. The LDAP Server information panel is displayed. Enter this information and click **Next**.

    LDAP Server Hostname
    The fully qualified host and domain name of the system where your LDAP-based directory is running.

    LDAP Port Number
    The port number on your LDAP server where the LDAP process is listening for communications.

    Support referrals to other LDAP servers
    The options are **ignore** or **follow** an LDAP referral. The default is **ignore**.

    Referral is the process by which an LDAP server, instead of returning a result to a client application, returns a reference (a referral) to another LDAP server which might contain further information related to the request for the client application.

    Bind Distinguished Name
    The distinguished name used to bind to the LDAP-based directory when performing a user or group search. This name is optional because some LDAP-based Directories allow anonymous binds and others require you to bind with a specific name and password.

    Bind Password
    The password associated with the Bind Distinguished Name.

15. The LDAP user and group information panel is displayed. Enter this information and click **Next**.

    Relative Distinguished Name for usernames
    Tells IBM Tivoli Storage Productivity Center where in the directory to start a search when performing user authentication.

    Attribute to use for usernames
    Tells IBM Tivoli Storage Productivity Center which attribute in a directory entry that contains the name for the user for authentication.

    Relative Distinguished Name for groups
    Tells IBM Tivoli Storage Productivity Center where in the directory to start a search when performing a group search for authentication or authorization.

    Attribute to use for groups
    Tells IBM Tivoli Storage Productivity Center which attribute in a directory entry that contains the name for the group for authentication or authorization.

16. The LDAP TPC Administrator user and group information panel are displayed. Enter information for these fields:

LDAP TPC Administrator username
The user ID that Tivoli Storage Productivity Center configures as the Tivoli Storage Productivity Center superuser during installation. This user name must exist in the directory before you start the Tivoli Storage Productivity Center installation. The user name must be a member of your designated LDAP Tivoli Storage Productivity Center Administrator group in the directory.

Note: If you install Tivoli Storage Productivity Center on a Windows system and have Tivoli Storage Productivity Center authenticate users against an LDAP-based repository, then the LDAP Tivoli Storage Productivity Center Administrator user name must **not** contain a space or blank character in it. See the WebSphere Application Server APAR PK77578.

Ensure that you use the correct case when entering the values for the LDAP TPC Administrator user name and LDAP TPC Administrator group. If you use the incorrect case when entering these values, then you cannot log into Tivoli Storage Productivity Center for Replication with the user name you expected to use.

LDAP TPC Administrator password
Password for the administrator.

LDAP TPC Administrator group
The group that Tivoli Storage Productivity Center maps to the Tivoli Storage Productivity Center superuser role during installation. This group must exist in the directory before you start the Tivoli Storage Productivity Center installation. The group must have the designated LDAP IBM Tivoli Storage Productivity Center Administrator as a member.

Note:

Ensure that you use the correct case when entering the values for the LDAP TPC Administrator user name and LDAP TPC Administrator group. If you use the incorrect case when entering these values, then you cannot log into Tivoli Storage Productivity Center for Replication with the user name you expected to use.

17. The Summary Information panel is displayed. Review the information. Click **Install**.
18. You see the installing panel. When Tivoli Integrated Portal has finished installing, Tivoli Storage Productivity Center starts the Tivoli Storage Productivity Center for Replication installation.
19. Tivoli Storage Productivity Center starts the IBM Tivoli Storage Productivity Center for Replication installation program.
    a. The Welcome panel is displayed. Click **Next**.
    b. The System prerequisites check panel is displayed. The installation wizard checks whether the prerequisites are installed, then confirms whether your operating system is supported and is at the appropriate fix pack or update level. Click **Next**.
    c. The License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.

d. The Directory Name panel is displayed. Accept the default installation directory by clicking **Next**, or specify a different installation directory and click **Next**.

e. The Tivoli Storage Productivity Center for Replication Administrator ID and Password panel is displayed. Enter the administrator user ID and password. This ID is typically the operating system administrator user ID. If you use a different ID, create it beforehand and ensure that it has administrator rights.

   Note: There is a limitation on the number of characters for the user name and password on AIX. At the time this document was written, AIX did not properly validate passwords that are longer than eight characters.

f. The Default ports panel is displayed. Accept the defaults. Click **Next**.

   The installation program checks for port conflicts on 5110 and 5120. If there is a conflict, a new port is assigned by increasing the port number, and then retesting for a port conflict until an available port is found. The default port numbers in the fields are set to the available ports when the panel is displayed.

   If you assign a port for the Default Client Port or Default Standby Server Port and there is a port conflict, an error message is displayed to let you know of the conflict.

   The client port must be set to the same number on both management servers (the active and the standby servers). The standby server port must also be set to the same number on both management servers (the active and the standby servers).

g. The settings panel is displayed. Review the settings and change them if needed by clicking **Back**. Otherwise, click **Install** to begin the installation.

h. You see the status panel. Wait for installation to complete.

i. You see the summary information panel. Review the information and click **Finish**.

j. Tivoli Storage Productivity Center for Replication is installed with a basic license. If you plan to install the Two Site or Three Site Business Continuity (BC) license, see "Installing the Two Site or Three Site Business Continuity license" on page 184.

20. The successfully installed panel is displayed. Click **Finish**.
21. To access the Tivoli Storage Productivity Center user interface, see "Starting the Tivoli Storage Productivity Center GUI using a web browser" on page 347. To access the Tivoli Storage Productivity Center for Replication user interface, go to the Tivoli Storage Productivity Center information center. Search for **Accessing the IBM Tivoli Storage Productivity Center for Replication GUI**.
22. After you have installed Tivoli Storage Productivity Center, you can use the Configure Devices wizard to configure your system. For information about the Configure Devices wizard, see "Configure Devices" on page 266.

Installing the IBM Tivoli Storage Productivity Center GUI

This procedure provides information about how to install the GUI. The GUI is installed where ever you might want to control IBM Tivoli Storage Productivity Center.

Note: If you are installing the GUI on Windows Vista, you must use the Windows setting "Windows Classic." To change your Windows setting, follow these steps:

1. Right-click your desktop.
2. Click **Personalize**.

3. The Appearance Settings window is displayed.
4. Under **Color scheme**, select **Windows Classic**.
5. Click **OK**.

To install the GUI, follow these steps:

1. Start the IBM Tivoli Storage Productivity Center installation program.
2. On the Select the type of installation panel, select **Custom installation**.
3. Go to the Select one or more components panel, select **GUI**. Click **Next**.
4. The Summary Information panel is displayed. Review the information. Click **Install**.
5. You see the installing panel. Wait for installation to complete.
6. The successfully installed panel is displayed. Click **Finish**.

Installing the IBM Tivoli Storage Productivity Center CLI

This procedure provides information about how to install the CLI.

To install the IBM Tivoli Storage Productivity Center CLI, complete the following steps:

1. Start the IBM Tivoli Storage Productivity Center installation program.
2. On the Select the type of installation panel, select **Custom installation**.
3. Go to the Select one or more components panel, select **CLI**. Click **Next**.
4. The Summary Information panel is displayed. Review the information. Click **Install**.
5. You see the installing panel. Wait for installation to complete.
6. The successfully installed panel is displayed. Click **Finish**.

Adding an installation license

If you have installed a package on your system, for example, IBM Tivoli Storage Productivity Center Basic Edition, and want to add another package to your system (for example, IBM Tivoli Storage Productivity Center Standard Edition), all you do is add the installation license on your system. This topic provides information about how to add the installation license.

To check the type of license you have, look for `.SYS2` files in the installation directory:

TPCBE0402.SYS2
: Indicates you have a Tivoli Storage Productivity Center Basic Edition license.

TPCSE0402.SYS2
: Indicates you have a IBM Tivoli Storage Productivity Center Standard Edition license.

TPCDATA0402.SYS2
: Indicates you have a IBM Tivoli Storage Productivity Center for Data license.

TPCDISK0402.SYS2
: Indicates you have a IBM Tivoli Storage Productivity Center for Disk license.

TPCSLCT0402.SYS2
: Indicates you have a IBM Tivoli Storage Productivity Center Select license.

TPCMRE0402.SYS2
Indicates you have a IBM Tivoli Storage Productivity Center for Disk Select license.

Note: was previously called IBM Tivoli Storage Productivity Center for Disk Midrange Edition.

For example, if you want to add the IBM Tivoli Storage Productivity Center Standard Edition package to your system, complete the following steps:

1. If you are adding a license on AIX or Linux, source the `db2profile`.
2. Start the Tivoli Storage Productivity Center installation program from the DVD or image for the license you want to add.
3. The Select a language panel is displayed. Select a language from the list box and click **OK**. The language you select is the language that is used to install the license.
4. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.
5. The Select the type of installation panel is displayed. Select **Installation Licenses**. Click **Next**.
6. The Database administrator information panel is displayed. Enter the database administrator user ID and password and click **Next**.
7. The Existing database schema information panel is displayed. Enter the information for where database is located. Click **Next**.
8. The successfully installed panel is displayed. Click **Finish**.
9. Stop the Tivoli Storage Productivity Center GUI.
10. Stop the Device server and Data Server.
11. Start the Data Server and the Device server.
12. Start the Tivoli Storage Productivity Center GUI.

Installing the agents

This section describes how to install the Storage Resource agents.

You can install the Storage Resource agent using the installation wizard, the Tivoli Storage Productivity Center user interface, or a command. For information about deploying Storage Resource agents, see “Storage Resource agent deployments” on page 308.

Before installing the Storage Resource agent, see “Deployment considerations for Storage Resource agents” on page 312.

If you are using IPv6 systems, see “Planning for Internet Protocol Version 6” on page 51.

For information about deploying Storage Resource agents from the Tivoli Storage Productivity Center user interface, go to the information center. Search for **Adding Storage Resource agents**.

If you receive an error message with a return code when installing Storage Resource agents, see the Tivoli Storage Productivity Center Information Center. Search for **Return codes used by Storage Resource agent**.

Note:

- If you are installing or deploying the Storage Resource agent on a system, you must first disable the firewall on that system before installing or deploying the agent.
- If the default port is not used for the agent port, then the Tivoli Storage Productivity Center agents need to be installed using custom installation.
- If the default ports were not used when the Tivoli Storage Productivity Center servers were installed, then the agents need to be installed using custom installation.
- **For PowerHA (previously known as HACMP):** A Storage Resource agent must be installed on each node of the cluster and all agents in a cluster must be configured to use the same listening port. To perform a scan on a cluster resource group, you must configure the cluster resource group to have at least one IP address that is accessible from the Tivoli Storage Productivity Center server.

To install the Storage Resource agents, you must log in as a user with the following authority:

On Windows systems
You must have Administrator authority and the following rights:

- Act as part of the operating system
- Log on as a service

On UNIX or Linux systems
You must be logged in as the root user.

Preparing to install Storage Resource agents

This section provides information about how to prepare for installing the Storage Resource agents.

General information

For information about deploying Storage Resource agents, see "Deployment considerations for Storage Resource agents" on page 312.

Installing the Storage Resource agent using a command

This section describes how to install the Storage Resource agent using a command.

You would typically install the Storage Resource agents using the Tivoli Storage Productivity Center GUI. However, if you need to install the Storage Resource agents locally, you can do so with limited support.

For example, if you use this method of installation, you get a return code of zero for a successful installation and a nonzero return code for an unsuccessful installation. If you have an unsuccessful installation, you need to go through the log files to determine what the problem was for the failure.

The **-force** option can be used when you have a Storage Resource agent pointing to multiple servers. If one server has installed an agent, and another server wants to install an agent in the same location, the second server can use the **-force** option to install the agent.

If you use this method of installation, make sure that when you specify a directory to install the Storage Resource agent into, do not specify an ending slash mark (\). For example, do not specify `C:\agent1\`. This action causes the installation to fail.

If you run the agent as a non-daemon service (On-Demand service), you must make sure that at least one protocol is valid for a successful connection from the server to the agent. See Table 35 for the required parameters for each protocol.

Table 35. Parameters required for each protocol

Protocol	Description
SSH	Requires the user ID and password or user ID, certificate, and passphrase.
Windows (SMB)	Requires the user ID and password.
REXEC	Requires user ID and password.
RSH	Requires the user ID.

The Storage Resource Agent image contains the installation images for the Storage Resource agents. The images are located in the following directory:

DVD_installation_image_location/data/sra/*operating_system*

See Table 36 for the Storage Resource agent installation images.

Table 36. Storage Resource agent installation images

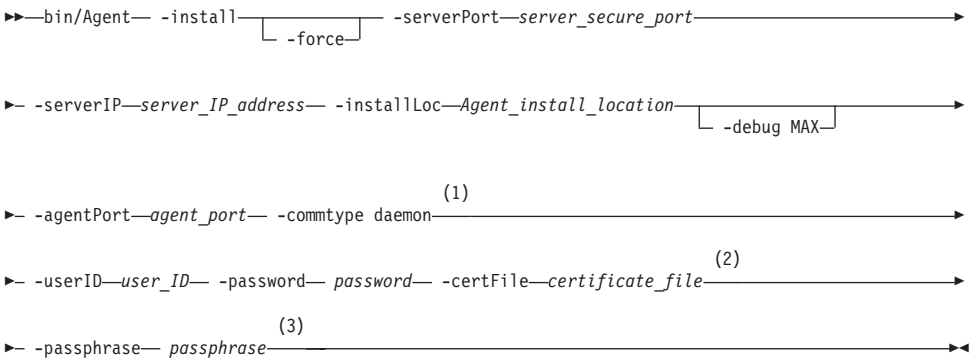
Operating system	Operating system name
AIX	aix_power
HP-UX	hp-ux_itanium
Linux x86	linux_ix86
Linux for Power Systems Servers	linux_power
Linux s390	linux_s390
Oracle Solaris	solaris_sparc
Windows	windows

To install the Storage Resource agents locally, follow these steps:

1. Go to the installation image location:

   ```
   cd DVD_installation_image_location
   ```

2. Run the following command:



Notes:

1 Parameters when the agent is run as a daemon service.

2 Parameters when the agent is run as a non-daemon service. See Table 35 on page 230 to determine which parameter is required for each protocol.

3 Parameters when the agent is run as a non-daemon service. See Table 35 on page 230 to determine which parameter is required for each protocol.

Parameters:

-install
Installs the Storage Resource agent.

-force
This forces the agent to be installed. There are two different situations in which this parameter must be specified:

- If an earlier installation failed and there is residue on the system which causes further installations to fail. You need to make sure that all the parameters provided are valid, such as the installation location, port, and so on.
- If the agent is already installed from one server and you now need to install the agent pointing to another server.

-serverPort *server_secure_port*
The default port for the Data server is 9549.

-serverIP *server_IP_address*
The IP address of the server. If the server can be reached through multiple IP addresses, then multiple IP addresses can be specified with IP addresses separated with a comma.

-installLoc "*Agent_install_location***"**
Location where the agent is installed. Enclose the directory name in quotation marks, for example, "C:\Program Files\IBM\TPC_SRA\".

-debug MAX
Optional parameter for debugging purposes.

-agentPort *agent_port*
If the agent is run as a daemon service, then the agent port must be specified.

-commtype daemon
If the agent is run as a daemon service, then this parameter must be specified.

-userID *user_ID*
For non-daemon service. The user ID defined on the agent system. This user ID is used by the server to connect to the agent system.

-password *password*
For non-daemon service. Password for the user ID.

-certFile "*certificate_file***"**
For non-daemon service. The certificate used for SSH communication between the server and agent. This certificate needs to be stored on the server system. Enclose the directory name in quotation marks, for example, "c:\keys\id_sra".

For information about SSH protocol and SSL protocol certificates, see “Creating a certificate for SSH protocol” on page 317 and “Replacing a custom certificate for SSL protocol” on page 322.

-passphrase *passphrase*
For non-daemon service. The passphrase defined for the certificate used in SSH communication.

If the installation fails, see the return codes in the Information Center. Search for **Return codes used by Storage Resource agent**.

Installing the agent using the installation program

This section describes how to install the Storage Resource agent on the local system using the installation program.

Follow these steps:

1. Start the Tivoli Storage Productivity Center installation program. For information about starting the installation program, see "Starting the installation program" on page 169.
2. The Select a language panel is displayed. Select a language from the list box and click **OK**. This is the language that is used for installing this product.
3. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.
4. The select the type of installation panel is displayed. On the Select the type of installation you want to run panel, select **Typical installation** > **Storage Resource Agent**. Click **Next**.

   Note: You cannot install the Storage Resource agent from the disk 1 image on machines that do not have a Data server installed. If you are installing only a Storage Resource agent, use the Tivoli Storage Productivity Center user interface. You can install the Storage Resource agent if you are also installing the Data server.
5. The User ID and password panel are displayed. Enter the user ID and password that has administrator and database administrator authority.

   Enter the server name, server port, and agent port or accept the defaults. Click **Next**.

   Note: The default server port is 9549.
6. The Summary Information panel is displayed. Review the information. Click **Install**.
7. You see the installing panel. Wait for installation to complete.
8. The Successfully installed panel is displayed. Click **Finish**.

Installing the agent using custom installation

This topic provides information about how to install the Storage Resource agent locally on a server to be managed.

To install the Storage Resource agent, complete the following steps:

1. Start the IBM Tivoli Storage Productivity Center installation program. For information about starting the installation program, see "Starting the installation program" on page 169.
2. The Select a language panel is displayed. Select a language and click **OK**.
3. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select **I accept the terms of the license agreement**. Click **Next**.
4. The Select the type of installation panel is displayed. Select **Custom installation**. Click **Next**.

5. The select one or more components to install panel is displayed. Select **Storage Resource Agent**. Click **Next**.

   Note: You cannot install the Storage Resource agent from the disk 1 image on machines that do not have a Data server installed. If you are installing only a Storage Resource agent, use the Tivoli Storage Productivity Center user interface. You can install the Storage Resource agent if you are also installing the Data server.

6. The Storage Resource Agent Information panel is displayed. Enter the following information:

   Daemon
   The Storage Resource agent is installed as a daemon service. For more information about the daemon service, see "Storage Resource agent protocol support" on page 50.

   Port The Storage Resource agent port. The default port is 9510.

   Windows Service Information (Windows only)
   Click this button to enter an agent service name, user ID, and password information that the installation program uses to create a Windows service for the agent.

   RXA The Storage Resource agent is installed as a non-daemon service. For information about the non-daemon service, see "Storage Resource agent protocol support" on page 50.

   Note: If you are using the RXA access method to install or deploy the Storage Resource agent on a UNIX or Linux system, you must use the root user ID.

   User ID
   The user ID used to connect to the server.

   Password
   The password for the user ID to connect to the server.

   Certificate
   The certificate used to connect to the server. For information about the SSH protocol certificates, see "Creating a certificate for SSH protocol" on page 317.

   Passphrase
   The passphrase used to connect to the server.

7. The Data server, Device server, and Storage Resource Agent Information panel is displayed. Enter the following information:

   Data Server Name
   Enter the fully qualified host name of the Data Server.

   Data Server Port
   Enter the Data Server port. The default is 9549.

   Device Server Name
   Enter the fully qualified host name of the Device server.

   Device Server Port
   Enter the Device server port. The default is 9550.

   TPC Superuser
   Enter an operating system group name to associate with the TPC Superuser role. This group must exist in your operating system before

you install Tivoli Storage Productivity Center. Membership in this group provides full access to the Tivoli Storage Productivity Center product. You can assign a user ID to this group on your operating system and start the Tivoli Storage Productivity Center GUI using this user ID.

Note: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the value you enter for LDAP TPC Administrator group overrides the value you entered here for the TPC Superuser.

Host Authentication Password
The password used for the Storage Resource agent to communicate with the Device server.

Data Server Account Password
For Windows only. You supply this password. When you supply this password, the Tivoli Storage Productivity Center installation program creates a user ID and the password you supply to create the Data server account on Windows. This user ID and password are used by the Data server service.

The display name for the Data Server is:

```
IBM Tivoli Storage Productivity Center - Data Server
```

WAS admin ID and Password
The user ID and password required by the Device server to communicate with embedded WebSphere. This is only used at installation time.

Note: If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the values you enter for LDAP TPC Administrator user name and password override the value you entered here for WAS Admin ID and password.

Click **Next**.

If you click the Security roles button:

The Advanced security roles mapping panel is displayed. Enter the operating system group for each Tivoli Storage Productivity Center role you want to make an association with. The operating group must exist before you can associate a Tivoli Storage Productivity Center role with a group. Click **OK**. Click **Next**.

Note: You do not have to assign security roles at installation time. You can assign these roles after you have installed Tivoli Storage Productivity Center.

If you select LDAP authentication later in the Tivoli Storage Productivity Center installation, then the values you enter here for Security Roles Mapping are deleted and you will have to assign roles to individuals after the installation is complete (use the **Role-to-Group Mapping** node).

If you click the NAS discovery button:

The NAS discovery information panel is displayed. Enter the NAS filer login default user name and password. If you want to enter the SNMP communities to be used for NAS discovery, click **Add**, enter the community name, then click **OK**. Click **Next** on the Data server, Device server, and Storage Resource agent information panel.

8. The Summary Information panel is displayed. Review the information. Click **Install**.
9. You see the installing panel. Wait for installation to complete.
10. The successfully installed panel is displayed. Click **Finish**.

Installing the Storage Resource agent on the Virtual I/O Server

Use the Storage Resource agent to gather information about the Virtual I/O Servers.

Local installation

To install the Storage Resource agent on the Virtual I/O Server, follow these steps:

1. Log in to the Virtual I/O Server using the padmin user ID.
2. Set up the AIX environment. Run the following command:
   ```
   oem_setup_env
   ```
3. From the directory where the Storage Resource agent installation image is located, run the Storage Resource agent command.

   Note: You can either mount the Storage Resource agent installation DVD or FTP to the Storage Resource agent installation image. Mount the installation image in the padmin home directory.
4. Run the Storage Resource agent installation command. For information about the Storage Resource agent installation command, see "Installing the Storage Resource agent using a command" on page 229.

Deploying the Storage Resource agent remotely

You can deploy the Storage Resource agent remotely for the Virtual I/O Server from your local workstation.

Remote deployment

To deploy the Storage Resource agent remotely, follow these steps:

1. Enable remote access for the **root** user ID. Follow these steps:
   a. Connect to the Virtual I/O Server using ssh/telnet with the padmin user ID.
   b. Run the following command to set up the AIX environment:
      ```
      oem_setup_env
      ```
   c. Run the following command:
      ```
      smitty login_user
      ```
   d. In the **User NAME** field, type **root** and press **Enter**.
   e. In the **User can LOGIN REMOTELY(rsh,tn,rlogin)?** field, set the value to **true**.
2. Save the changes.
3. Optional: Create an SSH certificate for the **root** user ID. For information about creating the SSH certificate for the **root** user, see "Creating an SSH certificate for the root user ID" on page 376.
4. Start the remote deployment.

Installing the Tivoli Storage Productivity Center Monitoring Agent

The 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. The Monitoring Agent uses an existing Tivoli Storage Productivity Center server in the environment to obtain information.

You will be installing the Tivoli Storage Productivity Center Monitoring Agent in the Tivoli Storage Productivity Center environment. You will also be installing the agent-specific support files needed by the Tivoli Enterprise Monitoring Server.

Installing and configuring the Tivoli Storage Productivity Center Monitoring Agent on Windows:

This topic describes how to install the Tivoli Storage Productivity Center Monitoring agent on Windows.

Installing the Monitoring Agent in the Tivoli Storage Productivity Center environment

This procedure describes how to install the Monitoring Agent in the Tivoli Storage Productivity Center environment. You can install the agent on the Tivoli Storage Productivity Center server or on another computer. This example assumes that the Tivoli Monitoring environment is installed and running, This example also assumes you are installing the agent on a separate computer from the Tivoli Storage Productivity Center server.

To install the Tivoli Storage Productivity Center Monitoring Agent on Windows, complete the following steps:

1. Gather information required by the Tivoli Storage Productivity Center Monitoring agent installation wizard. You will need to know the host name and authentication password of the Tivoli Storage Productivity Center server.
2. Log on to the system with administrator authority on Windows.
3. If you are using the installation DVD, the files are located in the itm folder. If you are using the installation image, use the Disk 1 part 4 image.
4. Go to the installation directory and run **setup.exe**.
5. The Welcome panel is displayed. Click **Next**.
6. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms, select **I accept the terms in the license agreement** and click **Next**. If you do not agree with the terms and click **Cancel**, the program will not be installed.
7. The Choose Destination Location panel is displayed. Accept the default location or select a location by clicking **Browse**. Click **Next**.
8. The User Data Encryption Key panel is displayed. The Tivoli Enterprise Monitoring server uses a data encryption key for Secure Socket (SSL) communication between the agent and server. Enter a data encryption key for encrypted communication. Write down the key so that you can configure the Tivoli Monitoring Enterprise server with the encryption key. Click **Next** when you have specified the key.

   Note: For information about the data encryption key, see http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/index.jsp?topic=/com.ibm.itcamwas_wr.doc_6.2/KHT62IG39.htm. See **Step 4: Define the node's security environment**.
9. The Select Program Folder panel is displayed. Enter a Program Folder name or select an exiting folder from the Existing Folders list. Click **Next**.

10. The Start Copying Files panel is displayed. Review your current settings. When you are ready to proceed, click **Next**. Wait for the installation to complete.
11. The Tivoli Storage Productivity Center Agent Configuration panel is displayed. In the navigation tree, you see these configuration steps:
    - Log Configuration
    - TPC Configuration
    - Server and Data Sources Node (optional)
    - Computers and Hypervisors Node (optional)
    - Storage Systems Node (optional)
    - Fabrics and Switches Node (optional)

    The last four nodes are optional. Specify information for these nodes if you want Tivoli Storage Productivity Center to collect information about those nodes.
12. In the step for **Log Configuration**, enter this information:

    Log File
    Directory and file name where the Monitoring Agent log file is created. The default is `$ITM_INSTALL_DIR/tmaitm6/logs/TPCAgent.log.`.

    Log Level
    Specify the debugging level for the log file.

    Click **Next**.
13. In the step for **TPC Configuration**, enter this information:

    Address
    IP address or host name of the Device server.

    TPC Device Server Port
    Port for connecting to the Device server. The default port is 9550.

    TPC Host Authentication Password
    Password to use to connect to the Device server.

    Confirm TPC Host Authentication Password
    Enter the password again.

    Click **Next**.
14. Configuring the node names is optional. You can enter any node name.
    a. For **Server and Data Sources Node**, click **New** to add a new data node. Enter a node name for this node. This is an internal Tivoli Storage Productivity Center node name and will not be displayed in the Tivoli Storage Productivity Center user interface. Click **Next**.
    b. For **Computers and Hypervisors Node**, click **New** to add a new data node. Enter a node name for this node. This is an internal Tivoli Storage Productivity Center node name and will not be displayed in the Tivoli Storage Productivity Center user interface. Click **Next**.
    c. For **Storage Systems Node**, click **New** to add a new data node. Enter a node name for this node. This is an internal Tivoli Storage Productivity Center node name and will not be displayed in the Tivoli Storage Productivity Center user interface. Click **Next**.
    d. For **Fabrics and Switches Node**, click **New** to add a new data node. Enter a node name for this node. This is an internal Tivoli Storage Productivity

Center node name and will not be displayed in the Tivoli Storage Productivity Center user interface. Click **OK**.

Note: If you configure more than one Monitoring Agent to collect information for these optional nodes, use unique node names for the different Monitoring Agents. Using identical names for the same node on more than one agent can result in Tivoli Monitoring Server displaying information for only the nodes on the most recently refreshed Monitoring Agent. For example, do not use the same node name for **Server and Data Sources Node** on more than one Monitoring Agent.

15. If the TEMS connection panel is not displayed, then open the Manage Tivoli Enterprise Monitoring Services window. Select **Monitoring Agent for TPC**. Right-click **Monitoring Agent for TPC** and click **Reconfigure**.
16. The Configure Monitoring Agent for TPC panel is displayed. For the TEMS Connection panel, enter information for the protocol. Enter this information for the Protocol 1 tab:

    TEMS Hostname
    Host name of the Tivoli Enterprise Monitoring Server.

    Port Number
    The default port is 1918.

    Protocol
    Enter the protocol type. For IP.PIPE settings, enter the port number.

    For information about the TEMS connection parameters, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.omegamon.mes.doc/zosconfigguide601263.htm.
17. Your configuration is now complete.
18. To enable the Eclipse Help Server for Tivoli Monitoring Services, configure the Eclipse Help Server. Specify port 80 and restart TEMS and TEPS if not done automatically.

Installing the Monitoring Agent support files in the IBM Tivoli Monitoring Server environment

Tivoli Storage Productivity Center has support files that are installed on each of these Tivoli Monitoring components:

- Tivoli Enterprise Monitoring Server (TEMS)
- Tivoli Enterprise Portal Server (TEPS)
- Tivoli Enterprise Portal Desktop Client (TEPDC)

Follow this procedure to install the Monitoring Agent in the IBM Tivoli Monitoring Server environment:

1. Log on to the IBM Tivoli Monitoring system with administrator authority on Windows.
2. If you are using the installation DVD, the files are located in the itm folder. If you are using the installation image, use the Disk 1 part 4 image.
3. Go to the installation directory and run **setup.exe**.
4. The Welcome panel is displayed. Click **Next**.
5. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms, select **I accept the terms in the**

license agreement and click **Next**. If you do not agree with the terms and click **Cancel**, the program will not be installed.

6. The Select Features panel is displayed. Select the features you want to install.

   Expand **Tivoli Enterprise Monitoring Agents - TEMA**. Make sure that you expand the node and check all boxes. Click **Next**.

```
Tivoli Enterprise Monitoring Server - TEMS
     TPC
Tivoli Enterprise Portal Server - TEPS
     TPC
TEP Desktop Client - TEPD
     TPC
```

7. The IBM Tivoli Productivity Center Monitoring - InstallShield Wizard panel is displayed. Select all the components related to Tivoli Storage Productivity Center.

   Note: This panel is displayed when the Monitoring Agent is installed on the Tivoli Monitoring Server.

```
Tivoli Enterprise Monitoring Agents - TEMA
     Tivoli Enterprise Monitoring Agent Framework
     TPC
     Tivoli Enterprise Services User Intrface Extensions
Tivoli Enterprise Monitoring Server - TEMS
     TPC
Tivoli Enterprise Portal Server - TEPS
     TPC
TEP Desktop Client - TEPD
     TPC
```

   Make sure that you expand all nodes and check all boxes. Click **Next**.

8. In the next panel, select the appropriate program folder for the agent to be installed in. Click **Next**.
9. The summary panel is displayed. Review the summary information and click **Next** to start the installation.
10. After installation is complete, the Configuration Defaults for Connecting to a TEMS panel is displayed. This panel lets you enter connection information to the Tivoli Monitoring server and the Tivoli Storage Productivity Center Device server.

    Note: This panel is displayed when the Monitoring Agent is installed on the Tivoli Monitoring server.

    Enter this information for the Protocol 1 tab:

    TEMS Hostname
    Host name of the Tivoli Enterprise Monitoring Server.

    Protocol
    Enter the protocol type. For IP.PIPE settings, enter the port number.

    For information about the TEMS connection parameters, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.omegamon.mes.doc/zosconfigguide601263.htm.

    When you are finished, click **OK**.

11. Your configuration is now complete. To see the reports generated, see the Tivoli Storage Productivity Center Information Center. Search for **monitoring agent reports**.

Installing and configuring the Tivoli Storage Productivity Center Monitoring Agent on AIX or Linux:

This topic describes how to install the Tivoli Storage Productivity Center Monitoring Agent on AIX or Linux.

Installing the Monitoring Agent and support files in the Tivoli Storage Productivity Center environment

This example assumes that the Tivoli Monitoring environment is installed and running and that you are installing the agent in the Tivoli Storage Productivity Center environment.

Tivoli Storage Productivity Center install support files on each of these Tivoli Monitoring components:

- Tivoli Enterprise Monitoring Server (TEMS)
- Tivoli Enterprise Portal Server (TEPS)
- Tivoli Enterprise Portal Desktop Client (TEPDC)

Note:

- IBM Tivoli Monitor does not support the Tivoli Enterprise Portal Desktop Client (TEPDC) on AIX.
- If you have the Tivoli Enterprise Portal Browser Client installed, you must also install the Monitoring Agent support files on the Browser Client.

To install the Tivoli Storage Productivity Center Monitoring agent on AIX or Linux, complete the following steps:

1. Gather information required by the Tivoli Storage Productivity Center Monitoring agent installation wizard. You will need to know the host name and IP address of the Tivoli Enterprise Monitoring server.
2. Log on to the system with the root user ID or a user ID that has root access rights.
3. If you are using the installation DVD, the files are located in the itm folder. If you are using the installation image, use the Disk 1 part 4 image.
4. Go to the installation directory and run **./install.sh**.
5. Here is an example of the prompts and responses you will see. This example assumes that you do not have a directory and are creating a directory and are installing the agent on a local host.

```
UPGRADE

Enter the name of the IBM Tivoli Monitoring directory
[ default = /opt/IBM/ITM ]:

"/opt/IBM/ITM" does not exist
Try to create it [ 1-yes, 2-no; "1" is default ]? 1

Select one of the following:

1) Install products to the local host.
2) Install products to depot for remote deployment (requires TEMS).
3) Install TEMS support for remote seeding
4) Exit install.

Please enter a valid number: 1
```

6. This example is a continuation from the previous step and lets you read the license agreement.

```
International Program License Agreement

...
...
```

   Read the license agreement and make a selection. Press **Enter**. If you do not accept the license agreement, the program will end and you will not install the program.

7. This example is a continuation from the previous step and lets you enter an encryption key for a secure connection between the monitoring agent and the Tivoli Enterprise Monitoring Server. You can use the default key or enter a new key. Write down and save the key so that you can configure the Tivoli Enterprise Monitoring Server later. Press **Enter**.

```
Enter a 32-character encryption key, or just press Enter to use the default
              Default = IBMTivoliMonitoringEncryptionKey

GSkit encryptin key has been set.
Key File directory: /opt/IBM/ITM/keyfiles
```

   Enter the encryption key and press **Enter**.

   Note: For information about the encryption key, see http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/index.jsp?topic=/com.ibm.itcamwas_wr.doc_6.2/KHT62IG66.htm. See **Step 6: Provide an encryption key**.

8. This example is a continuation from the previous step and lets you select the product packages to install. Make a selection and press **Enter**. For an agent only installation, select 1.

```
Product packages are available for this operating system and component support
categories:

1) IBM Tivoli Monitoring components for this operating system
2) Tivoli Enterprise Portal Browser Client support
3) Tivoli Enterprise Portal Desktop Client support
4) Tivoli Enterprise Portal Server support
5) Tivoli Enterprise Monitoring Server support
6) Other operating systems
```

9. This example is a continuation from the previous step and lets you select the product to install. Follow these steps:
   a. Enter **3** (for Tivoli Enterprise Portal Desktop Client support) on the Product packages panel and press **Enter**.

```
The following products are available for installation:

1) Monitoring Agent for TPC  V06.22.00.00
2) All of the above
```

      Select **2** and press **Enter**. Wait for this step to complete.

   b. Enter **4** (for Tivoli Enterprise Portal Server support) on the Product packages panel and press **Enter**.

```
The following products are available for installation:

1) Monitoring Agent for TPC  V06.22.00.00
2) All of the above
```

Select **2** and press **Enter**. Wait for this step to complete.

c. Enter **5** (for Tivoli Enterprise Monitoring Server support) on the Product packages panel and press **Enter**.

```
The following products are available for installation:

1) Monitoring Agent for TPC  V06.22.00.00
2) All of the above
```

Select **2** and press **Enter**. Wait for this step to complete.

10. You see this prompt. Enter **2** and press **Enter**.

```
Do you want to install additional products or product support packages?
```

You see the following message:

```
Installation step complete.
```

11. You see this prompt.

```
You may now configure any locally installed IBM Tivoli Monitoring product
via the "/opt/IBM/ITM/bin/itmcmd config" command.
```

Go to the following directory:

```
cd /opt/IBM/ITM/bin
```

Run this command:

```
./itmcmd manage &
```

12. The Manage Tivoli Enterprise Monitoring Services window is displayed. Select **Monitoring Agent for TPC**. Right-click **Monitoring Agent for TPC** and click **Configure**.
13. The Tivoli Storage Productivity Center Agent Configuration panel is displayed. In the navigation tree, you see these configuration steps:
    - Log Configuration
    - TPC Configuration
    - Server and Data Sources Node (optional)
    - Computers and Hypervisors Node (optional)
    - Storage Systems Node (optional)
    - Fabrics and Switches Node (optional)

    The last four nodes are optional. Specify information for these nodes if you want Tivoli Storage Productivity Center to collect information about those nodes.
14. In the step for **Log Configuration**, enter this information:

    Log File
    Directory and file name where the Monitoring Agent log file is created. The default is `$ITM_INSTALL/li6263/p1/bin/logs/TPCAgent.log`.

Log Level
Specify the debugging level for the log file.

Click **Next**.

15. In the step for **TPC Configuration**, enter this information:

    Address
    IP address or host name of the Device server.

    TPC Device Server Port
    Port for connecting to the Device server.

    TPC Host Authentication Password
    Password to use to connect to the Device server.

    Confirm TPC Host Authentication Password
    Enter the password again.

    Click **Next**.

16. Configuring the node names is optional. You can enter any node name.
    a. For **Server and Data Sources Node**, click **New** to add a new data node. Enter a node name for this node. This is an internal Tivoli Storage Productivity Center node name and will not be displayed in the Tivoli Storage Productivity Center user interface. Click **Next**.
    b. For **Computers and Hypervisors Node**, click **New** to add a new data node. Enter a node name for this node. This is an internal Tivoli Storage Productivity Center node name and will not be displayed in the Tivoli Storage Productivity Center user interface. Click **Next**.
    c. For **Storage Systems Node**, click **New** to add a new data node. Enter a node name for this node. This is an internal Tivoli Storage Productivity Center node name and will not be displayed in the Tivoli Storage Productivity Center user interface. Click **Next**.
    d. For **Fabrics and Switches Node**, click **New** to add a new data node. Enter a node name for this node. This is an internal Tivoli Storage Productivity Center node name and will not be displayed in the Tivoli Storage Productivity Center user interface. Click **OK**.

    Note: If you configure more than one Monitoring Agent to collect information for these optional nodes, use unique node names for the different Monitoring Agents. Using identical names for the same node on more than one agent can result in Tivoli Monitoring Server displaying information for only the nodes on the most recently refreshed Monitoring Agent. For example, do not use the same node name for **Server and Data Sources Node** on more than one Monitoring Agent.

17. Open the Manage Tivoli Enterprise Monitoring Services window. Select **Monitoring Agent for TPC**. Right-click **Monitoring Agent for TPC** and click **Configure**.
18. The Configure Monitoring Agent for TPC panel is displayed. For the TEMS Connection panel, enter information for the protocol. Enter this information for the Protocol 1 tab:

    TEMS Hostname
    Host name of the Tivoli Enterprise Monitoring Server.

    Protocol
    Enter the protocol type. For IP.PIPE settings, enter the port number.

For information about the TEMS connection parameters, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.omegamon.mes.doc/zosconfigguide601263.htm.

19. Your configuration is now complete.

    Note: You might have to reconfigure the **Tivoli Enterprise Portal Server**, **Tivoli Enterprise Monitoring Server**, and **Tivoli Enterprise Portal Desktop Client**. To reconfigure these components, follow these steps:

    a. Open the Manage Tivoli Enterprise Monitoring Services console by running the command:

       ```
       itmcmd manage
       ```

    b. Right-click **Tivoli Enterprise Portal Server** and click **Rebuild Configuration**.

    c. Right-click **Tivoli Enterprise Monitoring Server** and click **Rebuild Configuration**.

    d. Right-click **Tivoli Enterprise Portal Desktop Client** and click **Rebuild Configuration**.

20. To enable the Eclipse Help Server for Tivoli Monitoring Services, open Manage Tivoli Monitoring Services and configure the Eclipse Help Server. Specify port 80 to be used and restart TEMS and TEPS if not done automatically.

Installing the Language Pack

The Language Pack allows you to see messages, online help, and text in a different language other than English in Tivoli Storage Productivity Center.

Before you install the Language Pack, you must first have successfully installed Tivoli Storage Productivity Center. You can install the Language Pack for the components in the following list:

- Data server
- Device server
- GUI
- CLI

You can install the Language Pack for the languages in the following table:

Czech	Hungarian	Russian
German	Italian	Brazilian Portuguese
Spanish	Japanese	Simplified Chinese
French	Korean	Traditional Chinese
Polish		

You must have administrator authority on Windows or root authority on UNIX or Linux to install the Language Pack.

Installing the Language Pack

To install the Language Pack, complete the following steps:

1. Log on to the system with administrator authority on Windows or root authority on UNIX or Linux.

2. Insert the Language Pack CD or go to the directory where you have downloaded the Language Pack.

   Run the following program for Windows: **TPCLP.exe**.

   Run the following program for UNIX or Linux: **./setup.sh**. Run the db2profile before running TPCLP.

   Install the language packs for all the components (servers, agents, GUI, and so forth).

   Install Tivoli Storage Productivity Center for Replication into a directory other than the National Language Version (NLV) directory because the installation might fail in some languages.
3. For the Storage Resource agent deployed on the server system, follow the instructions in "Installing the Language Pack for the Storage Resource agents."

   Note: Applying the Language Pack for the local Storage Resource agent is not available using the Language Pack installation program.
4. The Select a Language panel is displayed. Select a language from the list box and click **OK**.
5. The Welcome panel is displayed. Click **Next**.
6. The Select one or more languages panel is displayed. Select one or more languages and one or more Tivoli Storage Productivity Center components to install. Click **Next**.
7. The Summary panel is displayed indicating the choices you have made on the previous panel. Review the information. Click **Install**.
8. You see the Installing status panel displayed. Wait for installation to complete.
9. The Installation successful panel is displayed. Click **Finish**.

For the changes to take effect, you must stop and restart the Device server, Data Server, and GUI (or whatever component you installed the Language Pack for). For information about stopping and restarting the components, go to the Tivoli Storage Productivity Center information center. Search for **Starting and stopping the Tivoli Storage Productivity Center services**.

Installing the Language Pack for the Storage Resource agents

To install the Storage Resource agents in a different language, follow these steps:

1. Install the language packs for all the components on the Tivoli Storage Productivity Center server. See the procedure for installing the language pack above.
2. Deploy a Storage Resource agent on your agent system. Install the agent as a daemon or non-daemon.
3. From the Tivoli Storage Productivity Center GUI, create an upgrade job. Click **Administrative Services** > **Configuration**. Right click **Data/Storage Resource Agent Upgrades**. Click **Create Data/Storage Resource Agent Upgrade**. Also check the box for **Force upgrade on agents already at this release level**.
4. In the content pane, click the Options tab. Click **Maintenance**. Select the languages for the Language Packs you want. Save and run the job. Wait until the job completes.
5. After the job completes, verify that the Storage Resource agent has been upgraded. Click **Administrative Services** > **Data Sources** > **Data/Storage Resource Agents**. Refresh the panel to verify that the agent upgrade job has completed. The Status column should indicate "Up" after completing the job.

6. After you have upgraded the Storage Resource agent with the language you want, verify that the Storage Resource agent language pack has been deployed on the agent system. See Table 40 on page 248 for the files that should be deployed.

Configuring the environment after installing the Language Pack

After you have installed the Tivoli Storage Productivity Center Language Pack, you need to set the regional and language settings for the environment.

Here is an example of configuring the environment for AIX and changing the language from English to German (de_DE).

1. Install the language packages for the language you are changing to.
2. Run the following commands:

```
chlang -m de_DE de_DE
export LANG=de_DE
```

3. Restart all Tivoli Storage Productivity Center services or reboot the system.

Here is an example of configuring the environment for Windows 2008.

1. Install the language packages for the language you are changing to.
2. Open the Windows Control Panel window.
3. From the control panel, select **Regional and Language Options**. Make the required changes on the Regional and Language Options window. Click **OK**.
4. Restart all Tivoli Storage Productivity Center services or reboot the system.

Location of Language Pack files

To check for the files that have been installed for the Language Pack, see Table 37.

Table 37. Language Pack files

Component	Directory	Language Pack files
Data server	*TPC_install_dir*\data\server\lib	TSRMsrv_*locale*.zip where *locale* is the abbreviation for the language you installed deviceSrv_*locale*.jar where *locale* is the abbreviation for the language you installed
Device server	*TPC_install_dir*\device\apps\was\profiles\deviceServer\installedApps\DefaultNode\DeviceServer.ear\DeviceServer.war\WEB-INF\lib	deviceSrv_*locale*.jar where *locale* is the abbreviation for the language you installed
GUI	*TPC_install_dir*\gui	TSRMgui_*locale*.jar where *locale* is the abbreviation for the language you installed
Storage Resource agent (see Table 38 on page 247)		

Table 37. Language Pack files (continued)

Component	Directory	Language Pack files
CLI	*TPC_install_dir*\cli\libs	tpccli_*locale*.jar where *locale* is the abbreviation for the language you installed deviceSrv_*locale*.jar where *locale* is the abbreviation for the language you installed

Note: The Device server language packs are copied to the Data server library (*TPC_install_dir*\data\server\lib) and CLI library (*TPC_install_dir*\cli\lib) so that the Data server and CLI can process messages that come from the Device server.

The Storage Resource agent files shown in Table 38 are copied to the *TPC_install_dir*\data\sra\SRA_platform as a part of the Language Pack installation.

Table 38. Storage Resource agent files copied to the Data server after the Language Pack installation

Locale	File that is copied
Czech	agent_cs.zip
French	agent_fr.zip
German	agent_de.zip
Hungarian	agent_hu.zip
Italian	agent_it.zip
Japanese	agent_jp.zip
Korean	agent_ko.zip
Polish	agent_pl.zip
Portuguese Brazil	agent_pt_BR.zip
Russian	agent_ru.zip
Spanish	agent_es.zip
Simplified Chinese	agent_zh_CN.zip
Traditional Chinese	agent_zh_TW.zip

The run-time jar files and their installed locations are shown in Table 39.

Table 39. Run-time jar files and installed locations

Run-time jar files	Installed location
TSRMgui_*locale*.jar	*TPC_install_dir*/gui
TSRMsrv_*locale*.zip	*TPC_install_dir*/data/server/lib
deviceSrv_*locale*.jar	*TPC_install_dir*/device/apps/was/profiles/deviceServer/installedApps/DefaultNode/DeviceServer.ear/DeviceServer.war/WEB-INF/lib *TPC_install_dir*/data/server/lib *TPC_install_dir*/cli/libs

Table 39. Run-time jar files and installed locations (continued)

Run-time jar files	Installed location
tpccli_*locale*.jar	*TPC_install_dir*/cli/libs

The *locale* can be one of the following abbreviations:

cs - Czech	it - Italian	pt_BR - Brazilian Portuguese
de - German	ja - Japanese	zh_CN - Simplified Chinese
es - Spanish	ko - Korean	zh_TW - Traditional Chinese
fr - French	pl - Polish	
hu - Hungarian	ru - Russian	

The Storage Resource agent files are shown in Table 40.

Table 40. Location of Windows, UNIX, or Linux agent after an upgrade

Locale	Windows agent location	UNIX or Linux agent location
English (this file gets copied by default)	nls\tpcnaenu.txt	nls/en_US/tpcna.cat
Simplified Chinese	nls\tpcnachs.txt	nls/zh_CN/tpcna.cat
Traditional Chinese	nls\tpcnacht.txt	nls/zh_TW/tpcna.cat
Czech	nls\tpcnacsy.txt	nls/cs_CZ/tpcna.cat
French	nls\tpcnafra.txt	nls/fr_FR/tpcna.cat
German	nls\tpcnadeu.txt	nls/de_DE/tpcna.cat
Hungarian	nls\tpcnahun.txt	nls/hu_HU/tpcna.cat
Italian	nls\tpcnaita.txt	nls/it_IT/tpcna.cat
Japanese	nls\tpcnajpn.txt	nls/ja_JP/tpcna.cat
Korean	nls\tpcnakor.txt	nls/ko_KR/tpcna.cat
Poland	nls\tpcnaplk.txt	nls/pl_PL/tpcna.cat
Portuguese Brazil	nls\tpcnaptb.txt	nls/pt_BR/tpcna.cat
Russian	nls\tpcnarus.txt	nls/ru_RU/tpcna.cat
Spain	nls\tpcnaesp.txt	nls/es_ES/tpcna.cat

Installing and configuring Tivoli Common Reporting

Use this procedure to install and configure Tivoli Common Reporting.

Preparing for Tivoli Common Reporting

Tivoli Common Reporting requires the following libraries to be installed on a Linux or AIX system.

If you are installing Tivoli Common Reporting on a 32-bit or 64-bit system, you must install the following 32-bit libraries on a Linux system:

- Red Hat Enterprise Linux 5.0 32-bit and 64-bit (Tivoli Storage Productivity Center runs in 32-bit mode on a 64-bit system):

  - compat-libstdc++-33.3.2.3
  - compat-glibc-2.3.4-2.26
  - openmotif22-2.2.3-18
- Red Hat Enterprise Linux 5.0 - zSeries or System z 31-bit and 64-bit (Tivoli Storage Productivity Center runs in 32-bit mode on a 64-bit system):
  - libXmu-1.0.2-5
  - libXp-1.0.0-8
  - openmotif22-2.2.3-18

For information about the hardware and software requirements for Tivoli Common Reporting, see http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_21/rtcr_soft_and_hard_reqs.html.

For information about installing Tivoli Common Reporting on a local server, see "Installing Tivoli Common Reporting ."

If you are installing Tivoli Common Reporting on a system that does not have the Tivoli Storage Productivity Center database (TPCDB) installed, you must install the DB2 client on the remote system and catalog TPCDB to point to the Tivoli Storage Productivity Center server. For information about installing the DB2 client, see the DB2 information center. For information about cataloging the database, see the DB2 information center. Search for CATALOG DATABASE. For information about installing Tivoli Common Reporting on a remote server, see "Installing Tivoli Common Reporting on a remote server" on page 251.

Installing Tivoli Common Reporting

If you are installing Tivoli Common Reporting 2.1.1 on the Tivoli Storage Productivity Center system, a Tivoli Integrated Portal profile is created. If you are installing Tivoli Common Reporting 2.1.1 on a different system, Tivoli Integrated Portal is also installed.

Before installing Tivoli Common Reporting, gather information that is required by Tivoli Common Reporting. You must know the administrative user ID and password and the port to use when creating a Tivoli Integrated Portal profile. The default port is 16310.

To install Tivoli Common Reporting, complete the following steps:

1. Log on to the system with administrator authority on a Windows system or root authority on a UNIX or Linux system.
2. Insert the Tivoli Common Reporting DVD into a drive or download the electronic image.
3. Go to the installation directory and start the `launchpad.exe` file.
4. From the IBM Tivoli Common Reporting page, click **Install IBM Tivoli Common Reporting 2.1.1**.
5. Select the language you want to use for installation. For example, select **English**.
6. On the Welcome to IBM Tivoli Common Reporting 2.1.1 page, click **Next**.
7. On the Software License Agreement page, read the terms of the license agreement. If you agree with the terms, click **I accept the terms in the license agreement** and click **Next**. If you do not agree with the terms and click **Cancel**. The program is not installed.

8. On the Deployment Engine Initialization page, after the deployment engine is installed, click **Next**.
9. On the Installation Mode page, select **Install a new instance of Tivoli Common Reporting**. Click **Next**.
10. On the Installation Scenario Selection page, select **Single-computer installation**. Click **Next**.
11. On the Installation Directory Selection page, you can use the default installation directory or enter a different directory. The default directory is: `C:\IBM\tivoli\tcr` on a Windows system and `/opt/IBM/tivoli/tcr` on a UNIX or Linux system. Click **Browse** to select a directory from a list of directories. Click **Next**.
12. On the WebSphere Information page, enter this information:

    User ID
    Enter an administrative user ID. This user ID is used to create a Tivoli Integrated Portal profile and to log on to the Tivoli Integrated Portal console.

    Password and Confirm Password
    Enter the password associated with the User ID.

    Port Number
    Enter the beginning range of ports. Tivoli Integrated Portal uses 15 ports starting with the beginning range of ports. The default beginning range of ports is 16310. If you use the default beginning range of ports, the ports used are 16310 - 16324. Click **Next**.

    Restriction: If you are installing Tivoli Common Reporting on the Tivoli Storage Productivity Center server, you must change the default port number. For example, if you specify the beginning port of 17310, ports 17310 - 17324 must be available.

    If there is a port conflict, the following message is displayed:

    ```
    The specified base port is in use. New base port has been selected.
    New base port port_number has been selected.
    ```

    where *port_number* is the new port. Click **Next**, and then click **Quit**
13. On the IBM Cognos® Content Database page, enter the port number for the IBM Cognos database. The default port number is 1527. Click **Next**.
14. After Tivoli Common Reporting is configured, click **Next**.
15. The Pre-Installation Summary page displays the components to be installed, the disk space required, the log file location, and the ports to be used. Review the summary page to ensure that you have disk space required and do not have any port conflicts. Click **Install**. The process can take from 20 minutes up to one hour to install.
16. Configure Tivoli Common Reporting. See "Configuring Tivoli Common Reporting" on page 252.

Any time you start the Tivoli Common Reporting server on a UNIX or Linux system, you must source the DB2 profile.

For more information about configuring Tivoli Common Reporting, go to the Tivoli Common Reporting information center at http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_21/ttcr_configure.html.

Installing Tivoli Common Reporting on a remote server

If you are installing Tivoli Common Reporting on a remote server, follow these steps:

1. Install the DB2 client on the remote server. For information about installing the DB2 client, see the DB2 information center.
2. Install Tivoli Common Reporting on the remote server. See "Installing Tivoli Common Reporting " on page 249.
3. Create the data source in DB2 for the remote server.
   a. Catalog the DB2 TCP/IP node.

      Run the following command:

      ```
      CATALOG TCPIP NODE db2_node_name
      REMOTE host_name
      SERVER port_number
      REMOTE_INSTANCE db2_instance_name
      SYSTEM system_name
      OSTYPE operating_system_type
      ```

      In the preceding command, specify the following values:

      CATALOG TCPIP NODE *db2_node_name*
      This value specifies the node name of the DB2 server.

      REMOTE *host_name*
      This value specifies the host name or IP address of the DB2 server.

      SERVER *port_number*
      This value specifies the port number of the DB2 server. The default port is 50000.

      REMOTE_INSTANCE *db2_instance_name*
      This value specifies the DB2 instance name on the remote server.

      SYSTEM *system_name*
      This value specifies the DB2 system name that is used to identify the server system.

      OSTYPE *operating_system_type*
      This value specifies the operating system type of the server system. Valid values for Tivoli Common Reporting are: AIX, WIN, and LINUX.

      An example for this command is:

      ```
      CATALOG TCPIP NODE db2node
      REMOTE system53.ibm.com
      SERVER 50000
      REMOTE_INSTANCE db2 system53
      SYSTEM system53
      OSTYPE win
      ```

   b. Catalog the Tivoli Storage Productivity Center database.

      Run the following command:

      ```
      CATALOG DATABASE tpc_database_name AS tpc_database_alias_name
      AT NODE db2_node_name AUTHENTICATION SERVER
      ```

      In the preceding command, specify the following values:

      CATALOG DATABASE *tpc_database_name*
      This value specifies the name of the remote database to catalog.

AS *tpc_database_alias_name*
This value specifies the alias for the remote database being cataloged. If you do not specify an alias, the database manager uses the name of the remote database as the alias.

AT NODE *db2_node_name*
This value specifies the name of the node that you previously cataloged.

AUTHENTICATION SERVER
This value specifies that authentication takes place on the DB2 data source node.

An example for this command is:

```
CATALOG DATABASE tpc_database_name AS tpc_database_alias_name
AT NODE db2node AUTHENTICATION SERVER
```

4. Configure Tivoli Common Reporting. See "Configuring Tivoli Common Reporting."

Configuring Tivoli Common Reporting

After you install Tivoli Common Reporting, follow these steps to configure Tivoli Common Reporting.

Note: To run the Tivoli Common Reporting commands, run the commands from the following default directory (version 2.1.1).

On a Windows system, the default directory is:

```
C:\IBM\Tivoli\tipv2Components\TCRComponent\bin
```

On a UNIX or Linux system, the default directory is:

```
/opt/IBM/Tivoli/tipv2Components/TCRComponent/bin
```

1. Create the data source connection in Tivoli Common Reporting.

   Run the following command:

   ```
   trcmd.bat (for Windows) or trcmd.sh (for UNIX or Linux)
   -user tip_admin_user
   -password password
   -dataSource -add TPCDB
   -dbType DB2
   -connectionName TPCDB
   -dbName TPCDB
   -dbLogin database_user_login
   -dbPassword password
   -groups group_access
   ```

   In the preceding command, specify the following values:

   -user *tip_admin_user*
   This value specifies the administrative user name for Tivoli Integrated Portal. This is the same user name that you use to log on Tivoli Common Reporting.

   -password *password*
   This value specifies the password for the *tip_admin_user*.

   -dataSource -add TPCDB -dbType DB2
   These values specify the name of the Tivoli Storage Productivity Center database and type.

   -connectionName TPCDB
   This value specifies the name of the connection.

-dbName TPCDB
This value specifies the name of the database.

-dbLogin *database_user_login*
This value specifies the user name for logging in to DB2.

-dbPassword *password*
This value specifies the password for the database user ID.

-groups *group_access*
This value specifies which groups can access the data. Enter **Everyone**.

An example for this command on a Windows system is:

```
trcmd.bat -user tipadmin -password tippassword
-dataSource -add TPCDB
-dbType DB2 -connectionName TPCDB -dbName TPCDB
-dbLogin databaseuser
-dbPassword db2password
```

An example for this command on a UNIX or Linux system is:

```
trcmd.sh -user tipadmin -password tippassword
-dataSource -add TPCDB
-dbType DB2 -connectionName TPCDB -dbName TPCDB
-dbLogin databaseuser
-dbPassword db2password
```

For more information about the **trcmd** command, see the Tivoli Common Reporting information center. Search for **trcmd -datasource**.

2. Test the data source connection.

   Run the following command:

   ```
   trcmd.bat (for Windows) or trcmd.sh (for UNIX or Linux)
   -user tip_admin_user
   -password password
   -dataSource -test TPCDB
   -dbLogin database_user_login
   -dbPassword password
   ```

   In the preceding command, specify the following values:

   -user *tip_admin_user*
   This value specifies the administrative user name for Tivoli Integrated Portal. This is the same user name that you use to log on Tivoli Common Reporting.

   -password *password*
   This value specifies the password for the *tip_admin_user* user.

   -dataSource -test TPCDB
   This value specifies the name of the Tivoli Storage Productivity Center database for testing.

   -dbLogin *database_user_login*
   This value specifies the user name that you use to log on DB2.

   -dbPassword *password*
   This value specifies the password for the database user login.

   An example for this command on a Windows system is:

   ```
   trcmd.bat
   -user tipadmin
   -password tippassword
   -dataSource -test TPCDB
   -dbLogin databaseuser
   -dbPassword db2password
   ```

   An example for this command on a UNIX or Linux system is:

```
trcmd.sh
-user tipadmin
-password tippassword
-dataSource -test TPCDB
-dbLogin databaseuser
-dbPassword db2password
```

3. Import the deployment package into Tivoli Common Reporting. The deployment package has the Tivoli Common Reporting reports. The deployment package is in the Tivoli Storage Productivity Center installation image in the following directory:

   ```
   tool\report
   ```

   Run the following command:

   ```
   trcmd.bat (for Windows) or trcmd.sh (for Linux or UNIX)
   -import -bulk path_of_zip_file
   -user tip_admin_user
   -password password
   ```

   In the preceding command, specify the following values:

   -import -bulk *path_of_zip_file*
   This value specifies the location of the compressed file. The compressed file name is `TPC_Reports.zip`. The compressed file is located in the Tivoli Storage Productivity Center installation image in the following directory: `tool\report\TPC_Reports.zip`.

   -user *tip_admin_user*
   This value specifies the administrative user name for Tivoli Integrated Portal. This is the same user name that you use to log on Tivoli Common Reporting.

   -password *password*
   This value specifies the password for the *tip_admin_user* user ID.

   An example of this command on a Windows system is:

   ```
   trcmd.bat
   -import -bulk c:\temp\TPC_Reports.zip
   -user tipadmin
   -password tippassword
   ```

   An example of this command on a UNIX or Linux system is:

   ```
   trcmd.sh
   -import -bulk /tmp/TPC_Reports.zip
   -user tipadmin
   -password tippassword
   ```

4. Log on to the Tivoli Integrated Portal console.

   For the Cognos reports, open the Tivoli Integrated Portal console.

   For example, if you installed Tivoli Common Reporting using port 17310:

   ```
   https://host_name:17310/ibm/console
   ```

   where *host_name* is the host name or IP address of the server where Tivoli Common Reporting is installed.

   Tip: In this example, change the port number (17310) to match what you used in step 12 on page 250 for installing Tivoli Common Reporting.

Chapter 3. Configuring IBM Tivoli Storage Productivity Center

After you have successfully installed IBM Tivoli Storage Productivity Center, you must configure it. You can use the Configuration Utility to learn how to configure and use Tivoli Storage Productivity Center. This section also describes each node under **Administrative Services > Configuration**.

The Data server hosts the control points for product scheduling functions, configuration, event information, reporting, and graphical user interface support. The Data server coordinates communication with the agents for data collection. The agents scan file systems and databases to gather storage demographics and populate the database with results. Automated actions can be defined to drive functions like data deletion, IBM Tivoli Storage Manager backup or archiving, or event reporting when defined thresholds are encountered. The Data server is the primary contact point for all user interface functions. The user interface functions also include actions that schedule data collection and discovery for the Device server.

The Device server discovers objects, gathers information from, analyzes performance of, and controls storage subsystems and SAN fabrics. The Device server coordinates communication with the agents to collect data about the SAN fabrics.

Configuration overview

After you have installed IBM Tivoli Storage Productivity Center, you must configure your system to get the information you want. This section provides information about the agents required, the jobs that must be run, and the alerts you can set.

To enable data gathering from the devices and servers that are managed and monitored, different configuration steps are required to enable these different functions in Tivoli Storage Productivity Center.

Tivoli Storage Productivity Center uses four different types of agents to gather data about the devices and servers that are managed and monitored. Different combinations of these agents are required to effectively enable the functions of Data Manager, Fabric Manager, Disk Manager, and Tape Manager. In addition to these manager functions, the topology viewer is greatly affected by the proper discovery of all the managed entities in the management scope of Tivoli Storage Productivity Center.

Note: CIM agents are not required for DS8000, the XIV system, SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 storage subsystems.

Agents

The agents are:

Storage Management Initiative Specification (SMI-S) providers

SMI-S providers or Common Information Model (CIM) agents are provided by the vendor of the storage device, fabric switch, or tape library. For

© Copyright IBM Corp. 2001, 2012

storage, they are needed for storage asset information, provisioning, alerting, and performance monitoring.

Note: CIM agents are not required for the DS8000, SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, and the XIV system storage systems. Tivoli Storage Productivity Center uses native interfaces for these storage systems.

For the fabric switch, they are used for performance monitoring. The CIM agents are also used for certain switch models to collect topology information, port state information, zoning information, zone control, and alerting. For tape libraries, they are used for asset and inventory information. Each vendor of the storage, fabric switch, or tape library supplies unique CIM agent code for their family of devices. These agents implement an SMI-S provider to provide a communication transport between Tivoli Storage Productivity Center and the managed devices. After a CIM agent is installed and configured, Tivoli Storage Productivity Center can be configured to communicate with it.

When setting up your CIM agents for managing devices, read the CIM agent documentation first. This documentation can provide guidelines for how many devices the CIM agent can be configured to manage. If the CIM agent documentation does not include such guidelines, limit three subsystems per CIM agent.

Data agents

Note: The Storage Resource agents now perform the functions of the Data agents.

The Data agents are installed on all the computer systems that you want Tivoli Storage Productivity Center to manage. These agents collect information from the server they are installed on. Asset information, file and file system attributes, and any other information needed from the computer system is gathered. Data agents can also gather information about database managers installed on the server, Novell NDS tree information, and NAS device information. In Tivoli Storage Productivity Center, you can create ping, probe, and scan jobs to run against the servers that have Data agents installed.

Storage Resource agents

The Storage Resource agents now perform the functions of the Data agents and Fabric agents. You are no longer required to install and maintain the Agent Manager, Data agents, and Fabric agents.

For computers that have storage that should be monitored, the Storage Resource agents are used to monitor that storage. These agents collect information from the server they are installed on. Asset information, file and file system attributes, and any other information needed from the computer system is gathered. The Storage Resource agents can also gather information about database managers installed on the server, and NAS device information. In Tivoli Storage Productivity Center, you can create ping, probe, and scan jobs to run against the servers that have Storage Resource agents installed.

For fabric information, the Storage Resource agents are installed on computer systems that have fiber connectivity (through HBAs) into the SAN fabrics you want to manage and monitor. The Storage Resource agents use scanners to collect information. The scanners communicate

through the HBA to collect fabric topology information, port state information, and zoning information. They also can identify other SAN-attached devices (if they are in the same zone). Storage Resource agents are discovered during the agent installation process and are not required to be discovered separately. A best practice is to have two agents connected to each switch for redundancy.

Fabric agents

Note: The Storage Resource agents now perform the functions of the Fabric agents.

Fabric agents are installed on computer systems that have fiber connectivity (through HBAs) into the SAN fabrics you want to manage and monitor. Fabric agents use scanners to collect information. The scanners communicate through the HBA to collect fabric topology information, port state information, and zoning information. They also can identify other SAN-attached devices (if they are in the same zone). Fabric agents are discovered during the agent installation process and are not required to be discovered separately. A best practice is to have two agents connected to each switch for redundancy.

Out-of-band agents

The out-of-band agents are used to collect topology information from fabric switches through the IP network using SNMP queries to the switches. A best practice is to have Tivoli Storage Productivity Center communicating with the out-of-band agent (SNMP agent) on each switch in each SAN fabric you are managing.

Jobs

Tivoli Storage Productivity Center uses data collection jobs to provide information about your environment. The types of data collection jobs are:

Discovery jobs

Discovery jobs locate data sources and collect basic information about these data sources.

When the discovery job is run for the DS8000, SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, or XIV system storage system, Tivoli Storage Productivity Center interfaces directly with the storage system.

For other storage systems, the discovery job runs against a storage subsystem CIM agent or tape library CIM agent. The discovery job locates all the storage systems or tape libraries for the CIM agent and retrieves basic information that the CIM agent holds for these systems. When a discovery job is run against a fabric switch CIM agent, the job retrieves basic available information about the fabric and switches.

Probe jobs

Probe jobs collect detailed statistics on all the assets of the managed infrastructure, such as computers, disk controllers, hard disks, clusters, fabrics, storage subsystems, LUNs, tape libraries, and file systems. Probe jobs can also discover information about new or removed disks and file systems. Probe jobs can be directed against any element in the managed infrastructure.

There are four types of probe jobs: for computers, fabrics, storage subsystems, and tape libraries.

Computer probe jobs collect local host information such as file systems, directories, clusters, and hard disks. Storage subsystem probe jobs collect information about the storage subsystems, such as disk, volumes, and LUNs. Fabric probe jobs collect information about the fabric. Tape library probe jobs collect information about the tape libraries.

Computer probe jobs for a Storage Resource agent also perform a fabric discovery. The information collected includes fabric WWNs of the fabrics connected to the host on which the Storage Resource agent is running, fabric to switch relationships, and some key switch attributes. Computer probe jobs for a Storage Resource agent also collect HBA information.

When probing storage subsystems that are registered with the same CIM agent, do not probe more than three storage subsystems within the same probe job because the increased load on the agent would increase the likelihood of timeouts. Instead, spread the storage subsystems across multiple probe jobs with different starting times.

Scan jobs
Scan jobs collect statistics about the actual storage consumption. Scan jobs are always directed against a Data agent or Storage Resource agent and deliver detailed information about the file systems, files, and databases of a computer.

Ping jobs
Ping jobs gather statistics about the availability of the managed computers. Ping jobs generate TCP/IP pings and consider the computer available if it gets an answer. Like scan jobs, ping jobs can be directed against computers only.

Performance monitoring jobs
Performance monitoring jobs collect statistics about the performance of storage subsystems and switches. Performance monitoring jobs can be run against storage subsystems, SAN Volume Controllers, Storwize V7000 Unified, or Storwize V7000s, and switches.

Note: Storage Subsystem Groups are not supported for the performance monitor.

Alerts

You can define alerts so that 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 depends on the storage resource that you are monitoring.

The alerts you can define are:

Computer
Use this alert to be notified when a condition is detected on a computer during a probe job.

Examples of computer alerts are when RAM is increased or decreased, virtual storage is increased or decreased, or when a new disk is detected or not found.

This alert is a Data Manager alert.

Filesystem
Use this alert to be notified when a condition is detected on a file system during a probe job or upon an SNMP event.

Examples of file system alerts are when the file system is not found or reconfigured or is low on free space.

This alert is a Data Manager alert.

Directory
Use this alert to be notified when a condition is detected on a directory during a scan job.

Examples of directory alerts are when the directory is not found or when a user or directory storage quota is exceeded.

This alert is a Data Manager alert.

Hypervisor
Use this alert to be notified when a hypervisor is discovered or missing, or a virtual machine is added or deleted from a VMware system.

This alert is a Data Manager alert.

Other NAS
Use this alert to be notified when a condition is detected on a network-attached storage (NAS) device during a probe job or upon an SNMP event. This alert type is for NAS devices that are not included in other alert types.

This alert is a Data Manager alert.

Instance
Use this alert to be notified when a condition is detected for a database.

Examples of instance alerts are when a new database or table space is discovered for a specific RDBMS.

This alert is a Data Manager for Databases alert.

Database-tablespace
Use this alert to be notified when a condition is detected for a database-tablespace.

Examples of database-tablespace alerts are when Tivoli Storage Productivity Center detects that free space is low or the database or table space is offline or dropped for a specific RDBMS.

This alert is a Data Manager for Databases alert.

Table Use this alert to be notified when a condition is detected for a table in a database.

Examples of table alerts are when a managed table is dropped or exceeds its storage quota.

This alert is a Data Manager for Databases alert.

Storage subsystem
Use this alert to be notified when a condition is detected on a storage subsystem during a probe job.

Examples of storage subsystem alerts are as follows:

- When a storage subsystem is discovered or not found
- A subsystem cache is increased or decreased

- A disk is detected or not found,
- A subsystem goes offline or online
- A port is discovered or missing
- A back-end controller goes offline or online
- A volume goes offline or online
- A pool is missing or discovered
- A node goes offline or online,
- A SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 node is missing or discovered
- A disk utilization percentage threshold is reached

This alert is a Disk Manager alert.

Note: Storage Subsystem Groups are not supported for threshold alerts.

Fabric
Use this alert to be notified when a condition is detected on a fabric during a probe job.

Examples of fabric alerts are if a fabric is missing or discovered, a fabric goes online or offline, a zone is missing or discovered, or a zone set is missing or discovered.

This alert is a Fabric Manager alert.

Switch
Use this alert to be notified when a condition is detected on a switch during a probe job.

Examples of switch alerts are if a switch is missing or discovered, a switch goes online or offline, a switch blade is missing or discovered, a switch blade goes online or offline, a total port data rate threshold is reached, or a total port packet rate threshold is reached.

This alert is a Fabric Manager alert.

Endpoint device
Use this alert to be notified when a condition is detected on an endpoint device during a probe job. Endpoint devices are devices discovered by the Fabric Manager but cannot be identified as a specific device type, such as a computer or subsystem.

Examples of endpoint device alerts are when an endpoint device is missing or discovered, a peripheral entity to node association is missing or discovered, or the version for an endpoint device changes.

This alert is a Fabric Manager alert.

After you have defined alerts, you can choose to be notified through different methods.

SNMP trap
Generate an SNMP trap message to any console or terminal to indicate the occurrence of an alert.

TEC/OMNIbus event
Send alerts to the IBM Tivoli Enterprise Console or OMNIbus console. The Tivoli Enterprise Console administrator can write correlation and

automation rules to analyze Data Manager events and perform responses such as send further notification, create or update trouble tickets, or run programs.

Login notification
Send alerts to a specified user upon logging on to the system.

Windows Event Log, UNIX Syslog
Record alerts to the OS Event Log. If you already have an administrator monitoring OS logs, this action is an easy way to have all of your priority messages centralized for quick notification and viewing.

Run script
Run a script in response to an alert. (This option is not valid for the Storage Resource agent.)

Email Send alerts through to email addresses that you specify.

Configuration Utility

Use the Configuration Utility to help you learn about how to configure your IBM Tivoli Storage Productivity Center system.

Configuration utility tasks

The configuration utility is a tool that guides you through the steps and prerequisite steps required to configure your system.

All of the functions in the configuration utility exist on their own within the IBM Tivoli Storage Productivity Center GUI. However, the configuration utility provides you with a single place for new users to find all these common functions and information to help you decide what agents you need and which jobs to run to get the information you want. It provides links into the Tivoli Storage Productivity Center GUI to display reports, add agents, create jobs, create alerts, and so on.

If you are a new user of Tivoli Storage Productivity Center, use the configuration utility to learn what the common functions are and to help you with the task of deciding what agents you need and which jobs to run to get the information you want.

If you are an experienced Tivoli Storage Productivity Center user, you can go to the GUI directly to perform all of the configuration tasks. However, you might also want to use the configuration utility to provide you with a quick summary review of your system.

Depending on what information you want to see and collect, you will go to a different tab in the configuration utility for information and configuration.

How to manage the servers and agents

This section describes how to use the configuration utility to manage your servers and agents.

To manage your servers, storage subsystems, and agents, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** > **Configuration Utility** > **Services** tab.
2. You can see the following items to manage:

- Data Server - indicates whether the server is up or down. Provides detailed information about the Data server.
- Device Server - indicates whether the server is up or down. Provides detailed information about the Device server.
- Data/Storage Resource Agents - lists the Data agents and Storage Resource agents. Provides detailed information about the agents and their status.
- CIM Agents - lists the CIM Agents installed and connected to Tivoli Storage Productivity Center. You can add CIM agents, edit CIM agent details, test the CIM agent connections, run a CIMOM discovery, and show the managed devices for the CIM agent.
- Inband Fabric Agents - lists the Inband Fabric agents installed. Provides detailed information about the agents and their status.
- Out of Band Fabric Agents - lists the Out of Band Fabric agents. Provides detailed information about the agents and their status.
- Storage Subsystem Interfaces - lists the storage systems. You can add, modify, or remove a storage system. You can also update or configure the credentials for the storage system, or test the connection to the storage system.

Show managed devices

This section describes how to use the configuration utility to show the devices (storage subsystems, switches, or tape libraries) that are managed by the CIM agent.

To show the devices, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** > **Configuration Utility** > **Services** tab.
2. Under the CIM Agents section, click **Show Managed Devices**. A dialog box displays the managed devices.

How to display availability reports

This section describes how to use the configuration utility to display availability reports.

To display availability reports of storage assets, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** > **Configuration Utility** > **Data Manager** tab.
2. In the Ping section, you can run, create, or display information about the ping job. You can then display or define the storage availability system report.

How to monitor your storage assets, storage systems, and fabrics

You can use the configuration utility to monitor your storage assets, storage systems, and fabrics in your enterprise.

To monitor your storage assets, storage systems, and fabrics, complete the following steps:

1. For information collected from the Data Manager, expand **IBM Tivoli Storage Productivity Center** > **Configuration Utility** > **Data Manager** tab. Use the Ping section to collect information about the availability of the hosts managed by Storage Resource agents in your enterprise. The ping jobs provide information for the availability reports.

Use the Computer Probes section to collect statistics about the storage assets in your enterprise. The computer probe jobs provide information for different reports.

Use the Scan section to collect statistics about the files on the hosts where Storage Resource agents are installed. The scan jobs provide information for capacity, usage, usage violations, and backup reports.

2. For information collected from the Disk Manager, expand **IBM Tivoli Storage Productivity Center** > **Configuration Utility** > **Disk Manager** tab. Use the Storage Subsystems Probes section to collect information about the configuration of the storage system. Some of the information collected are storage pools, storage volumes, FC ports, disks, and host-to-volume assignments. A probe job must be run before volumes can be created or monitored for performance on a storage system.

   Use the Storage Subsystem Performance Management section to create and edit a performance monitor job. The performance monitor job collects information about the performance statistics for the storage system.

3. For information collected from the Fabric Manager, expand **IBM Tivoli Storage Productivity Center** > **Configuration Utility** > **Fabric Manager** tab.

   You can have Storage Resource agents installed on systems in each fabric that you are managing, CIM agents that are configured to manage each fabric, and out-of-band Fabric agents that communicate with each switch in the fabric. It is a good practice to have multiple agent types configured. CIM agents are the preferred agent type for Brocade fabrics. For Cisco switches, out-of-band agents are required for VSAN information. Zoning information is provided only by the Storage Resource agents for Cisco and QLogic switches.

   Before you can run a switch performance monitoring job in Tivoli Storage Productivity Center, you must have a switch CIM agent installed and configured in Tivoli Storage Productivity Center. This CIM agent must be SMI-S V1.1 compliant.

   For more information about the supported agent types for switch performance management and fabric zone configuration, see “Collecting data with the Fabric Manager” on page 105.

   Use the Fabric Probes section to collect information about the configuration of the fabric. A probe job must be run before information about performance for a fabric can be collected.

   Storage Resource agents perform fabric discovery during a computer probe and in response to events. The information gathered during a fabric discovery performed by the Storage Resource agent includes fabric and fabric to switch relationships.

   Storage Resource agent discovery gathers the WWN of the fabric and WWN of the switches. For topology information, the agent discovers the topology for the entire fabric. A fabric probe with the Storage Resource agent collects other information such as zoning, topology, and so on. For more information about what information is gathered, see “Planning for the Fabric Manager” on page 98.

How to create alert definitions

This section describes how to use the configuration utility to create alert definitions.

You can create alert definitions to notify you of an alert condition. When you receive an alert notification, you can take the appropriate action when a certain condition occurs on storage subsystems, computers, file systems, directories, fabrics, switches, or endpoint devices.

To create an alert definition, complete the following steps:

1. If you want to see the alerts for Data Manager, expand **IBM Tivoli Storage Productivity Center** > **Configuration Utility** > **Data Manager** tab to create alerts for a computer resource.

   In the Alerting section, you can create alerts for computers, file systems, and directories. You can also see the alert logs for computers, file systems, and directories.

2. If you want to see the alerts for Disk Manager, expand **IBM Tivoli Storage Productivity Center** > **Configuration Utility** > **Disk Manager**. tab

   In the Alerting section, you can create alert definitions for the storage system and view the alert logs.

   There are two categories of alerts: device alerts and performance alerts. Device alerts use triggering conditions such as actions on the device itself. These alerts include discovery or state changes. Performance alerts use triggering conditions that are thresholds on measured metrics.

   Performance thresholds are only available with alert definitions when the IBM Tivoli Storage Productivity Center package contains the performance capability.

   Performance alerts must be defined before the performance monitor job is run. Because these alerts are based on thresholds violated from measured metrics, they can only be triggered when defined in an alert and when those metrics are being collected by a monitoring job.

3. If you want to see the alerts for Fabric Manager, expand **IBM Tivoli Storage Productivity Center** > **Configuration Utility** > **Fabric Manager** tab. In the Alerting section, you can create fabric, switch, and endpoint point device alerts. You can also view the logs for these alerts.

   There are two categories of alerts: the device alert and performance alert. Device alerts use triggering conditions which are actions on the device itself. Examples of these alerts are discovery, state change, and so on. Performance alerts use triggering conditions that are thresholds on measured metrics. Device alerts are composed of three categories: fabric alerts, switch alerts, and endpoint alerts. A best-practice procedure would be to define more than the three default alerts that are present. In addition, the three default alerts should be modified to include a "Triggered Action".

   Performance thresholds are only available for alert definitions when the IBM Tivoli Storage Productivity Center package contains the performance capability. For Fabric Manager, the only performance alerts are switch.

   Performance alerts must be defined before a performance monitor job is run. Because these alerts are based on thresholds violated from measured metrics, they can only be triggered when they are defined in an alert and when those metrics are being collected in a monitoring job.

How to manage tape library reports

This section describes how to use the configuration utility to manage tape library reports.

A CIM agent must be installed and configured to collect information from the tape libraries. After the CIMOM discovery job completes, you can create a tape library probe job that collects asset information.

To manage your tape library reports, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** > **Configuration Utility** > **Tape Manager** tab.
2. Use the CIMOM Discovery section to run a CIMOM Discovery job.

3. Use the Tape Library Probes section to create and run a probe job for the tape library.
4. After you have run the probe job, you can see information in the tape library reports. The probe job collects information about the configuration of the tape library. Some of the information collected are drives, changers, cartridges, slots, and I/O ports.

Managing element managers and external tools

This section describes how to use the Configuration Utility to add and manage element managers and external tools.

Element managers are programs that allow you to configure and maintain your storage devices. An example of an element manager is DS8000 Storage Manager, which is the default element manager that is provided with IBM System Storage DS8000. Examples of external tools are Web-based user interfaces and local executable files. For example, you might add the address of a website that you frequently access or an application that you often use (a text editor, a spreadsheet application, and so on). You can add and manage connections to element managers and external tools from the Configuration Utility.

To add and manage connections to element managers and external tools, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** > **Configuration Utility** > **Element Manager**.
2. See the following sections on the Element Manager page:
   - **DS8000 Element Manager**. This section shows the element managers that have been added for DS8000.
   - **Other Device Element Manager**. This section shows element managers that have been added for devices other than DS8000. For example, element managers for other DS series systems, IBM TotalStorage Enterprise Storage Server, SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, and IBM XIV Storage System.
   - **External Tools**. This section shows external tools such as Web-based interfaces and local executable files that have been added.
   - **Advanced Settings**. This section defines default settings for the element managers and external tools that are defined in the preceding sections. For example, you can define the default Web browser that you want to use to start Web-based element managers and external tools.

How to access IBM Tivoli Storage Productivity Center for Replication through IBM Tivoli Storage Productivity Center

This section describes how to access IBM Tivoli Storage Productivity Center for Replication through the IBM Tivoli Storage Productivity Center GUI.

Complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** > **Configuration Utility** > **Replication Manager** tab. For information about how to administer IBM Tivoli Storage Productivity Center for Replication, see the online help.
2. On the Replication Manager tab, you see these buttons:

   Replication Health Overview
   The **Replication Health Overview** button displays the Health Overview panel for IBM Tivoli Storage Productivity Center for Replication. This panel provides the following information:

- Overall session status: indicates the session status, which can be normal, warning, or severe.
- Overall storage system status: indicates the connection status of the storage system.
- 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.

Replication Sessions Overview
The **Replication Sessions Overview** button lists all sessions defined within the IBM Tivoli Storage Productivity Center for Replication environment including their state and status.

Replication Storage Systems Overview
The **Replication Storage Systems Overview** button displays the Storage Systems panel. The Storage Systems panel lists all the known storage systems, and indicates whether the storage systems are communicating normally with the active and remote servers, if enabled.

Replication Paths Overview
The **Replication Paths Overview** button displays the ESS/DS Paths panel. The ESS/DS Paths panel summarizes all the known ESS/DS series paths, listing them by storage system.

Replication Management Servers Overview
The **Replication Management Servers Overview** button displays the Management Servers panel. The Management Servers main panel displays the status of the management servers configuration, lists the management servers in operation (up to two), and enables you to define a standby server, or to define the local server as a standby server to an alternate server. This panel has two variations, depending on whether you have defined a standby.

Note: In order to see and access the Management Servers option on the menu, you must have IBM Tivoli Storage Productivity Center for Replication Two Site Business Continuity installed.

Replication Administration
The **Replication Administration** button displays the Administration panel. The Administration panel displays a list of IBM Tivoli Storage Productivity Center for Replication users and groups and their access privileges, and allows administrators to take actions on users and groups.

Replication Advanced Tools
The **Replication Advanced Tools** button displays the Advanced Tools panel. The Advanced Tools panel enables you to create a diagnostic package and change the automatic refresh rate of the GUI.

When you click a button, a new browser window is opened to the IBM Tivoli Storage Productivity Center for Replication page.

Configure Devices

Use the Configure Devices wizard to set up storage devices for monitoring by IBM Tivoli Storage Productivity Center. The wizard guides you through the steps for adding a device as a data source, running a discovery, including devices in groups, specifying alerts, and setting up data collection schedules.

In previous versions of IBM Tivoli Storage Productivity Center, setting up a storage device for monitoring required performing tasks in different locations of the user interface, and each type of device required different steps. In Tivoli Storage Productivity Center v4.2 and later, all the tasks for setting up and configuring a device are included in the Configure Devices wizard.

Use the wizard to configure the following device types for monitoring:

- Storage subsystems
- Fabrics and switches
- Computers
- Tape libraries

You can add multiple devices of the same type using a single session of the wizard. For example, you can add an IBM DS8000, IBM XIV, IBM SAN Volume Controller, IBM Storwize V7000, and storage subsystems associated with CIM agents using a single session of the wizard. You cannot configure multiple devices of different types at the same time. For example, you cannot add a storage subsystem and a fabric using a single session of the wizard.

The pages in the wizard are divided into two sections:

Steps (left side)
Displays the steps for completing in the wizard. An icon is displayed next to the current step. The wizard includes the following steps for configuring a device:

1. Select the device type that you want to configure.
2. Select the specific devices that you want to configure.
3. Determine how data is collected from the devices.
4. Review your selections and run the wizard.
5. Review the results of the wizard.

How you access the wizard determines on what step you begin. For example, if you click the **Configure Devices** icon in the toolbar or select the **Configure Devices** node in the navigation tree, you start by selecting the device type that you want to configure. If you click the **Add Storage Subsystem** button on the Disk Manager > Storage Subsystem panel, you begin by entering connection information about the storage systems that you want to configure.

Options (right side)
Displays the options and fields that you must select and complete for each step.

Configure storage subsystems

Use the Configure Devices wizard to set up storage subsystems for monitoring by Tivoli Storage Productivity Center. Storage subsystems that you configure in the wizard are automatically included in probe schedules and associated with alert conditions.

You can configure one or more storage systems for monitoring by Tivoli Storage Productivity Center using a single session of the Configure Devices wizard. All the storage systems that you configure in a single pass through the wizard share the same probe schedule and alert conditions.

Depending on the type of storage system that you want to configure for monitoring, you must define connection information for its native interface or associated CIMOM. Enter information in the wizard about the native interfaces for the following storage systems:

- IBM System Storage DS8000 (native interface)
- IBM System Storage SAN Volume Controller (native interface)
- IBM XIV Storage System (native interface)
- IBM Storwize V7000 (native interface)
- IBM Storwize V7000 Unified (native interface)

For all other storage systems, you must enter information about the CIM agents that manage them.

To configure storage systems using the wizard, complete the following steps:

1. Access the wizard based on the task you want to perform:
   - If you want to add new or configure existing storage systems, click the configure devices icon in the toolbar or select the **Configure Devices** node in the navigation tree. You cannot configure existing storage systems if they are already included in a probe schedule. If you use this method, you must also complete the following steps:
     a. Select **Storage Subsystem** on the Select Device Type page.
     b. Determine whether to add new storage systems or configure existing storage systems that are known to Tivoli Storage Productivity Center.
        – To add storage systems, select **Add and configure new storage subsystems** and go to step 3.
        – To configure existing storage systems, click **Configure existing storage subsystems**. The page displays the storage systems that are known to Tivoli Storage Productivity Center. Go to step 2.
   - If you only want to add new storage systems, use any of the following methods and then go to step 3:
     – In the navigation tree, go to **Disk Manager** > **Storage Subsystems** and click the **Add Storage Subsystem** button.
     – In the navigation tree, go to **IBM Tivoli Storage Productivity Center > Configuration Utility > Services > Storage Subsystem Interfaces** and click the **Add** button.
2. Select the existing storage systems that you want to configure. Storage systems that are already included in probe schedules are not displayed.

   If the Credentials column indicates **Missing** or **New Credentials Required** for the storage systems that you select, you must update their connection properties before continuing. See Configure storage subsystem connections page for descriptions of the connection properties.

   Go to step 6 after entering valid connection properties for each storage system that you want to configure.
3. Configure storage system or CIMOM connections. Tivoli Storage Productivity Center uses the information you enter in this step to connect to a storage system or CIMOM. See Configure storage subsystem connections page for descriptions of the fields and options on the Configure storage system connections page.

   After entering information for a storage system or CIMOM, click **Add** to add it to the wizard session. Repeat the process to configure multiple storage systems or CIMOMs in the same session of the wizard.

4. Discover the storage systems that are managed by the CIMOMs you entered in the previous step. Tivoli Storage Productivity Center collects initial configuration information for these entities.
5. Select the storage systems that you want to configure. Any storage systems that were added in previous steps are automatically selected. Deselect storage systems that you do not want to configure.
6. Specify how data is collected for storage systems by including them in a monitoring group or template. Use monitoring groups and templates to determine when data is collected and what alert conditions are checked for the storage systems that you are configuring.
   - Select a monitoring group if you want to manage a collection of storage systems in the same manner. By default, the monitoring group that you select is associated with an existing probe schedule and set of alerts. When you select the group, its data collection schedule and alerts are automatically applied to all the storage systems that you are configuring. After completing the wizard, the storage systems remain in the group and you can use that group when working in other parts of the user interface, such as in reports.

     Tip: Tivoli Storage Productivity Center provides a number of for storage systems. If you want to include a storage system in a user-defined group rather than in a default group, that group must be associated with a data collection schedule. If the user-defined group is not associated with a data collection schedule, it is not displayed in the list of monitoring groups from which you can select.
   - Select a template if you want to create a unique probe schedule and set of alerts for a storage system rather than using the settings of an existing group.

   The option that you select in this step determines the next step in the wizard.
   - If you select a monitoring group, the following actions occur:
     - The storage systems are automatically included in a probe schedule to which the group is associated. For example, if you select Subsystem Basic Group, the storage systems are included in the Subsystem Basic Probe schedule. The Subsystem Basic Probe schedule runs one time per week on Tuesday at 1:00 am.
     - A set of alert conditions are applied to the storage systems in a group. When any of these conditions are detected during data collection, an alert is triggered. The group that you select determines the alerts that are associated with the storage systems. For example, if you select Subsystem Basic Group, 14 alert conditions are automatically checked for the storage systems.
     - The storage systems are included in the group. You can use that group in other parts of the user interface, such as in reports.
   - If you select a monitoring template, the following actions occur:
     - The storage systems are automatically included in a probe schedule that is created for the template. For example, if you select the Basic Subsystem Monitoring template, the storage systems are automatically included in a new probe that runs one time per week on Tuesday. The probe is created after the wizard is completed. Tivoli Storage Productivity Center determines the time when the probe is run based on server load to ensure optimal performance.
     - A set of alert conditions are applied to storage systems in the template. When any of these conditions are detected during data collection, an alert is triggered. The template you select determines the alerts that are

associated with the storage systems. For example, if you select to the Basic Subsystem Monitoring template, 14 alert conditions are automatically checked for the storage systems.

– You must enter a prefix for the probe schedule and alerts that are created for the template. This prefix is used to identify the probe and alerts in the user interface. For example, if you select the Basic Subsystem Monitoring template and enter the prefix "ABC", a probe schedule named "*userID*.ABCssBasicProbe" and alerts such as "*userID*.ABCPool Discovered (Basic)" are created, where *userID* represents the ID of the user that is logged in toTivoli Storage Productivity Center.

  The new probe and alerts are displayed under the corresponding node in the navigation tree for you to view and edit. For example, a probe schedule created from a template in the wizard appears under the **Tivoli Storage Productivity Center >Monitoring > Probes** node; new storage system alerts appear under the **Disk Manager > Alerting** node.

See for a list of groups and templates with their associated schedules and alerts.

7. Review the selections that you made in the wizard. The Review user selections page includes the following information:
   - A list of devices that you are configuring.
   - The name of the monitoring group that you selected. This value is not displayed if you selected a template.
   - The name of the probe schedule that is created based on the template you selected. This value is not displayed if you selected a monitoring group.
   - Information about the probe schedule created for a template.
   - The names of the alerts that are created based on the template you selected. This value is not displayed if you selected a monitoring group.
8. View the results of the device configuration and click **Finish** to close the wizard and probe the devices that you configured. The Review results page includes the following information:
   - A list of the actions performed by the wizard. The page displays a row for the successful actions. If the configuration failed, a row is displayed for each failed action. For example, if the wizard expects to assign a specific alert to five devices, but the operation only succeeds for the three of the devices, this page displays one row for the successful actions and two rows each failed actions.
   - Error messages for any failed actions.

Configure storage subsystem connections page

Use this page in the Configure Devices wizard to enter or update connection properties for the storage systems and CIMOMs that you are configuring. Tivoli Storage Productivity Center uses this information to collect data about the storage systems.

Use the Configure storage system connections page to enter or update connection properties for the following storage entities. Depending on the type of storage system that you want to configure for monitoring, you must define connection information for the native interface or associated CIMOM for the storage system.

- System Storage DS8000 (native interface)
- SAN Volume Controller (native interface)
- XIV system (native interface)
- Storwize V7000 (native interface)

- Storwize V7000 Unified (native interface)
- Other (CIMOM)

The entity that is selected from the **Device Type** list determines the information that Tivoli Storage Productivity Center requires to connect to that entity. The following sections describe the fields for each entity.

Requirement: If you are adding or updating connection properties for a SAN Volume Controller or Storwize V7000 storage system, select **IBM SAN Volume Controller / IBM Storwize V7000** in the **Device Type** list. If you are adding or updating connection properties for a Storwize V7000 Unified storage system, select **IBM Storwize V7000 Unified**.

IBM DS8000

Enter or update the following information to have Tivoli Storage Productivity Center connect to a System Storage DS8000 storage system.

HMC Address
Enter the IP address or host name for the Hardware Management Console (HMC) that is used to manage the System Storage DS8000 storage system.

HMC2 Address (Optional)
Enter the IP address or host name for a second HMC that is used to manage the storage system.

Username
Enter the user name for logging in to the System Storage DS8000 element manager. The default user name is `admin`.

Note: This user name is the same as the ESSNI user name.

Password
Enter the password for the user name.

IBM XIV

Enter or update the following information to have Tivoli Storage Productivity Center connect to the XIV system.

IP Address
Enter the IP address or host name for the XIV system that you want to configure.

Username
Enter the user name for the XIV system.

Password
Enter the password for the user name.

IBM SAN Volume Controller / IBM Storwize V7000

Enter or update the following information to have Tivoli Storage Productivity Center connect to a SAN Volume Controller or Storwize V7000 storage system.

Software Version
Enter the version of the SAN Volume Controller or Storwize V7000 software that you want to configure.

IP Address

Enter the IP address or host name of the SAN Volume Controller or Storwize V7000. 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

A Secure Shell (SSH) key is required 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 SAN Volume Controller or Storwize V7000 cluster. If you select this action, the **Admin Username** and **Admin Password** fields are required.
- **Use existing uploaded key:** Select this option if you want to use a public key that was already uploaded by using a method other than through Tivoli Storage Productivity Center, such as the SAN Volume Controller or Storwize V7000 web interface. Enter the location of the private key file that corresponds to the public key in the **Private SSH Key** field. The **Admin Username** and **Admin Password** fields are not required if you select this option.

Admin Username

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 SAN Volume Controller or Storwize V7000 cluster that contains the storage system. Tivoli Storage Productivity Center uses this value during the setup process to contact the cluster 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.

Note: The **Admin Username** only applies to **SVC5.*x*+**; it does not apply to **SVC 4.*x***.

Admin Password

Enter the password for the user that you entered in the **Admin Username** field.

Username

Enter an authorized user for the storage system. Tivoli Storage Productivity Center uses this value when monitoring and configuring 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 have selected **Upload new key** in the **Select Key** list and have entered a user name and password in the **Admin Username** and **Admin Password** fields.
Important: If you upload a new key for the existing SAN Volume Controller or Storwize V7000 user by using Tivoli Storage Productivity Center, that new key overrides the existing key for the user. If you have another application that is using that user (for example Tivoli Storage Productivity Center for Replication), the application cannot log on to SAN Volume Controller or Storwize V7000 because the key associated with the user is overwritten.
If you have this situation, create a user ID specifically for Tivoli Storage

Productivity Center or Tivoli Storage Productivity Center for Replication on the SAN Volume Controller or Storwize V7000 cluster. Add the SAN Volume Controller or Storwize V7000 to Tivoli Storage Productivity Center for Replication using the new user ID and password (or change the user ID and password while the storage system is in the disconnected state).

Select User
Click this button to select from a list of existing users that are defined on the SAN Volume Controller or Storwize V7000.

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. Click **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 your private SSH key file. If you do not have a passphrase, leave this field blank.

Storwize V7000 Unified

Enter or update the following information to have Tivoli Storage Productivity Center connect to a Storwize V7000 Unified storage system.

Storwize V7000 Unified is composed of Storwize V7000 and Storwize V7000 File Module. You can connect to only Storwize V7000 or to both Storwize V7000 and Storwize V7000 File Module. Storwize V7000 provides block-level data and Storwize V7000 File Module provides file-level data.

Storwize V7000
Use the information in this section to connect Storwize V7000.

IP Address
Enter the IP address or host name of the Storwize V7000 storage 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
An SSH key is required to access Storwize V7000. 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 Storwize V7000 cluster. If you select this action, the **Administrator User Name** and **Administrator Password** fields are required.
- **Use existing uploaded key:** Select this option if you want to use a public key that was already uploaded by using a method other than through Tivoli Storage Productivity Center, such as the Storwize V7000 web interface. Enter the location of the private key file that corresponds to the public key in the **Private SSH Key** field. The **Administrator User Name** and **Administrator Password** fields are not required if you select this option.

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 Storwize V7000 cluster that contains the storage system. Tivoli Storage Productivity Center uses this value during the setup process to contact the cluster and configure the SSH keys for the user. This user name must belong to the Security Administrator role, which has privileges to modify other user accounts. 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 to monitor and configure the storage system. If the user name does not exist, it is created. This user name must belong to 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. **Important:** If you upload a new key for the existing Storwize V7000 user by using Tivoli Storage Productivity Center, that new key overrides the existing key for the user name. If you have another application that is using that user name (for example Tivoli Storage Productivity Center for Replication), the application cannot log on to Storwize V7000 because the key associated with the user is overwritten.

Select User
Select a user name from the list. The names are defined on Storwize V7000.

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. Click **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 your private SSH key file. If you do not have a passphrase, leave this field blank.

Storwize V7000 File Module (Optional)
To connect to the Storwize V7000 Unified system, select this check box.

Depending on the configuration of Storwize V7000 Unified, the information that is required to connect to this system can be the same as or different from the information for Storwize V7000. The fields and field descriptions are the same as the fields in the Storwize V7000 pane with the following exceptions:

- The **IP Address** field is not displayed because the IP address is retrieved automatically.
- The **User Name** field is always required.
- The private SSH key must be in OpenSSH format. The PuTTY format is not supported. You can use the puttygen utility to convert the key to the OpenSSH format.
- The **Select Key** list contains the option **Use user password** in addition to the **Upload new key** and **Use existing uploaded key** options. If you select **Use user password**, 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 **Administrator User Name** and **Administrator Password** fields are optional.

Other (CIMOM)

Enter or update connection properties for a CIMOM. Tivoli Storage Productivity Center uses this information to discover the storage systems that are managed by a CIMOM. You can then configure those storage systems as part of the wizard session.

The connection information that you enter for a CIMOM is similar to the information that you enter when adding a CIMOM on the Add CIMOM page. For information about the fields on this page, go to the information center. Search for **Adding a CIM agent connection**.

Discovered Storage Systems

After you run a discovery job, a table is displayed at the bottom of the "Configure storage subsystem connections" page. This table provides information about the storage systems that were discovered for the selected storage system type.

You can remove the connection to a storage system by clicking the row for the system in the table, and then clicking **Remove**.

Configure fabrics and switches

Use the Configure Devices wizard to set up fabrics and switches for monitoring by Tivoli Storage Productivity Center. Fabrics and switches that you configure in the wizard are automatically included in probe schedules and associated with alert conditions.

You can use the wizard to configure the following fabric and switch types for monitoring:

- Brocade
- McData
- Cisco
- QLogic
- Brocade/McData (heterogeneous fabric)

You can configure one or more fabrics using a single session of the Configure Devices wizard. All the fabrics and switches that you configure in a single pass through the wizard share the same probe schedule and alert conditions.

To configure fabrics and switches using the wizard, complete the following steps:

1. Access the wizard based on the task you want to perform:
   - If you want to add new or configure existing fabrics and switches, click the configure devices icon in the toolbar or select the **Configure Devices** node in the navigation tree. If you use this method, you must also complete the following steps:
     a. Select **Fabrics/Switches** on the Select Device Type page.
     b. Determine whether to add new fabrics or configure existing fabrics that are known to Tivoli Storage Productivity Center.
        – To add fabrics and switches, select **Add and configure new fabrics** and go to step 3.
        – To configure existing fabrics and switches, select **Configure existing fabrics**. The page displays the fabrics that are known to Tivoli Storage Productivity Center. Go to step 2.
   - If you only want to add new fabrics and switches, go to **Fabric Manager > Fabrics** and click the **Add Fabric** button. Go to step 3.
2. Select the existing fabrics that you want to configure. Fabrics that are already included in probe schedules are not displayed. Go to step 6 after selecting the fabrics that you want to configure.
3. Configure CIMOM or agent connections. The fields for entering connection information for a fabric changes based on the type of fabric you are adding. Tivoli Storage Productivity Center uses the information you enter in this step to connect to a fabric or CIMOM. Click **Add** to configure multiple fabrics in a session of the wizard.

Note: Depending on the type of fabric you are adding, the wizard recommends whether to enter CIMOM or agent information. For example, if you are configuring a Brocade fabric, the recommended way to administer it is with a CIMOM. If a CIMOM is not available, enter connection information for an out of band fabric agent.

4. Discover the fabrics and fabrics managed by the CIMOMs that you entered in the previous step. Tivoli Storage Productivity Center collects initial configuration information for these entities.
5. Select the fabrics that you want to configure for monitoring. Any fabrics that were added in previous steps are automatically selected. Deselect fabrics that you do not want to configure.
6. Specify how data is collected for fabrics by including them in a monitoring group or template. Use monitoring groups and templates to determine when data is collected and what alert conditions are checked for the fabrics that you are configuring.
   - Select a monitoring group if you want to manage a collection of fabrics in the same manner. Each monitoring group that you can select is associated with an existing probe schedule and set of alerts. When you select a group, its data collection schedule and alerts are automatically applied to all the fabrics that you are configuring. After completing the wizard, the fabrics remain in the group and you can use that group when working in other parts of the user interface, such as in reports.

     Tip: Tivoli Storage Productivity Center provides a number of for fabrics. If you want to include a fabric in a user-defined group rather than in a default group, that group must be associated with a data collection schedule. If the user-defined group is not associated with a data collection schedule, it is not displayed in the list of monitoring groups from which you can select.
   - Select a template if you want to create a unique probe schedule and set of alerts for a fabric rather than using the settings of an existing group.

   The option that you select in this step determines the field and pages that appear next in the wizard.
   - If you select a monitoring group, the following actions occur:
     - The fabrics are automatically included in a probe schedule to which the group is associated. For example, if you select Fabric Basic Group, the fabrics are included in the Fabric Basic Probe schedule. The Fabric Basic Probe schedule runs two times per week on Monday and Wednesday at 3:00 am.
     - A set of alert conditions are applied to the fabrics in a group. When any of these conditions are detected during data collection, an alert is triggered. The group that you select determines the alerts that are associated with the fabrics. For example, if you select Fabric Basic Group, nine alert conditions are automatically checked for the fabrics.
     - The fabrics are included in the group. You can use that group in other parts of the user interface, such as in reports.
   - If you select a monitoring template, the following actions occur:
     - The fabrics are automatically included in a probe schedule that is created for a template. For example, if you select the Basic Fabric Monitoring template, the fabrics are automatically included in a new probe that runs two times per week on Monday and Wednesday. The probe is created after the wizard is completed.Tivoli Storage Productivity Center determines the time when the probe is run based on server load to ensure optimal performance.

– A set of alert conditions are applied to the fabrics in a template. When any of these conditions are detected during data collection, an alert is triggered. The template you select determines the alerts that are associated with the fabrics. For example, if you select to the Basic Fabric Monitoring template, nine alert conditions are automatically checked for the fabrics.
– You must enter a prefix for the probe schedule and alerts that are based on a template. This prefix is used to identify the probe and alerts in the user interface. For example, if you select the Basic Fabric Monitoring template and enter the prefix "ABC", a probe schedule named "*userID*.ABCfabBasicProbe" and alerts such as "*userID*.ABCFabric State Change (Basic)" are created, where *userID* represents the ID of the user that is logged in toTivoli Storage Productivity Center.

  The new probe and alerts are displayed under the corresponding node in the navigation tree for you to view and edit. For example, a probe schedule created from a template in the wizard appears under the **Tivoli Storage Productivity Center >Monitoring > Probes** node; new storage system alerts appear under the **Fabric Manager > Alerting** node.

See for a list of groups and templates with their associated schedules and alerts.

7. Review the selections that you made in the wizard. The Review user selections page includes the following information:
   - A list of devices that you are configuring.
   - The name of the monitoring group that you selected. This value is not displayed if you selected a template.
   - The name of the probe schedule that is created based on the template you selected. This value is not displayed if you selected a monitoring group.
   - Information about the probe schedule created for a template.
   - The names of the alerts that are created based on the template you selected. This value is not displayed if you selected a monitoring group.
8. View the results of the device configuration and click **Finish** to close the wizard and probe the devices that you configured. The Review results page includes the following information:
   - A list of the actions performed by the wizard. The page displays a row for the successful actions. If the configuration failed, a row is displayed for each failed action. For example, if the wizard expects to assign a specific alert to five devices, but the operation only succeeds for the three of the devices, this page displays one row for the successful actions and two rows each failed actions.
   - Error messages for any failed actions.

Configure computers

Use the Configure Devices wizard to set up computers for monitoring by Tivoli Storage Productivity Center. Computers that you configure in the wizard are automatically included in probe and ping schedules and associated with alert conditions.

Use the Configure Devices wizard to configure computers that meet the following conditions:

- A Storage Resource agent has been deployed on the computer. See Storage Resource Agent Deployments for information about deploying Storage Resource agents.
- The computer is not included in an existing probe schedule.

You can configure one or more computers using a single session of the Configure Devices wizard. All the computers that you configure in a single pass through the wizard share the same probe schedule, ping schedule, and alert conditions.

To configure computers using the wizard, complete the following steps:

1. Click the configure devices icon in the toolbar or select the **Configure Devices** node in the navigation tree.
2. Select **Computers** on the Select Device Type page.
3. Select the computers that you want to configure. Computers that are already included in an existing probe schedule are not displayed.
4. Specify how data is collected for computers by including them in a monitoring group or template. Use monitoring groups and templates to determine when data is collected and what alert conditions are checked for the computers that you are configuring.
   - Select a monitoring group if you want to manage a collection of computers in the same manner. Each monitoring group that you can select is associated with existing probe and ping schedules and a set of alerts. When you select a group, its data collection schedules and alerts are automatically applied to all the computers that you are configuring. After completing the wizard, the computers remain in the group and you can use that group when working in other parts of the user interface, such as in reports.

     Tip: Tivoli Storage Productivity Center provides a number of default monitoring groups for computers. If you want to include a computer in a user-defined group rather than in a default group, that group must be associated with a data collection schedule. If the user-defined group is not associated with a data collection schedule, it is not displayed in the list of monitoring groups from which you can select.
   - Select a template if you want to create a unique probe schedule, ping schedule, and set of alerts for a computer rather than using the settings of an existing group.

   The option that you select in this step determines the field and pages that appear next in the wizard.
   - If you select a monitoring group, the following actions occur:
     - The computers are automatically included in probe and ping schedules to which the group is associated. For example, if you select Computer Basic Group, the computers are included in the Computer Basic Probe and Computer Basic Ping schedules. The Computer Basic Probe schedule runs two times per week on Monday and Wednesday at 3:00 am. The Computer Basic Ping schedule runs every 3 hours.
     - A set of alert conditions are applied to the computers in a group. When any of these conditions are detected during data collection, an alert is triggered. The group that you select determines the alerts that are associated with the computers. For example, if you select Computer Basic Group, seven alert conditions are automatically checked for the computers.
     - The computers are included in the group. You can use that group in other parts of the user interface, such as in reports.
   - If you select a monitoring template, the following actions occur:
     - The computers are automatically included in a probe and ping schedule that are created for the template. For example, if you select the Basic Computer Monitoring template, the computers are automatically included

in a new probe that runs two times per week on Monday and Wednesday and a ping that runs every three hours. The probe and ping are created after the wizard is completed. Tivoli Storage Productivity Center determines the time when the probe is run based on server load to ensure optimal performance.

– A set of alert conditions are applied to the computers in a template. When any of these conditions are detected during data collection, an alert is triggered. The template you select determines the alerts that are associated with the computers. For example, if you select to the Basic Computer Monitoring template, seven alert conditions are automatically checked for the computers.

– You must enter a prefix for the probe schedules, ping schedules, and alerts that are based on a template. This prefix is used to identify these definitions in the user interface. For example, if you select the Standard Computer Monitoring template and enter the prefix "ABC", the following definitions are created in the user interface after the wizard is completed:

  - a probe schedule named "*userID*.ABCcompStandardProbe"
  - a ping schedule named "*userID*.ABCcompStandardPing"
  - seven alerts that follow this naming convention: *userID*.ABCStatus Change Online (Standard)

  Where *userID* represents the ID of the user that is logged in toTivoli Storage Productivity Center.

  The new schedule definitions and alerts are displayed under the corresponding node in the navigation tree for you to view and edit. For example, a probe schedule created from a template in the wizard is displayed under the **Tivoli Storage Productivity Center >Monitoring > Probes** node; a ping schedule is displayed under the **Data Manager > Monitoring > Pings** node; new storage system alerts appear under the **Data Manager > Alerting** node.

See "Alerts and schedules associated with monitoring groups and templates" for a list of groups and templates with their associated schedules and alerts.

5. Review the selections that you made in the wizard. The Review user selections page includes the following information:
   - A list of devices that you are configuring.
   - The name of the monitoring group that you selected. This value is not displayed if you selected a template.
   - The name of the probe and ping schedules that are created based on the template you selected. This value is not displayed if you selected a monitoring group.
   - Information about the schedules created for a template.
   - The names of the alerts that are created based on the template you selected. This value is not displayed if you selected a monitoring group.
6. View the results of the device configuration and click **Finish** to close the wizard and probe the devices that you configured. The Review results page includes the following information:
   - A list of the actions performed by the wizard. The page displays a row for the successful actions. If the configuration failed, a row is displayed for each failed action. For example, if the wizard expects to assign a specific alert to five devices, but the operation only succeeds for the three of the devices, this page displays one row for the successful actions and two rows each failed actions.

- Error messages for any failed actions.

Configure tape libraries

Use the Configure Devices wizard to set up tape libraries for monitoring by IBM Tivoli Storage Productivity Center. Tape libraries that you configure in the wizard are automatically included in probe schedules.

You can configure one or more tape libraries using a single session of the Configure Devices wizard. All the tape libraries that you configure in a single pass through the wizard share the same probe schedule.

To configure tape libraries using the wizard, complete the following steps:

1. Access the wizard based on the task you want to perform:
   - If you want to add new or configure existing tape libraries, click the configure devices icon in the toolbar or select the **Configure Devices** node in the navigation tree. If you use this method, you must also complete the following steps:
     a. Select **Tape Libraries** on the Select Device Type page.
     b. Determine whether to add new tape libraries or configure existing tape libraries that are known to Tivoli Storage Productivity Center.
        – To add tape libraries, select **Add and configure new tape libraries** and go to step 3.
        – To configure existing tape libraries, select **Configure existing tape libraries**. The page displays the tape libraries that are known to Tivoli Storage Productivity Center. Go to step 2.
   - If you only want to add new tape libraries, go to **Tape Manager > Tape Libraries** and click the **Add Tape Library** button. Go to step 3.
2. Select the existing tape libraries that you want to configure. Tape libraries that are already included in probe schedules are not displayed. Go to step 6 after selecting the tape libraries that you want to configure.
3. Configure CIMOM connections for tape libraries. Tivoli Storage Productivity Center uses the information you enter in this step to connect to a CIMOM. Click **Add** to configure multiple CIMOMs in a session of the wizard.

   The connection information that you enter for a CIMOM is like the information that you enter on the Add CIMOM panel.
4. Discover the tape libraries that are managed by the CIMOMs you entered in the previous step. Tivoli Storage Productivity Center uses the information you enter in this step to connect to a CIMOM. Click **Add** to configure multiple CIMOMs in a session of the wizard.
5. Select the tape libraries that you want to configure for monitoring. Any tape libraries that were added in previous steps are automatically selected. Deselect tape libraries that you do not want to configure.
6. Specify how data is collected for tape libraries by including them in a monitoring group or template. Use monitoring groups and templates to determine when data is collected for the tape libraries that you are configuring.
   - Select a monitoring group if you want to manage a collection of tape libraries in the same manner. Each monitoring group that you can select is associated with an existing probe schedule. When you select a group, its data collection schedule is automatically applied to all the tape libraries that you are configuring. After completing the wizard, the tape libraries remain in the group and you can use that group when working in other parts of the user interface, such as in reports.

Tip: Tivoli Storage Productivity Center provides a number of "Alerts and schedules associated with monitoring groups and templates" for tape libraries. If you want to include a tape library in a user-defined group rather than in a default group, that group must be associated with a data collection schedule. If the user-defined group is not associated with a data collection schedule, it is not displayed in the list of monitoring groups from which you can select.

- Select a template if you want to create a unique probe schedule for a tape library rather than using the settings of an existing group.

The option that you select in this step determines the field and pages that appear next in the wizard.

- If you select a monitoring group, the following actions occur:
  - The tape libraries are automatically included in a probe schedule to which the group is associated. For example, if you select Tape Basic Group, the tape libraries are included in the Tape Basic Probe schedule. The Tape Basic Probe schedule runs one time per week on Sunday at 4:00 am.
  - The tape libraries are included in the group. You can use that group in other parts of the user interface, such as in reports.
- If you select a monitoring template, the following actions occur:
  - The tape libraries are automatically included in a probe schedule that is created for a template. For example, if you select the Basic Tape Monitoring template, the tape libraries are automatically included in a new probe that runs one time per week on Sunday. The probe is created after the wizard is completed.Tivoli Storage Productivity Center determines the time when the probe is run based on server load to ensure optimal performance.
  - You must enter a prefix for the probe schedule and alerts that are based on a template. This prefix is used to identify the probe and alerts in the user interface. For example, if you select the Basic Tape Monitoring template and enter the prefix "ABC", a probe schedule named "*userID*.ABCtapeBasicProbe" is created, where *userID* represents the ID of the user that is logged in toTivoli Storage Productivity Center.

    The new probe is displayed under the corresponding node in the navigation tree for you to view and edit. For example, a probe schedule created from a template in the wizard appears under the **Tivoli Storage Productivity Center >Monitoring > Probes** node.

See "Alerts and schedules associated with monitoring groups and templates" for a list of groups and templates with their associated schedules and alerts.

7. Review the selections that you made in the wizard. The Review user selections page includes the following information:
   - A list of devices that you are configuring.
   - The name of the monitoring group that you selected. This value is not displayed if you selected a template.
   - The name of the probe schedule that is created based on the template you selected. This value is not displayed if you selected a monitoring group.
   - Information about the probe schedule created for a template.
   - The names of the alerts that are created based on the template you selected. This value is not displayed if you selected a monitoring group.
8. View the results of the device configuration and click **Finish** to close the wizard and probe the devices that you configured. The Review results page includes the following information:

- A list of the actions performed by the wizard. The page displays a row for the successful actions. If the configuration failed, a row is displayed for each failed action. For example, if the wizard expects to assign a specific alert to five devices, but the operation only succeeds for the three of the devices, this page displays one row for the successful actions and two rows each failed actions.
- Error messages for any failed actions.

Managing schedules, jobs, and runs

Use the Job Management panel to view and manage all the schedules, runs, and jobs related to the storage devices that are monitored by IBM Tivoli Storage Productivity Center.

Specifically, you can use the Job Management panel to perform the following tasks:

- View information about all 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 all the schedules that include a particular storage device.
- Identify all the schedules, runs, and jobs that are encountering problems when executing.
- Ensure that a storage device 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 perform 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 the **View Job History** button to view the runs and jobs related to the schedule. For example, in the navigation tree, click **IBM Tivoli Storage Productivity Center** > **Monitoring** > **Probes** > **TPCUser.Default Probe** to view the probe 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, probes, pings, quotas, discovery, and performance monitors. For example, define a probe schedule to collect data about storage subsystems; define a scan schedule to collect file-level information about hosts.
- **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 subsystem probe schedule has one job created per subsystem for which it collects data.

The Job Management panel is divided into the following sections:

Filters section

Use this section to perform the following tasks:

- Determine the type of storage entities whose schedules are displayed on the panel.
- View recommendations for storage entities that are not fully monitored by Tivoli Storage Productivity Center.

Filtering the contents of this panel based on entity type is optional. You select **Show All** from the **Entity Type** menu to display all the schedules defined in Tivoli Storage Productivity Center.

Schedules section

Use this section to manage the schedules related to the entities you selected in the **Filters** section. You can perform 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, probe, 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 associated with a run. You can perform the following tasks for jobs and runs:

- View the status of a run and the date and time when it was executed.
- View information about the jobs 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.

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 storage entity 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 entity type, the following actions take place:

- The schedules, runs, and jobs related to the entity type are displayed in the appropriate sections on the panel.
- The names of the devices that belong to that entity type are displayed in the **Entity Name** list box.
- The text that is displayed above the **Show Recommendations** button indicates if the devices are fully monitored.

You can select from the following entity types in the **Entity Type** menu:

Show All
(default) View all the schedules 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.

Tape Library
View the schedules for tape libraries 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 devices contained within a particular storage resource group.

Entity Name Filter
Enter text to filter the names of entities 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 devices 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 devices such as DS8000-B34-IBM, DS8000-B36-IBM, and DS5000-B35-IBM in the **Entity Name** list box.

Text that is displayed above the Show Recommendations button
Warning text is displayed above the **Show Recommendations** button if the devices selected in the **Entity Name** list box are not fully monitored.

A device 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. For example, if using the Standard edition of Tivoli Storage Productivity Center to monitor a host, that host must appear in the following schedule definitions to be considered fully monitored: ping, probe, and scan. If using the Disk edition of Tivoli Storage Productivity Center to monitor a storage subsystem, that storage subsystem must appear in the following schedule definitions to be considered fully monitored: probe.

If a device 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 device that is not fully monitored. From the Recommendations panel, you can perform the following tasks:

- View the schedules that are recommended for a device.
- Click the **Take Action** button to create the recommended schedule for a device.

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 based on the current server load and network speed.

action drop down menu
Use the options in this menu to perform actions against the selected schedule. The menu includes **View Detail**, **Run Now**, **Stop Now** (for performance monitor schedules) and the following actions:

Rename
Rename a schedule. You must be the owner of a schedule or the Tivoli Storage Productivity Center superuser to perform this action. 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 associated with the schedule are also deleted. You must be the owner of a schedule or the Tivoli Storage Productivity Center superuser to perform this action. This action is not available for the default schedules that are owned by TPCUser.

Only show schedules with status
Filter the schedules that appear in the schedules list based on their status. 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 filter schedules using the following statuses:

- 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 selected in this menu.

schedules list

Use this list to view the schedules for the storage entities 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 associated with them. For example, jobs that are run through the command-line interface (CLI) do not have a schedule 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 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 schedule, such as Probe, Ping, Scan, Aggregator Job, CIMOM Discovery, or Subsystem Performance Monitor.

Last Run Finished
The data and time when the most recent run of a schedule completed. This column is blank if the schedule has never been 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 in the lower right corner of 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 one 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 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 specified on 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 have started within the last four weeks. Tivoli Storage Productivity Center automatically updates the runs list if you change the value selected in these fields.

Only show jobs with status

Determine the runs that are displayed in the runs list based on their status. For example, use this menu to indicate that only runs with a status of **Completed with warning(s)** are displayed. The default setting displays all runs for the selected schedule.

You can filter the display of runs using the following statuses:

- 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 selected in the **Schedules** section of the panel. This list displays runs based on the criteria 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 associated with a run. Click the name of a column to sort the contents of this list based on that column.

The following information is displayed for each run or job:

Run Status
The status of a run.

Run The run number 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 executed. For database probes, this column displays the name of the Storage Resource agent that is performing the probe. For storage subsystem probes, this column displays **localhost** because the probe is performed by the Tivoli Storage Productivity Center server.

Device
(jobs only) The name of the device to which a job applies. This column is blank if the job is not related to a specific device.

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
Appears in the lower right corner of the runs list and indicates 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 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

Use this window to view and apply the recommendations for storage devices that are not fully monitored by IBM Tivoli Storage Productivity Center.

A device 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. For example, if using the Standard Edition of Tivoli Storage Productivity Center to monitor a host, that host must appear in the following schedule definitions to be considered fully monitored: ping, probe, and scan. If

using the Disk Edition of Tivoli Storage Productivity Center to monitor a storage subsystem, that storage subsystem must appear in a probe schedule to be considered fully monitored.

Recommendation Severity
Indicates whether a recommendation is a **Warning** or **Informational.**

- A **Warning** message indicates an action that you can take to more fully use Tivoli Storage Productivity Center to monitor a device. For example, when using a probe schedule in the Standard Edition of Tivoli Storage Productivity Center to collect data about a storage system, a warning message is displayed. This message recommends that you also include the storage subsystem in a performance monitor schedule. You can then use the data 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 additional monitoring functions. For example, when using probe schedules in the Basic Edition of Tivoli Storage Productivity Center to collect data about a fabric, an informational message is displayed. This message indicates that if upgraded to the Standard Edition, Tivoli Storage Productivity Center can also collect performance data about the fabric.

For **Warning** messages, select the corresponding row and click **Take Action** to perform the recommended action.

Entity Displays the name of a storage entity to which the recommendation applies.

Recommendation
Indicates the recommended actions for a storage entity. One row is displayed per recommendation. Click **Take Action** to resolve recommendations that have a **Warning** severity.

If multiple recommendations are displayed for a device, sometimes performing one of the recommended actions might resolve the other recommendations. For example, consider if these two recommendations are displayed for a fabric:

- `In order to collect information about device configuration, define a probe schedule for fabric {0}`
- `To enable zone control, add a data source that is capable of zone control for fabric {0}`

If you click **Take Action** for the first recommendation and include the fabric in a probe, the second recommendation is also resolved. Both recommendations are then automatically removed from the **Recommendation** list

Take Action
Perform the recommended action for a warning. When you click this button, the corresponding panel for the action is displayed. For example, if the recommended action is to include a storage subsystem in a performance monitor, the Create Subsystem Performance Monitor panel is displayed.

Cancel
Exit the window without taking action.

Configuration nodes

This section is organized to help you associate the information in the IBM Tivoli Storage Productivity Center GUI with an explanation of each node.

Expand **Administrative Services** > **Configuration**. You see these nodes.

Role-to-Group Mappings

One of the first tasks to perform after installing IBM Tivoli Storage Productivity Center is to assign roles to individuals who use the product. From the **Role-to-Group Mapping** node, you can map Tivoli Storage Productivity Center roles, such as Tape Operator or Fabric Administrator, to user groups that you create in your operating system or in your LDAP-compliant repository.

For example, on a Windows computer you can create user groups using the Administrative Tools control panel. When a user name is used to authenticate with Tivoli Storage Productivity Center, the group membership for the user determines the authorization level.

Role-based authorization

Operating system groups or LDAP groups (for example, groups contained in your LDAP-compliant repository) are associated with predefined roles. 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.

Table 41 shows the association between roles and authorization level.

If you select OS authentication for your Tivoli Storage Productivity Center installation, then you do not have to create any of the groups before installation. The Tivoli Storage Productivity Center superuser role automatically gets mapped to the Administrators group on Windows, to the system group on AIX, or to the root group on Linux.

Note: If you plan to select LDAP authentication during your Tivoli Storage Productivity Center installation, then the group you intend to map to the Tivoli Storage Productivity Center Superuser role must exist in the LDAP-compliant directory before you start your Tivoli Storage Productivity Center installation.

Table 41. Roles and authorization levels

Role	Authorization level
Superuser	Has full access to all Tivoli Storage Productivity Center functions.
Productivity Center administrator	Has full access to operations in the Administration section of the GUI.
Disk administrator	Has full access to Tivoli Storage Productivity Center disk functions.
Disk operator	Has access to reports only for Tivoli Storage Productivity Center disk functions. These reports include information about tape devices.
Fabric administrator	Has full access to Tivoli Storage Productivity Center Fabric Manager functions.
Fabric operator	Has access to reports only for Tivoli Storage Productivity Center Fabric Manager functions.

Table 41. Roles and authorization levels (continued)

Role	Authorization level
Data administrator	Has full access to Tivoli Storage Productivity Center Data Manager functions.
Data operator	Has access to reports only Tivoli Storage Productivity Center Data Manager functions.
Tape administrator	Has full access to Tivoli Storage Productivity Center tape functions.
Tape operator	Has access to reports only for tape functions.

Notes:

1. If a user has multiple roles, the authorization level is a combination of the levels for each of the roles.
2. If a user is not a member of any of the roles listed, no access is granted to that user.
3. For enterprise-rollup reports, you need superuser or Tivoli Storage Productivity Center 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 user can generate enterprise-rollup reports.
4. When you create and save role-to-group mappings in Tivoli Storage Productivity Center, these mappings get propagated into Tivoli Integrated Portal, where the groups are given the operator authorization. Occasionally, you might find that after creating and saving the role-to-group mappings in Tivoli Storage Productivity Center, you are unable to access Tivoli Integrated Portal as a valid user (in a valid group that is mapped to the operator authorization in Tivoli Integrated Portal). If this condition occurs, stop and restart the Tivoli Integrated Portal server.

How to associate a user and group with IBM Tivoli Storage Productivity Center

This section provides information about how to associate a user and group with IBM Tivoli Storage Productivity Center.

Creating a system group on Windows:

You must first create an operating system group on your computer, and then assign the users to the group. Then you can use the GUI to associate the group with the IBM Tivoli Storage Productivity Center role.

For Windows, first create the users you want to have access to the various Tivoli Storage Productivity Center roles. Complete the following steps:

1. Go to **Start** > **Settings** > **Control Panel** > **Administrative Tools** > **Computer Management** > **Local Users and Groups**.
2. Right-click **Users** and click **New Users**.
3. The New User panel is displayed. Enter a user name, description (optional), and password (twice). Click **Create**.

4. Repeat steps 3 and 4 for each new user you want to create. Then create a system group on Windows:
   a. Go to **Start** > **Settings** > **Control Panel** > **Administrative Tools** > **Computer Management** > **Local Users and Groups**.
   b. Right-click **Groups** and click **New Group**.
   c. The New Group panel is displayed. Enter a Group Name and Description (optional). Click **Add**.
   d. The Select Users or Groups panel is displayed. Enter or select the users or groups you want to add to the new group you have created. Click **Add**. Then click **Create**.

Creating a system group on UNIX or Linux:

For UNIX or Linux, to add new users to the system, follow the steps in this topic.

Complete the following steps:

1. Create a name for the user.
2. The new user must have a home directory allocated for them.
3. That user must be placed in the group file. Refer to your UNIX or Linux documentation for information about how to create a system group.

Associating the user groups to IBM Tivoli Storage Productivity Center roles:

After you have created the operating system group and assigned users to that group, you must assign the groups to IBM Tivoli Storage Productivity Center roles.

Complete the following steps:

1. Open the Tivoli Storage Productivity Center GUI.
2. Go to **Administrative Services** > **Configuration** > **Role-to-Group Mappings**.
3. On the content pane, click **Edit** for the role you want to associate with the group you have created.
4. The Edit Group dialog is displayed. Enter the group you want to associate with the Tivoli Storage Productivity Center role. Click OK. This associates the operating system group with the Tivoli Storage Productivity Center role.
5. Click **File** > **Save** to save the mappings.

License Keys

This section provides information about how to administer the database licenses for Data Manager.

You must have the following permissions set to monitor the databases with Data Manager:

Table 42. Database permissions

Database	Permissions
DB2	db2admin
Microsoft SQL Server	public
Oracle	DBA
Sybase	SA level

Note: To monitor Oracle databases with IBM Tivoli Storage Productivity Center, the Oracle user needs DBA authority. A non-DBA Oracle user with "create session", "select any dictionary", "analyze any", and "analyze any dictionary" roles can still monitor the Oracle database through Tivoli Storage Productivity Center. However, Tivoli Storage Productivity Center is not able to obtain free space information for database objects. During the database registration, when a non-DBA Oracle user is used with the roles mentioned in the preceding list, the following warning message displays:

```
RDBMS login does not have dba privilege. Freespace will not be calculated for
system objects.
```

Clicking **OK** continues with the registration process. This message also appears in the scan logs when scan jobs are run.

Through the License Keys node, you can administer the license keys for Data Manager for Databases. Specifically, you can:

- Assign Data Manager for Databases licenses to your Data agents or Storage Resource agents.
- Edit, add, and delete the instances within your organization that you want to monitor.
- View the number of Data Manager for Databases licenses.
- View the number of Data Manager for Databases licenses that are not currently assigned (unused) to agents.

Before you can use Data agents or Storage Resource agents to manage the storage for your instances, you must do the following steps:

1. Assign Data Manager and Data Manager for Databases licenses to the agents that are monitoring RDBMS instances.
2. Register the instances on the machines that contain licensed agents.

Assigning Data Manager database licenses to installed agents

This topic describes how to assign Data Manager database licenses to an agent.

To assign a Data Manager database license to an agent, complete the following steps:

1. Expand **Administrative Services** > **Configuration** > **License Keys**. The License Editor window is displayed.
2. Click the icon to the left of **IBM Tivoli Storage Productivity Center for Data - Databases**. The Tivoli Storage Productivity Center for Data - Databases License Editor window is displayed. The field and button descriptions for the Licensing tab are:

   Select All
   Click this button to select all the **Licensed** boxes.

   Deselect All
   Click this button to remove all the licenses for Data Manager - Databases.

   Computer
   Displays all the computers on which a Data Manager agent is installed.

   OS Type
   Displays the operating system of the machine where the agent is installed.

Domain
Displays the domain of a machine where the agent is installed.

Tree Name
For NetWare only.

Licensed
Contains a check box that indicates whether a computer is licensed for use with Data Manager - Databases.

If you have unused Data Manager - Databases licenses, continue to the next step. To unassign licenses from agents, see "Unassigning Data Manager - Databases license" on page 298.

3. Click the check box in the **Licensed** column next to the computer containing the instance you want to monitor.
4. Click **File** > **Save** to save the updated license settings. If an instance is registered on the machine where you licensed the agent, you are ready to set up your Data Manager - Databases jobs to monitor its storage. If an instance is not registered on the machine where you licensed the agent, see "Registering instances on machines that contain licensed agents."

Registering instances on machines that contain licensed agents

This topic describes how to register instances on machines that contain licensed agents.

1. Click the **RDBMS Logins** tab on the IBM Tivoli Storage Productivity Center for Data - Databases License Editor window.
2. The RDBMS Logins window is displayed. Use this window to edit, add, and delete the instances within your organization that you want to monitor. The field and button descriptions for the RDBMS Logins tab are:

   Edit Highlight a row and click this button to edit the login information for the corresponding instance. The RDBMS Login Editor window is displayed.

   Add New
   Click this button to add login information for a new instance. The RDBMS Login Editor window is displayed.

   Delete Highlight a row for an instance and click this button to delete that instance from Data Manager - Databases. Once you delete an instance, all the previously gathered statistics for that instance are automatically deleted from the database repository and that instance is no longer available for selection in the reporting section of the navigation tree.
3. Click **Add New** to add a new instance. The RDBMS Login Editor window is displayed. Use this window to enter information about the instance that you want Data Manager - Databases to monitor.

   Note: In an Oracle Parallel Server (OPS) environment, you only need to register one of the instances within that environment. OPS is a resource-sharing system that increases availability and performance by partitioning the workload across multiple servers of a cluster (nodes). Databases installed on clustered servers or clustered database servers (for example, Oracle Real Application Cluster (RAC) environment) are not supported for monitoring.
4. Select the name of the machine where the instance is running from the **Host Name** list box.

   Note: You can only register instances on machines that contain licensed agents.
5. Enter the following information in the next field:

For Oracle
Enter the Oracle SID and host for the instance.

For Microsoft SQL/Server
Enter the name of the instance you want to register in the **Instance** field.

Note: When you use the Storage Resource agent to monitor the RDBMS, the JDBC field is not shown. The Storage Resource agent does not require this field.

For Sybase
Enter the server name in the **Server** field.

For IBM UDB
Enter the name of the instance you want to register in the **Instance** field.

6. Enter a user ID that has the appropriate database privileges within the instance in the **User** field. See "License Keys" on page 293.

   For Oracle
   The following privileges are required for this user ID:
   - CREATE SESSION
   - SELECT ANY DICTIONARY
   - ANALYZE ANY

   Note: For Oracle 9i or 10g, specify ANALYZE ANY DICTIONARY

   For Microsoft SQL Server
   The login ID that Data Manager uses to log in to Microsoft SQL Server instances that you want to probe must have "permit" access.
7. Enter a password for the user ID in the **Password** field.
8. Enter the port on which the instance is listening in the **Port** field.

   For Oracle
   The default port is 1521.

   For Microsoft SQL/Server
   The default port is 1433. You must also provide the fully qualified path to the JDBC driver in the **JDBC Driver** field.

   Note: When you use the Storage Resource agent to monitor the RDBMS, the JDBC field is not shown. The Storage Resource agent does not require this field.

   For Sybase
   The default port is 5000.

   For IBM UDB
   When monitoring multiple UDB instances within your environment, you must ensure that the port numbers you choose are open (unique for each instance) for JDBC and Java connections to those instances. To open up a port, run the following DB2 command on the machine where the instance is located:

   ```
   db2jstrt
   ```

   The default port number is 6789. You can change this default by indicating the port number you want to use when running the **db2jstrt** command. For example:

```
db2jstrt 6790
```

If you enter an incorrect port number, an error occurs. To verify the port number for an instance, do the following steps:

UNIX or Linux
View the `etc/services` file to confirm the correct port number.

AIX Run the following command:

```
ps -ef | grep db2jd
```

The output from this command indicates the port on which the instance is listening. Use this port number when running the **db2jstrt** command.

Windows
Use the information in the IBM DB2 configuration tools to confirm the port number.

9. Click **File** > **Save** to save the instance configuration information. You can now run a probe job against the registered instances. You must run a probe job against an instance before you can select any databases or table spaces against which you want to run a scan job.

Configuring Microsoft SQL Server 2005 or 2008:

Before you can monitor a Microsoft SQL Server 2005 or Microsoft SQL Server 2008 database, you must make some configuration changes to the Microsoft SQL Server.

Note: Before registering an instance of Microsoft SQL Server 2005 or 2008 to be monitored by a Storage Resource agent, please make sure that the directory containing the sqlcmd utility (sqlcmd.exe) is in the system PATH environment variable on the Microsoft SQL Server system. The default location of sqlcmd is `C:\Program Files\Microsoft SQL Server\<version>\Tools\Binn`.

Before monitoring the Microsoft SQL Server database, follow these steps:

1. Install the Microsoft SQL Server and provide the required information about the installation panels. See the Microsoft SQL Server 2005 Installation and Configuration Guide for detailed information. For the installation and configuration guide, see http://msdn.microsoft.com/en-us/library/ms143516.aspx.
2. For Microsoft SQL Server 2005 and 2008, make sure that the Microsoft SQL Server is using Mixed Mode authentication.
3. After the installation, go to the SQL Server Configuration Manager and make sure that you set the Dynamic TCP/IP port to the default port 1433. (See the Microsoft SQL Server 2005 Installation and Configuration Guide for detailed information.) To, configure the Microsoft SQL Server, follow these steps:
   a. Open the SQL Server Configuration Manager.
   b. Go to SQL Server 2005/2008 Network Configuration.
   c. Select **Protocols for MSSQLINST** (name of the instance).
   d. Right-click **TCP/IP**.
   e. Select Enable: **Yes**.
   f. Go to the **IP Addresses / IP All** and add **TCP Dynamic Ports : 1433** (default port).
4. Launch the Microsoft SQL Server setup.exe.

5. Go to **Installation/Search** for product updates. You are redirected to the Microsoft update Web site that scans the computer for the components that need to be updated.
6. When finished, select **Express Install** to install the components found.
7. Install the Microsoft CoreXML Services (MSXML) 6.0 SP 2. Reboot the system.
8. **For Microsoft SQL Server 2008 only**: Install the Microsoft SQL Server 2008 Management Objects from the following Microsoft SQL Server 2008 Feature Pack link: http://www.microsoft.com/downloads/. Search for **Management Objects**.
9. **For Microsoft SQL Server 2008 only**: Install the Microsoft SQL Server 2005 Backward Compatibility Components from the link specified in step 8. Search for **2005 Backward Compatibility**.
10. Install the Storage Resource agent and JDBC driver on the Microsoft SQL Server system.
11. Open the RDBMS Login Editor. Expand **Administrative Services > Configuration > License Keys**. In the content pane, click **RDBMS Logins** tab.
12. Click **Add New**. The RDBMS Login Editor opens. Enter the following information:

    Database
    Microsoft SQL/Server.

    Instance name
    Name of the instance (`mssqlinst`).

    User User ID to logon to the Microsoft SQL Server.

    Password
    Password for the user ID.

    Port 1433

    JDBC Driver
    Enter the JDBC driver. For Microsoft SQL Server 2005 and 2008, open the Microsoft SQL Server 2005/2008 JDBC Driver package and install the driver to the wanted location. After the package is unpacked, in the RDBMS logon panel, enter the absolute path to the JDBC Driver (suggested version 1.2):

    ```
    %MS_SQL_DRIVER_HOME%/sqljdbc_1.2/enu/sqljdbc.jar
    ```
13. Click **Save**.

Unassigning Data Manager - Databases license

This topic describes how to unassign a Data Manager - Databases license.

To unassign a license, complete the following steps:

1. Expand **Administrative Services** > **Configuration** > **License Keys**.
2. The License Editor window is displayed.
3. Click the icon to the left of **IBM Tivoli Storage Productivity Center for Data - Databases**. The IBM Tivoli Storage Productivity Center for Data - Databases License Editor window is displayed. Use the **Licensed** column in this window to view the agents to which licenses are currently assigned.
4. Clear the **Licensed** check box next to a machine to remove the license for the agent on that machine.
5. Click **File** > **Save**. When you remove the license for an agent, the following actions occur:

- All the data gathered by that agent for the instance it monitors is removed from the database repository.
- You can no longer run monitoring, alerting, or policy management jobs against the instance on the machine where the agent was located.
- The number of unused licenses increases by one.

Alert disposition

You can configure IBM Tivoli Storage Productivity Center so that SNMP traps and Tivoli Enterprise Console events can be sent to other consoles.

Configuring Tivoli Storage Productivity Center to send SNMP traps and IBM Tivoli Enterprise Console events

This topic provides information about how to configure Tivoli Storage Productivity Center to send SNMP traps, Tivoli Enterprise Console events, or email notifications. These actions are taken when an event specified in an alert is detected. From this node, you can also specify the number of days after which records in the alert log are deleted.

To enable SNMP traps and Tivoli Enterprise Console event alert notifications, consider the following information:

SNMP traps

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.

TEC events

Send an alert to the Tivoli Enterprise Console. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events. The event definitions are specified in the `tivoliSRM.baroc` and `fabric.baroc` files. It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

The `tivoliSRM.baroc` and `fabric.baroc` files must be loaded into the current active rule base of the Tivoli Enterprise Console server. This action allows Tivoli Enterprise Console events sent by Tivoli Storage Productivity Center to display on the console.

Table 43. Location of MIB and Baroc files

Component	MIB file location	Baroc file location
Data Manager and Disk Manager	• Installation media, disk1, part 1 • Data server: `data\snmp\tivoliSRM.mib`	• Installation media, disk1, part 1 • Data server: `data\tec\tivoliSRM.baroc`
Fabric Manager	• Installation media, disk1, part 1 • Device server: `device\snmp\fabric.mib`	• Installation media, disk1, part 1 • Device server: `device\tec\fabric.baroc`
Tivoli Storage Productivity Center for Replication	• Installation media, disk1, part 2 • `tpcr\TPCRM\CSM-Client\etc\ibm-TPC-Replication.mib`	• Installation media, disk1, part 2 • `tpcr\TPCRM\CSM-Client\etc\ibm-TPC-Replication.baroc`

Table 43. Location of MIB and Baroc files (continued)

Component	MIB file location	Baroc file location
Note: • The Tape Manager displays only indications received from the tape CIM agent in the tape alert log. Because the tape alerts cannot be configured, there is no forwarding of SNMP traps and there are no MIB or baroc files for the Tape Manager. • Some alerts are generated by Tivoli Storage Productivity Center about Tivoli Storage Productivity Center for Replication. Examples of these alerts are communication failures between Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication. Other alerts generated are Tivoli Storage Productivity Center for Replication configuration changes. These alerts can be forwarded by Tivoli Storage Productivity Center as SNMP traps or through email. These alerts are generated directly by Tivoli Storage Productivity Center and not by SNMP traps from Tivoli Storage Productivity Center for Replication. If Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication are running on the same system, only one program can listen on the SNMP default port. All traps received on the default port are forwarded to the other server process using a different port number.		

Port 162

Tivoli Storage Productivity Center uses port 162 to listen for SNMP traps. This port is the default port. To receive SNMP traps from switches, you must configure the switch to send SNMP traps to the Device server's IP address and port.

Note: Port 162 is also used by the Windows SNMP Trap service. If the Windows SNMP Trap service is running, stop the service or change the default port that Tivoli Storage Productivity Center uses to listen for SNMP traps. If you need to change the Device server's default port number, use the **setdscfg** command. The attribute to set is **SNMPTrapPort**.

To configure Tivoli Storage Productivity Center to send traps and Tivoli Enterprise Console events, follow these steps:

1. Open the Tivoli Storage Productivity Center GUI.
2. Go to **Administrative Services** > **Configuration** > **Alert Disposition**.
3. On the content pane, you can enter information for two SNMP destinations and one Tivoli Enterprise Console event destination. For the SNMP destination, enter the community name, host, and port information. For the Tivoli Enterprise Console event destination, enter the server and port information. You can also enter information for email information. Click **File -> Save**.

Log File and Cached Batch Report Retention

You can set the retention period for log files.

You can specify the number of runs and the length of time to maintain the log files generated by IBM Tivoli Storage Productivity Center. We recommend that you set the value for **Maximum number of runs to keep of each schedule** to be less than or equal to the value in the **Maximum number of days' worth of log files to keep (regardless of schedule)** field. This setting helps avoid the situation where entries for schedule runs are displayed in the user interface, but the corresponding log files are deleted.

Note: By default, job log files for schedules are retained for 90 days, but you can change that value on the Log File and Cached Batch Report Retention panel. However, for CLI and Event Driven Jobs schedules, job log files are automatically retained for 90 days.

To set the retention period, complete the following steps:

1. Go to **Administrative Services** > **Configuration** > **Log file and Cached Batch Report Retention**.
2. In the **Maximum number of runs to keep of each schedule** field, enter the maximum number of runs for a schedule that you want to display in the user interface. The default is five runs per schedule. For example: if you enter **10** in this field and then run a probe schedule 11 times, job log and job information entries for the first run are deleted from the database repository. Additionally, only the 10 most recent runs of the schedule are displayed in the Job Management panel.
3. In the **Maximum number of days ' worth of log files to keep (regardless of schedule)** field, enter the maximum number of days to keep a log file for a run. A log file is generated for each job in a run and is stored in the appropriate log directory. Every log file that is older than the number of days you specify here is deleted. The default is 90 days.
4. In the **Cached Batch Report Retention** field, enter the maximum number of days to keep a batch report. The default is 90 days.

   Note: When the batch report is created on the Tivoli Storage Productivity Center server and copied to the Storage Resource agent, the copying of the report to the agent might fail. If a failure occurs, no attempt is made to send the report again. Instead, the report remains cached on the Tivoli Storage Productivity Center server. You can configure the time the report remains cached on the Tivoli Storage Productivity Center server.

Quota and Constraint e-mail Address Rules

You can specify rules for generating email addresses of users who break the quota and constraint email address rules based on their user ID, given name, or family name as they are registered within the operating system.

The user names are obtained as follows:

- For Windows: Full name field, from LDAP.
- For NetWare: `Surname` and `Given name` fields, from LDAP.
- For UNIX or Linux: User description from the Password file.

To set the rules, complete the following steps:

1. Expand **Administrative Services** > **Configuration** > **Quota and Constraint e-mail Address Rules**.
2. Click **Add After** or **Add Before** to include elements within an email address template. An email address template describes how to build the user ID. The user ID is notified in the event of a quota violation. This user ID represents the actual user who breaks the quota rules.
3. Select **USERNAME**, **FIRSTNAME**, **LASTNAME**, **Text**, or **Substring** from the pop-up menu to include as an element in the email address rule:
   - **USERNAME**: the login ID of the user who breaks the quota or constraint rules.

- **FIRSTNAME**: the given name of the user who breaks the quota or constraint rules.
- **LASTNAME**: the family name of the user who breaks the quota or constraint rules.
- Text: free form text that you want to appear within the email address.
- Substring: an element in the email address that is a substring of **USERNAME**, **LASTNAME**, or **FIRSTNAME**. For example,

  ```
  LASTNAME + SUBSTRING(USERNAME, 0, 3)
  ```

  includes the first three characters of the **USERNAME**. If the family name is Smith and **USERNAME** (as defined by file owner information) is 9A0723, then this substring example is equal to "SMITH9A0".

  To define a substring, complete the following steps:

  a. Select **Substring** after clicking **Add After** or **Add Before**.
  b. Select the substring variable: **USERNAME**, **LASTNAME**, or **FIRSTNAME**.
  c. Highlight the range for the substring.
  d. Click **OK**. That substring appears within the email address template.
  e. Click **File** > **Save** to save the email address rule.

Scan/Probe Agent Administration

You can assign Data agents or Storage Resource agents to run scan and probe jobs.

Assign the Data agents or Storage Resource agents to perform scans against the following objects:

- Volumes and file systems within the licensed NetWare servers of an NDS tree (only the Data agents support the NDS tree)
- File systems within NAS filers
- IBM Tivoli Storage SAN File Systems

The window associated with this node provides a complete listing of the licensed NetWare volumes, NAS filers, and SAN File System discovered by Data Manager.

Manual NAS/NetWare Server Entry

You can configure IBM Tivoli Storage Productivity Center to support Network Attached Storage or NetWare.

Use the Manual NAS/NetWare Server Entry node to:

- Manually enter information about the NetWare and Network Attached Storage (NAS) servers that you want to monitor within your environment. After you enter information about the NetWare or NAS servers, you can assign agents to those servers on the Scan/Probe Agent window.
- View a list of the NetWare server and NAS filers whose information was manually entered into Data Manager.
- Delete NetWare servers and NAS filers whose information entered into Data Manager.

You can manually set up individual NetWare and NAS servers for monitoring by Data Manager using this window, or you can use a "discovery" method for automatically adding multiple servers at the same time.

We recommend using the "discovery" method when working in large environments where you want to add multiple NAS Filers and NetWare servers for monitoring. Use the Manual NAS/NetWare Server Entry window when you want to manually add individual NAS Filers and NetWare servers for monitoring.

For more information about configuring NAS and NetWare, see http://www.redbooks.ibm.com/. Search for **sg247490**.

Manually adding a NAS filer or gateway

To manually add a NAS filer or gateway:

1. Expand **Administrative Services** > **Configuration** > **Manual NAS/Netware Server Entry**.
2. Click the **Add NAS Server** button. The Add NAS Server window is displayed.
3. Enter the following information:

   Network name

   Enter the network name of the NAS server you want to add.

   When manually adding a NAS Gateway or Filer that is monitored by a Data agent or Storage Resource agent on a UNIX computer, you must add the NAS using the same name that was used when mounting its file systems on that UNIX computer. You can mount file systems using the short name, fully qualified name, or IP address of a NAS.

   For example, if the file systems from a NAS Gateway were mounted to a UNIX computer (where the Data agent or Storage Resource agent is installed) using the short name of the NAS Gateway, you must add the NAS Gateway in IBM Tivoli Storage Productivity Center using the short name of the NAS Gateway. If the file systems from a NAS Gateway are mounted using an IP address, you must add the NAS Gateway in Tivoli Storage Productivity Center using the IP address as its name.

   Consider the following example:

   a. The NAS filer named "oxide" was mounted on the UNIX machine where a Data agent or Storage Resource agent is located. The following commands used the short name and IP address to identify the NAS filer during a file system mount:
      - oxide:/vol/john1 81920 22108 74% 123 5% /n3700_john1
      - 192.168.1.65:/vol/vol0 47566060 28715228 40% 1016085 52% /n3700_vol0

   b. To add this NAS filer to Tivoli Storage Productivity Center, enter the following values in the **Network Name** field:
      - oxide
      - 192.168.1.65

   If file systems from a NAS Gateway or Filer is mounted on a UNIX computer in different ways, you must add that NAS to Tivoli Storage Productivity Center with the names used in both methods. For example, if one file system is mounted using the fully qualified name of a NAS Gateway and the other file system is mounted using the IP address, you must add that NAS Gateway to Tivoli Storage Productivity Center twice: once with the fully qualified name and once with the IP address. Only one row appears for this NAS on Manual NAS/Netware Server Entry panel, but both file systems are listed on the Scan/Probe Agent Administration panel.

Data Manager Agent OS Type
Select the operating system of the computer that contains the agent that gathers information about the NAS filer.

Accessible from
Select the agent that you want to use to "discover" the NAS filer. This list box only displays agents that are:

- Running under the operating system selected in the **Data Manager Agent OS Type** field.
- Located on Windows or UNIX computers that are accessible to the NAS filers (Data Manager's agents are not located on the NAS filers themselves):
  - Windows: agents are located on Windows computers within the same domain as the NAS filers.
  - UNIX: agents are located on UNIX or Linux computers that have NFS imports for the file systems within the NAS filers.

SNMP Community
Enter the name of the SNMP communities that Data Manager uses when communicating with machines in your environment. If you do not enter the name of an SNMP community, the default community **public** is used. Data Manager uses the SNMP protocol to contact and identify NAS filers. This field is optional.

Login ID
(Windows only.) Enter the login ID that Data Manager uses when logging in to the NAS filer.

Password
(Windows only.) Enter the password that Data Manager uses when logging in to the NAS filer.

Add as Other NAS
Select this check box to add a NAS server as **Other NAS** filer. With this option, you can monitor and report on file system information about the NAS filer or gateway through Windows CIFS or UNIX NFS shares accessible to the scan or probe job for the agent. No controllers, disks, and logical volumes information are collected or reported.

NAS Server Vendor Name
Enter the vendor name (or manufacturer) hosting the file system of the NAS server. The default vendor is Network Appliance.

4. Click **OK** to have the Data Manager verify the filer for which you entered information. During this verification, Data Manager will:
   - Log in to the NAS filer.
   - Gather information about the file systems visible on those filers to the agent. For UNIX or Linux, it gathers information about the file systems that it can actually see (for example, file systems that are mounted to the UNIX boxes). By default, file systems are discovered at the root. For Windows, it finds all the NAS filers that are visible through CIFS.
   - Determine which file systems are visible to which agents.
   - Enter the NAS server information into the repository.
5. Expand **Administrative Services** > **Configuration** > **Scan/Probe Agent Administration**. This window allows you to assign agents to each file system of the NAS server.

Note: At any time, you can change the login ID and password for a NAS filer on the **Administrative Services** > **Configuration** > **License Keys** > **Filer Logins** window.

Deleting a manually added NAS filer

To delete a NAS filer whose information was manually entered into Data Manager, follow these steps:

1. Expand **Administrative Services** > **Configuration** > **Manual NAS/Netware Server Entry**.
2. Highlight a row that represents the device you want to delete.
3. Click **Delete**. Remember that when you delete a device from this window, all information about that device is removed from the repository.

Manually adding a NetWare server

To manually add a NetWare server, complete the following steps:

1. Expand **Administrative Services** > **Configuration** > **Manual NAS/Netware Server Entry**.
2. Click the **Add Netware Server** button. The Add Netware Server windows is displayed.
3. Enter the following information:

   NDS Trees
   Select an NDS Tree. The NDS Trees that appear in this list were discovered and saved to the repository by probe jogs that were automatically run upon agent installation.

   Note: Only the Data agents support NDS Trees.

   Network Name
   Enter the network name of the NetWare server you want to add.
4. Click **OK** to have the Data Manager verify the server for which you entered information. During this verification, the Data Manager gathers information about the volumes and file systems within the specified NetWare server.

Deleting a manually added NetWare server

To delete a NetWare server whose information was manually entered into Data Manager, complete the following steps:

1. Expand **Administrative Services** > **Configuration** > **Manual NAS/Netware Server Entry**.
2. Highlight a row that represents the device you want to delete.
3. Click **Delete**. Remember that when you delete a device from this window, all information about that device is removed from the repository.

Editing Data Manager configuration files

Tivoli Storage Productivity Center provides you with the ability to edit Data Manager configuration files to further customize the settings for a component according to the standards at your site.

Data Manager has configuration files for customizing the operation of the server and agent components within your environment. These files reside in the *TPC_install_directory*/config/ directory, where *TPC_install_directory* represents the directory where you installed the product.

The configuration file for the Tivoli Storage Productivity Center server is located in the following default installation directory:

Windows
```
C:\Program files\IBM\TPC\Data\config\
```

UNIX or Linux
```
/opt/IBM/TPC/Data/config
```

The agent configuration file for the Storage Resource agent is located in the following default installation directory:

Windows
```
C:\Program Files\IBM\TPC\agent\config
```

UNIX or Linux
```
/opt/IBM/TPC/agent/config/
```

The agent configuration file for the Data agent is located in the following default installation directory:

Windows
```
C:\Program files\IBM\TPC\ca\subagents\TPC\Data\
```

UNIX or Linux
```
/opt/IBM/TPC/ca/subagents/Data/config/
```

Note for UNIX: When you change the configuration files for the server component, you must stop and start the server before those changes take effect.

Edit the `agent.config` file to configure the Data Manager agents in your environment. This file is located in the agent installation directory on every computer where an agent is installed. For information about the `agent.config` file, go to the information center. Search for *agent.config file*.

Edit the `nas.config` file to configure the Data Manager NAS feature for your environment.

The `nas.config` file contains the following:

- On each line not beginning with #, the first blank-delimited field must contain the SNMP Enterprise code of a NAS filer that the agent discovers, probes, or scans.
- The second field contains identifying information about the filer.

Any remote host that cannot be reached by SNMP or whose enterprise code does not match one of these values is ignored.

For information about the server files to configure, go to the information center. Search for *server.config file*, *scheduler.config file*, and *TPCD.config file*.

For information about the `agent.config` file, go to the information center. Search for *agent.config file*.

Editing the NAS configuration file

This topic provides information about editing the `nas.config` file for the Data Manager NAS feature.

Edit the `nas.config` file to configure the Data Manager NAS feature for your environment.

The `nas.config` file contains the following information:

- On each line not beginning with #, the first blank-delimited field must contain the SNMP Enterprise code of a NAS filer that the agent discovers, probes, or scans.
- The second field contains identifying information about the filer.

Any remote host that cannot be reached by SNMP or whose enterprise code does not match one of these values are ignored.

Agent Manager Registration

Use the Agent Manager Registration page to register the Device server and Data server with the Agent Manager.

Note: This procedure does not install the Agent Manager. For information about installing the Agent Manager, see the Tivoli Storage Productivity Center 4.1.1 Information Center.

To register the Device server and Data server with the Agent Manager, follow these steps:

1. Navigate to **Administrative Services > Configuration > Agent Manager Registration**.
2. Click **Agent Manager Registration** in the content pane.
3. The Agent Manager Registration window is displayed. Enter the following information:

   Hostname or IP address
   Enter the host name for the Agent Manager. If you specify a host name, use the fully qualified host name. It is preferable to use the fully qualified host name rather than an IP address.

   Port (Public)
   The public port for the Agent Manager. The default port is 9513.

   Port (Secured)
   The secured port for the Agent Manager. The default port is 9511.

   User ID (for Data server and Device server registration)
   This is the user ID used to register the Data server and the Device server with the Agent Manager.

   Password (for Data server and Device server registration)
   This is the password used to register the Data server and the Device server with the Agent Manager.

   Password (Common agent registration password)
   Enter the agent registration password. This is the password used to register the common agents with the Agent Manager. The agent registration password was specified when you installed the Agent Manager. The default password is **changeMe**.
4. Click **OK**.
5. If the operation is successful, the Device server and Data server will be registered with the Agent Manager and you will see the information in the following table:

Field Name	Description
Agent Manager	IP address of the Agent Manager.

Field Name	Description
Port (Public)	The public port for the Agent Manager. The default port is 9513.
Port (Secured)	The secured port for the Agent Manager. The default port is 9511.
User ID	The user ID that the Device server and Data server used to register with the Agent Manager.
Unique ID	This is the unique ID returned by the Agent Manager.

History Aggregator

You can configure reports for data aggregation.

The History Aggregator in IBM Tivoli Storage Productivity Center defines and runs jobs to sum data in the database repository for historical reporting purposes. For example, you can view the sum of usage across multiple storage resources, by filesystem. Trending enables you to see patterns of your historical data across your entire network.

You have the option of turning aggregation off, although this action is not recommended. To turn off aggregation, access the History Aggregator window, clear the **Enabled** check box, and select **File** > **Save**.

To configure reports for data aggregation, complete the following steps:

1. Expand **Administrative Services** > **Configuration** > **History Aggregator**.
2. The Edit History Aggregator panel is displayed in the content pane. You can specify the following information:
   - How often to run the job
   - How to handle time zones
   - Triggering conditions
   - Triggered actions

Storage Resource agent deployments

Use this page to manage Storage Resource agent deployments.

Deploying Storage Resource agents

Deploy Storage Resource agents through the user interface rather than a separate installation wizard. You can have only one type of agent per host that points to the same server. For example, if you install a Storage Resource agent, and then later install a Data agent on the same host that points to the same Data server, the Storage Resource agent will automatically get uninstalled after the Data agent runs a successful probe schedule. See “Deployment considerations for Storage Resource agents” on page 312 for information about the features available for Data agents and Storage Resource agents.

You can enter host names from a Microsoft directory, enter host names manually, or import a host list file. You can define a schedule to add the Storage Resource agents at a time that is convenient for you. You can also specify the type of alerts to generate if a schedule fails. Once you enter information for the Storage Resource

agent, you can optionally validate the connection to the Data server. This step helps to eliminate most of the possible failure cases once you have submitted the deployment schedule.

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

To deploy Storage Resource agents, follow these steps:

1. See "Deployment considerations for Storage Resource agents" on page 312 for a list of considerations before deploying Storage Resource agents. If you are deploying Storage Resource agents on UNIX or Linux systems, you must use root as the user ID.
2. Open the Tivoli Storage Productivity Center user interface.
3. In the navigation tree pane, expand **Administrative Services > Configuration > Storage Resource agent Deployments**. Right-click **Storage Resource agent Deployments** and click **Create Storage Resource agent Deployments**.
4. The Create Storage Resource agent Deployment window opens displaying the Computers tab. Enter the following information:

   Top pane

   Creator
   Pre-defined creator name.

   Name Pre-defined name.

   Description (optional)
   Enter a user-defined description name for the schedule.

   Enabled
   Check this box to enable Storage Resource agent deployments.

   Add Host List
   Click this button to manually enter names and login credentials for the computers on which you want to deploy Storage Resource agents to.

   If you click this button, the Login Information window opens. You can enter information in 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.
   - Click **Get Agent List From file** to install Storage Resource agents on one or more computers listed in a file. The computers listed in the file must share the same administrative user ID and password.
   - Enter the computer names or IP addresses in the Remote Agent Machines table.
   - Enter the installation locations for the agents.

     Note: The default path for Linux and UNIX is /opt/IBM/TPC/ in the subfolder agent. The default path for Windows is C:/Program Files/IBM/TPC in subfolder /agent. The path is automatically provided by the GUI if you do not supply any value and have

checked the box for **Validate before save**. Otherwise, if you enter a path, the agent is installed in that path.

- 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 perform this action, you must create the deployment schedule from the additional Tivoli Storage Productivity Center server to which you want the Storage Resource agent to communicate.

  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 schedule.
  - 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.

Enter the following information:

User The user ID that Tivoli Storage Productivity Center uses when logging in to the computers listed in the **Remote Agent Machines** section. The value in this field is applied to all the computers that are included in the deployment schedule. If a computer listed in the **Remote Agent Machines** section uses a different user ID than what you specify in this field, you must create a separate deployment schedule for that computer.

Password
Password for the computers on which to deploy Storage Resource agents.

Certificate Location
For Storage Resource agents, certificates are required for SSH protocol communication between the server and agent. This certificate needs to be stored on the server. See “Creating a certificate for SSH protocol” on page 317.

Passphrase
For Storage Resource agents, enter the passphrase used for SSH protocol communication. See “Creating a certificate for SSH protocol” on page 317.

Port Port for the Storage Resource agent. The default is 9510.

Use Daemon Service for Runtime Operation
(default) Select this option to run a Storage Resource agent on the target computer as a daemon service. A daemon service runs in the background on a monitored computer and listens

for requests from the Tivoli Storage Productivity Center server on the port number defined in the **Port** field. Connectivity between the server and agent is established using SSL. The server and agent have their respective certificates and no additional information is required and the security provided by the SSL protocol.

Do not select this option if you want to run a Storage Resource agent as a stand-alone program on the monitored computer. If run as a non-daemon service, the agent responds to requests from the Tivoli Storage Productivity Center server on demand. Connectivity between the server and agent uses the same protocol that was used during the deployment of the agent: SSH, SMB, REXEC, or RSH.

See "Deployment considerations for Storage Resource agents" on page 312 for more information about deploying an agent as a daemon or non-daemon service.

Windows Service Information
Click this button 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. 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.

Change Authentication
Click this button to change the user ID, password, and certificate location for deploying Storage Resource agents. For more information, go to the information center. Search for **Changing authentication for a Storage Resource agent**.

Edit Selected Entries
Click this button to edit the host list.

Remove
Click this button to remove the Storage Resource agent from Tivoli Storage Productivity Center.

5. Click the When to run tab to enter the following information:

How often to run
You can specify a time to run:
- **Run now**
- **Run once at** (specify a date and time to run)

How to handle time zones
You can specify a time zone to use:
- **Use the time zone that the server runs in**
- **Use this time zone** (select a time zone)

6. Click the Alert tab to specify the following information:

Triggering-Condition
The triggering conditions you can specify are:
- Storage Resource agent Deployment Failed

Triggered-Actions

You can choose from the following check boxes:

- **SNMP Trap**
- **TEC / OMNIbusEvent**
- **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.

7. Click **File > Save**.

Deployment considerations for Storage Resource agents

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

Use the following information when deploying Storage Resource agents and migrating existing Data agents:

- Multiple Storage Resource agents that are probing or scanning the same storage resources
- Platforms that support Storage Resource agents deployments
- 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 and Data agents on the same computer
- Time zones for computers monitored by Storage Resource agents
- UNIX 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 Storage Resource agents for the same

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 GUI** in the **Search** box. Click the link to the document for the appropriate release.

Product functions that are not available for storage devices monitored by Storage Resource agents

Before you migrate an existing Data agent to a Storage Resource agent or deploy a new 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 storage devices that are monitored by Storage Resource agents:

- NetWare monitoring. You must use Data agents to monitor NetWare Servers.
- Monitoring of certain relational databases. For list of relational databases that can be monitored by 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 GUI** in the **Search** box. Click the link to the document for the appropriate release.
- Changing zone configuration and reporting of HBA, fabric topology, or zoning information for fabrics connected to hosts running Linux on IBM Power Systems or zSeries hardware. These limitations also apply to Storage Resource agents on all guest operating systems for VMware configurations.

If you are already using these functions in Tivoli Storage Productivity Center, it is recommended that you do not migrate the corresponding Data agents to Storage Resource agent. If you plan to use these functions in the future, you must install Data agents on the computers that you want to monitor.

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 **Productivity Center Administrator** role. See the Configuration - Role-to-Group Mappings topic for more information about user roles.
- 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 creating a deployment schedule. Tivoli Storage Productivity Center uses this ID to log in 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 Linux or UNIX, the user must have root authority. By default, an agent runs under the user 'root'.
- On Windows, 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. On the desktop of the Windows 2008 computer, click **Start** > **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 you include in a deployment schedule must share the same administrative user ID and password. IBM Tivoli Storage Productivity Center uses these user credentials to log in to the computers when installing Storage Resource agents. If the computers in a deployment schedule do not share the same administrative user credentials, you must create separate deployment jobs for them.

Note: 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 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 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 ID, 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 using the **Change Authentication** or **Update Storage Resource agent List** functions on the Data/Storage Resource agents panel located under the **Administrative Services** > **Data Sources** node in the navigation tree.

Storage Resource agents and Data agents on the same computer

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

Time zones for computers 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. The time zones of computers that are monitored by Data agents are shown as the local time zone in Tivoli Storage Productivity Center reports. For example, a computer located in Los Angeles shows the following time zones in the By Computer report in Asset reporting.

- If monitored by a Storage Resource agent: (GMT-8:00) GMT-8:00
- If monitored by a Data agent: (GMT-8:00) America/Los_Angeles Time zone

UNIX systems 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 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 - NetBIOS setting

If you want to install a Storage Resource agent or Data 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, click **Start > Settings > Control Panel > 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, click **Start > Settings > Control Panel > Administrative Tools**.

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

Locally stored policies
For locally stored policies, click **Start > Settings > Control Panel > Administrative Tools > 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 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**. 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 system, you must disable the User Account Control (UAC) remote restrictions on the Windows system. User Account Control is a security component on Windows.

Note: 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. Click **Start** > **Run**. Enter **regedit** and click **OK**.
2. Locate and click the following registry subkey:

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

3. 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. Click **Enter**.
   c. Right-click **LocalAccountTokenFilterPolicy**, then click **Modify**.
   d. 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.
   e. Exit the registry editor.

Additional information about Storage Resource agents

See "Planning for Storage Resource agents" on page 49 for more information about Storage Resource agents.

Creating a certificate for SSH protocol

Before installing the Storage Resource agents using the SSH protocol, you can optionally create a certificate.

Note: The Storage Resource agent only supports either DES-EDE3-CBC encryption or no encryption for the private key used in SSH protocol communication between the server and agent. The default encryption used in the **ssh-keygen** command on UNIX is always DES-EDE3-CBC. However, with Windows Cygwin, **ssh-keygen** generates a key with AES-128-CBC encryption if a passphrase is specified. If there is no passphrase, the private key is generated without encryption. For more information about encryption, see http://www.openssl.org/docs/apps/enc.html.

Creating a certificate for SSH protocol (non-Windows)

Note: The Storage Resource agent only supports either DES-EDE3-CBC encryption or no encryption for the private key used in SSH protocol communication between the server and agent. The default encryption used in ssh-keygen command on UNIX is always DES-EDE3-CBC but with Windows Cygwin, it is using AES-128-CBC encryption if a passphrase is specified. If there is no passphrase, the private key is generated without encryption.

To create a certificate for SSH protocol, complete the following steps:

To create a certificate for SSH protocol, complete the following steps:

1. Telnet to the remote machine using the root user ID.
2. To create an SSH certificate on AIX, you must first install the following packages (if not already installed):

   ```
   openssl.base.openssh.base.client
   openssh.base.server
   ```

3. Go to the directory where you want to create the certificate:

cd to **~/.ssh**

4. Enter **ssh-keygen**. Accept the default names (for example, **id_rsa**).
5. Enter the passphrase.
6. Two files are created:

   id_rsa The private key.

   id_rsa.pub
   The public key.
7. Create an `authorized_key` file in the same location as `id_rsa.pub` by entering the following command:

   ```
   cat id_rsa.pub >> authorized_keys
   ```
8. Copy the `id_rsa` (private key) to your server machine. For example, to copy the `id_rsa` file to `:\keys\id_rsa` on the IBM Tivoli Storage Productivity Center server (user responses are in boldface type):

   ```
   # ssh-keygen
   Generating public/private rsa key pair.
   Enter file in which to save the key (//.ssh/id_rsa):
   Enter passphrase (empty for no passphrase):
   Enter same passphrase again:
   You identification has been save in //.ssh/id_rsa.
   Your public key has been save in //.ssh/id_rsa.pub.
   The key fingerprint is:
   xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx root@server
   # cat id_rsa >> authorized_keys
   # ls -l
   total 24
   -rw-r--r--  1  root  system  1743 Oct 15 09:40 authorized_keys
   -rw---      1  root  system  1743 Oct 15 09:39 id_rsa
   -rw-r--r--  1  root  system   399 Oct 15 09:39 id_rsa.pub
   #
   ```

   Note: You must copy the file in binary mode.
9. To connect to the remote system using the private key, the following information needs to be entered on the Remote Agent Machines panel when you install the Storage Resource agent:
   - User
   - Certificate Location (`c:\keys\id_rsa`)
   - Passphrase

Setting up an SSH daemon on Windows

On Windows Server 2003, Windows XP, Windows Server 2008, or Windows Vista, you must run the **ssh-host-config** command.

Note: Cygwin is not a prerequisite for the Storage Resource agent on Windows. To use the SSH protocol on Windows, an SSH software program must be used because Windows does not come with an SSH service. Cygwin is a free software program providing SSH access to a Windows server. Cygwin can be used if you want to run the Storage Resource agent using the SSH protocol.

You must be in a Cygwin window or be an X term user to create the **sshd** service. In most cases, you click the **cygwin.bat** to start the Bash shell.

Run the following steps:

1. Install Cygwin.

2. Set up your **sshd** service in Cygwin.
3. Create the certificate (this step is optional).

Installing Cygwin

To install Cygwin, go to http://cygwin.com. This page contains a link that displays help for the setup program and a link to download the setup program. Read the help before running the setup program. Then download the Cygwin program by clicking the **Install Cygwin now** link. Start the setup program on your computer by running the **setup.exe** program. Select the appropriate download option (**Install from Internet**, **Download from Internet**, or **Install from Local Directory**) as described in the help files.

If you are upgrading from an older version of Cygwin to a newer version, you need to remove the **sshd** service before installing the new version of Cygwin.

Accept the default installation options as they are presented to you (Root Directory, Install For, Default Text File Type, and so on). Select a download mirror that is geographically close to your location. Some sites require an FTP account before you can install Cygwin. You can either request an account or simply select another mirror.

During the installation process, a Select Packages list is displayed. Expand the plus sign (+) next to the Admin category and select **cygrunsrv** and the **Bin** check box. Expand the plus sign (+) next to the Net category and select **openssh**. Expand the plus sign (+) next to the Util category and select **diffutils**. Click **Next** to resume the setup program. The time required to download the packages depends on how busy the mirror is, and on the speed of your internet connection. With **openssh** and **cygrunsrv**, the downloaded files require approximately 70 MB of disk space. Allow 20 minutes to 30 minutes for the download and installation to complete.

Setting up your sshd service in Cygwin

Here is an example of the sequence of steps and responses. The responses to the prompts are in boldfaced type.

1. Run the **ssh-host-config** command.

   Note: With Cygwin, you might experience permission problems when running the **ssh-host-config** command. If you have permission problems, run these commands:

   ```
   chmod +r /etc/passwd
   chmod +r /etc/group
   chmod 777 /var
   ```

```
$ ssh-host-config
Generating /etc/ssh_host_key
Generating /etc/ssh_host_rsa_key
Generating /etc/ssh_host_dsa_key
Generating /etc/ssh_config file
Should privilege separation be used? (yes/no) no

Generating /etc/ssh_config file

Warning: The following functions require administrator privileges!

Do you want to install sshd as service? yes (If sshd is already
installed as a service, answer no.)

You appear to be running Windows 2003 Server or later. On 2003 and
later systems, it's not possible to use the LocalSystem account, if
sshd should allow passwordless logon (e. g. public key authentication).

If you want to enable that functionality, it's required to create  a
new account sshd_server' with special privileges, which is then used
to run the sshd service under.

Should this script create a new local account 'sshd_server' which  has
the required privileges? (yes/no) yes

Please enter a password for new user 'sshd_server'. Please be sure that
this password matches the password rules given on your system.

Entering no password will exit the configuration.
PASSWORD=password (Specify a password for the sshd_server
account.)

User 'sshd_server' has been created with password 'password'.

If you change the password, please keep in mind to change the password
for the sshd service, too.

Also keep in mind that the user sshd_server needs read permissions on
all users' .ssh/authorized_keys file to allow public key authentication
for these users. (Re-)running ssh-user-config for each user will set
the required permissions correctly.

Which value should the environment variable CYGWIN have when sshd starts?
It's recommended to set at least "ntsec" to be able to change user
context without password.

Default is "ntsec".

CYGWIN=ntsec

The service has been installed under sshd_server account.

To start the service, call 'net start sshd' or 'cygrunsrv -S sshd'.

Note! If the service doesn't start because of a login failure

Host configuration finished. Have fun!
```

2. Start the **sshd** service:
   a. Open a command prompt window.
   b. Enter **net start sshd** or in a Bash prompt, enter **cygrunsrv -start sshd**.
   c. Verify that the daemon is running.
   d. Enter **ps -a**. Examine the output to see if /usr/sbin/sshd is contained in the list of running processes.

   To stop the service from a Windows command prompt, enter **net stop sshd**. Alternatively, you can change to the `C:\cygwin\bin` directory (or open a Bash shell) and enter **cygrunsrv -stop sshd**.

3. When you have started the **sshd** service, test it by entering the following command from a Bash shell prompt:

   ```
   ssh localhost -l user_ID
      or
   ssh host_name -l user_ID
   ```

   If **localhost** does not work, use the short host name. If you receive a message indicating that the authenticity of localhost cannot be established, answer **Yes** to the question "Are you sure you want to continue connecting?" When prompted for your account password on **localhost**, enter the password you use when logging in to the computer.

4. Create the accounts that can log in to the computer:
   a. Create the Windows accounts. Click **Start > Settings > Control Panel > User Accounts**. Make each user a member of the Administrators group. Perform this operation for each user you want to add before you create the corresponding Cygwin accounts.
   b. Make a backup copy of the /etc/passwd file.
   c. Create the Cygwin user accounts. Run the following command:

      ```
      mkpasswd -l>/etc/passwd
      ```

   d. Verify that a home directory has been created for each account that you have added. Change the ownership of the home directory to its owner (run the **chown** command). If a home directory for the user does not exist, create one. For example, enter the following command:

      ```
      mkdir home/account_name;chown account_name/home/account_name
      ```

   e. When you add users, you need to stop and start **sshd** before that account is recognized because **sshd** only reads the file /etc/passwd when the service starts.
   f. If you need to create groups of accounts, create the Windows groups first, then create the Cygwin groups. After creating the Windows groups, run the following command:

      ```
      mkgroup -l>/etc/passwd
      ```

      When you add groups, you also need to stop and start **sshd** before the new group is recognized.

5. Set the TEMP environment variable. For information about setting the environment variable, see http://www.cygwin.com/cygwin-ug-net/setup-env.html.

   Here is an example of setting the environment variable:
   a. Right-click **My Computer** > **Properties** > **Advanced** > **Environment Variables**.
   b. Under **System variables**, find out the value of TEMP. For example, "C:\WINNT\TEMP"
   c. Set the TEMP environment variable to point to the Cygwin format of TEMP in the ~/.bashrc file. For example run the following command:

      ```
      export TEMP=/cygdrive/c/WINNT/temp
      ```

      Uncomment and modify this line in the ~/.bashrc file from the default:

```
# export TEMP=/tmp
to
export TEMP=/cygdrive/c/WINNT/temp
```

The Cygwin **sshd** service must be added as a service that starts automatically. To verify this step, click **Start > Settings > Control Panel > Administrative Tools > Services**. Look for **CYGWIN sshd** in the name list. Verify that it is started and configured to start automatically.

Creating the certificate

To create a certificate for SSH protocol, complete the following steps:

1. Run this command:

   ```
   cd to ~/.ssh
   ```

2. Generate the public and private keys with a passphrase.From the Bash shell prompt, here is an example of the input and output (user responses are in boldface type):

```
Administrator:

~/.ssh
$ openssl genrsa -des3 -out key 1024

Response:

Generating RSA private key, 1024 bit long modulus
...
...
e is 65537 (0x10001)
Enter pass phrase for key: (enter pass phrase for key)
Verifying - Enter pass phrase for key: (enter pass phrase for key again)

Administrator:

~/.ssh
$ chmod 600 ~/.ssh/key

~/.ssh
$ ssh-keygen -y -f ~/.ssh/key > key.pub

Response:

Enter passphrase: (enter passphrase)

Administrator:

~/.ssh
$ cat key.pub >> authorized_keys
```

3. Once the `key.pub` file has been added to the `authorized_keys` directory, copy the private key **key** to the IBM Tivoli Storage Productivity Center server.

Replacing a custom certificate for SSL protocol

IBM Tivoli Storage Productivity Center provides default SSL certificates for communication between the Data server and Storage Resource agent. You can replace your custom certificates with the default certificates.

Overview of replacing a custom certificate for SSL protocol

IBM Tivoli Storage Productivity Center uses SSL certificates for communication between the Data server and Storage Resource agent (daemon service). IBM Tivoli Storage Productivity Center provides default SSL certificates for this communication. If you want to use your own certificates, you can replace the default certificates with your custom certificates.

Server certificate

The IBM Tivoli Storage Productivity Center Data server uses the `TPCDataServer.jks` and `server.pwd` files for communication with the Storage Resource agents. If you are using custom certificates, you need to replace these files.

There are two ways you can replace these certificates:

- Before installation of the Data server
- After installation of the Data server

Storage Resource agent certificate

The Storage Resource agent uses the certificate files `sra.pem` and `sra.pwd` for communication with the Data server. These two files are compressed into the `certs.zip` file on the server system for deployment purposes. If you are using custom certificates, you need to replace these files.

There are two ways you can replace the certificates:

- Before deployment of the agent
- After deployment of the agent

The general steps for replacing custom certificates are:

1. Generate the custom certificates.
2. Stop the Data server (and the Storage Resource agent, if the agent is already deployed) if applicable.
3. Replace the custom certificates on the Data server and Storage Resource agent or on the disk 1 or Storage Resource Agent installation image.
4. Start the Data server (and the Storage Resource agent, if the agent is already deployed) if applicable.

Note: When generating custom SSL certificates, the certificates have a start date, end date, and time when they are valid. These dates and times are related to the system where these custom certificates were generated (which is usually the server system). When installing a Storage Resource agent on a remote system, you must check the date and time on the Storage Resource agent system. If the server and agent systems are in the same time zone, they must have the same date and time. Otherwise, the time zone difference should be set.

For example, if the server system is 8:00 PM, the agent system should also be 8:00 PM. If the agent system is set at a different time (for example, 6:00 PM) at the time the SSL custom certificates are generated on the server system with a time of 8:00 PM, the deployment of the Storage Resource agent will fail.

How to generate custom certificates

The script file `createSRACerts.sh` (for Linux or UNIX) or `createSRACerts.bat` (for Windows) is located in the following directory:

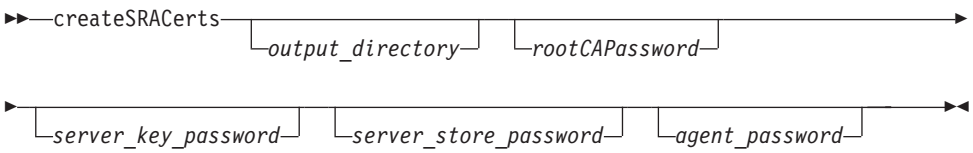
TPC_install_directory`/data/sra/tools/certs`

TPC_install_directory is where the IBM Tivoli Storage Productivity Center servers are installed. The default directory is `/opt/IBM/TPC` for Linux or UNIX or `C:\Program Files\IBM\TPC` for Windows.

To generate custom certificates, follow these steps:

1. Create the custom certificates.

   The `createSRACerts` script creates the custom certificates. The syntax is:



output_directory
Directory where the certificates are created. The default is `sra_certs_out` directory in the current working directory.

rootCAPassword
Root CA password (root common authority password). The default non-encrypted password is: s5umEvApR6cafruhustu.

server_key_password
Server key password. The default non-encrypted password is: drUtaxahaswefraf9uth.

server_store_password
Server store password. The default non-encrypted password is: wr4d5Xekaqafehet5u2a.

agent_password
Agent password. The default non-encrypted password is: jawUchezuthew6azEjef.

Note: You are prompted for the password when the script is run. Provide the value of the rootCAPassword on the command line (if specified on the command line). Otherwise, use the default values in the `createSRACerts.sh` or `createSRACerts.bat` script when prompted for the password (twice).
The following example creates the SSL certificate using the default output directory and default passwords.

```
createSRACerts
```

The following examples create the SSL certificates using the directory `/tmp` on UNIX and `C:\temp` on Windows. These examples use the default passwords. The passwords are stored in the following files:

Windows
`C:\temp\sra_certs_out`

Here is an example:

```
createSRACerts C:\temp
```

UNIX or Linux
`/tmp/sra_certs_out`

Here is an example:

```
createSRACerts /tmp
```

The following examples create SSL certificates in the directory `/tmp` on UNIX and `C:\temp` on Windows. These examples use non-default passwords for root CA and server key. The passwords are stored in the following files:

Windows
`C:\temp\sra_certs_out`

Here is an example of creating a non-default password for the root CA password and server key password:

```
createSRACerts C:\temp newpasswordforrootCA newpasswordforserver
```

UNIX or Linux

```
/tmp/sra_certs_out
```

Here is an example of creating a non-default password for the root CA password and server key password:

```
createSRACerts /tmp newpasswordforrootCA newpasswordforserver
```

2. Regenerate the certificates again if you have a failure. Delete the files in the out directory before rerunning the `createSRACerts` script.
3. Stop the Data server and Storage Resource agent. If the Data server is installed and running, stop the Data server. Make sure that the JVM for the Data server is not running.

   Stop the Storage Resource agent if the certificates for the agent are installed and running.
4. Replace the certificates. There are three scenarios:
   - Replacing the certificates after the server is installed.
   - Replacing the certificates before the server is installed.
   - Replacing the certificates after the agent is installed.

   Replacing the certificates after the server is installed

   The server certificates are in the following directory:

   TPC_install_directory/data/sra/tools/certs/out

   These are the server certificates:

   ```
   TPCDataServer.jks
   server.pwd
   ```

   For the server certificates, copy the server certificates to the following directory:

   TPC_install_directory/data/sra/certs

   The agent certificates are created in the `certs.zip` file in the following directory:

   TPC_install_directory/data/sra/tools/certs/out/agent

   Copy the `certs.zip` file into the following directory:

   TPC_install_directory/data/sra/*agent_operating_system*

   Extract the `certs.zip` file into the following directory:

   TPC_install_directory/data/sra/*server_operating_system*

   agent_operating_system is the operating system that the Storage Resource agent is running under. *server_operating_system* is the operating system on which the Data server is installed on.

   Replacing the certificates before the server is installed

   The server certificates are in the following directory:

   disk1_or_SRA_image_install_directory/data/sra/tools/certs/out

   These are the server certificates:

   ```
   TPCDataServer.jks
   server.pwd
   ```

   disk1_or_SRA_image_install_directory is the location of the disk1 or Storage Resource agent image.

For the server certificates, copy the server certificates into the following directory:

TPC_install_directory/data/sra/certs

The agent certificates are created in the certs.zip file in the following directory:

TPC_install_directory/data/sra/tools/certs/out/agent

Copy the certs.zip file into the disk1 or Storage Resource agent image directory:

disk1_or_SRA_image_install_directory/data/sra/*agent_operating_system*

Disk1 or SRA is the location of the disk1 or Storage Resource agent image that was extracted for installation of the server and agents.

Extract the certs.zip file into the following directory:

disk1_or SRA_image_install_directory/data/sra/*server_operating_system*

agent_operating_system is the operating system that the Storage Resource agent is running under. *server_operating_system* is the operating system on which the Data server is installed on.

Replacing the certificates after the agent is installed

The agent certificates are in the certs.zip file that gets created in the following directory:

TPC_install_directory/data/sra/tools/certs/out/agent

Extract the certs.zip file into the *storage_resource_agent_install_directory*/certs directory. *storage_resource_agent_install_directory* is the location where the Storage Resource agent has been installed.

5. Start the Data server and the Storage Resource agent. If the Data server has stopped for replacement of the certificates, then start the Data server after the replacement of the certificates.

   If the Storage Resource agent has stopped for replacement of certificates, then start the Storage Resource agent after the replacement of the certificates.

Storage Resource Agent Upgrades

You can configure IBM Tivoli Storage Productivity Center for a Storage Resource agent upgrade.

When you apply maintenance to Tivoli Storage Productivity Center, you can upgrade the Storage Resource agent now or upgrade the agent at a later time. To ensure that all your agents are at the current release level and to manage your network load, schedule upgrades on a regular basis.

To configure Tivoli Storage Productivity Center for a Storage Resource agent upgrade, complete the following steps:

1. In the navigation tree, expand **Administrative Services** > **Configuration** > **Storage Resource Agent Upgrades**.
2. Right-click **Storage Resource Agent Upgrades**. Click **Create Storage Resource Agent Upgrade**.
3. The Create Storage Resource Agent Upgrade panel is displayed in the content pane. You have the following tabs:

Computers
You can select Computer Groups or Computers.

When to Run
You can specify how often to run the job or how to handle time zones.

Options
You can specify to force an upgrade for an agent, maintenance upgrade for an agent, or language pack to apply to an agent.

Alerts You can specify the triggering conditions and triggering actions.

Enter your information in the appropriate fields.

4. Save and run the job.

Configuration - Migrate Data/Fabric Agent

Define schedules to migrate existing Data agents or Inband Fabric agents to Storage Resource agents. The migration process installs a Storage Resource agent on a target host and then uninstalls the existing agent.

IBM Tivoli Storage Productivity Center provides the following benefits when you migrate an existing Data agent or Inband Fabric agent to a Storage Resource agent:

- Storage Resource agents require less resources on a host computer than a Data agent or Inband Fabric agent that is based on Common Agent Services. Additionally, you can deploy Storage Resource agent as a daemon or non-daemon service.
- Reduced complexity when deploying agents. You can deploy Storage Resource agents directly from the Tivoli Storage Productivity Center user interface and they do not require you to install Agent Manager. For Data agents and Inband Fabric agents, you must use the Tivoli Storage Productivity Center installation program and ensure that Agent Manager is registered with the Data server and Device server.
- Improved interface when deploying, upgrading, and administering agents. You can manage Storage Resource agents using the nodes in the **Administrative Services** section of the navigation tree. You can deploy and upgrade Storage Resource agents on schedules that you define.

See "Deployment considerations for Storage Resource agents" on page 312 to view additional information about deploying and using Storage Resource agents.

To migrate a Data agent or Inband Fabric agent to a Storage Resource agent on a host, complete the following steps:

1. Expand **Administrative Services > Configuration** in the navigation tree.
2. Right-click **Migrate Data/Fabric Agent** and select **Create Data/Fabric Agent Migration** from the menu. The Create Data/Fabric Agent Migration panel is displayed.
3. Enter information about the schedule on the following pages:
   - **Computers:** Use this page to select the computers and groups of computers that contain the Data or Fabric agents that you want to migrate to Storage Resource agents.
   - **When to run:** Use this page to determine when the Data and Fabric agents in a schedule are migrated.

- **Options:** Use this page to specify the runtime settings of a Storage Resource agent, the method for migrating an agent, and the login credentials for an agent.
- **Alert:** Use this page to define the actions that occur if a Data/Fabric Agent Migration schedule fails.

4. Click **File** > **Save** from the menu bar to save the migration schedule. The migration runs according to the date and time that you defined in the schedule definition. Use the **Administrative Services > Data Sources > Data/Storage Resource Agents** panel to view a list of the Storage Resource agents that are deployed in your environment.

Update Storage Subsystem Credentials

If you have a System Storage DS8000, XIV system, or SAN Volume Controller storage system, you must update your storage system credentials to use the Tivoli Storage Productivity Center native interfaces.

To enter the storage subsystem credentials, follow these steps:

1. Start the Tivoli Storage Productivity Center migration tool to enter storage subsystem credentials. You can start the migration tool in one of the following ways:
   - Use the migration tool from the Tivoli Storage Productivity Center DVD. Use the DVD or download the migration tool: `MigrateUserInfo.bat` (for Windows) or `MigrateUserInfo.sh` (for UNIX or Linux). Run the migration tool from the directory `UserMigrationTool`.
   - Run the Tivoli Storage Productivity Center upgrade installation program. When Tivoli Storage Productivity Center detects the storage system, use the panel that is displayed to run the migration tool. Run the migration tool.
   - After upgrading Tivoli Storage Productivity Center, open the GUI. In the navigation tree, expand **Administrative Services** > **Configuration** > **Update Storage Subsystem Credentials**. This node is visible only when Tivoli Storage Productivity Center detects that you have a System Storage DS8000, XIV system, or SAN Volume Controller storage system.
2. The Update Storage Subsystem Credentials panel is displayed.

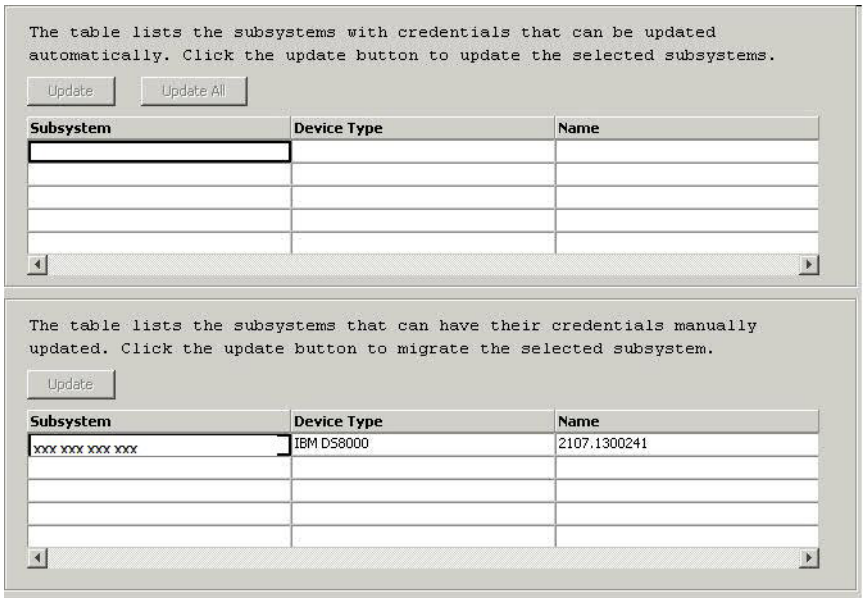


Figure 42. Update Storage Subsystem Credentials panel

Top table
All System Storage DS8000 storage systems are displayed if Tivoli Storage Productivity Center can use the existing CIM agent user ID and password to connect to the HMC. If you want to use the existing credentials, click **Update** or **Update All** (if you have more than one System Storage DS8000). Tivoli Storage Productivity Center updates the credentials and removes the System Storage DS8000 from the table.

Bottom table
All XIV system and SAN Volume Controller storage systems are displayed in this table. All other System Storage DS8000 storage systems are displayed in this table if they are not listed in the top table. All XIV system and SAN Volume Controller storage systems must have the storage system credentials migrated manually.

If you want to manually update the credentials for the System Storage DS8000, select a storage system and click **Update**. The Update Subsystem Credentials window is displayed. Enter the following information:

Username
Enter the user ID for the System Storage DS8000 element manager.

Password
Enter the password for the user name.

HMC2 Address (Optional)
You can enter a second HMC address. This parameter is optional.

Click **Save**. The System Storage DS8000 is removed from the table after you have updated the credentials.

If you want to update the credentials for the XIV system, select a storage system and click **Update**. The Update Subsystems Credentials window is displayed. Enter the following information:

Username
Enter the user ID for the XIV system.

Password
Enter the password for the user name.

Click **Save**. The XIV system is removed from the table after you have updated the credentials.

If you want to update the credentials for IBM SAN Volume Controller, select a storage system and click **Update**. The Update Subsystem Credentials window is displayed. Enter the following information:

Software Version
Enter the version of the SAN Volume Controller software that you want to configure.

IP address
Enter the IP address or host name of the SAN Volume Controller. You can enter an IPv4 or IPv6 address, for example:

- 9.47.97.128
- [2001:DB8::1234:0000:0000:5678:ABCD]

Select Key
Select one of the following actions for the Secure Shell (SSH) key:

- **Use existing uploaded key:** Select this option if you want to use a public key that has already been uploaded using a method other than through Tivoli Storage Productivity Center, such as the SAN Volume Controller Web console. Enter the location of the private key file that corresponds to the public key in the **Private SSH Key** field. The **Admin Username** and **Admin Password** fields are not required if you select this option.
- **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 SAN Volume Controller cluster. If you select this action, the **Admin Username** and **Admin Password** fields are required.

Admin Username
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 SAN Volume Controller cluster that contains the device. Tivoli Storage Productivity Center uses this value during the setup process to contact the cluster 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.

Note: The **Admin Username** only applies to **SVC5.x+**; it does not apply to **SVC 4.x**.

Admin Password
Enter the password for the user that you entered in the **Admin Username** field.

Username
Enter an authorized user for the device. Tivoli Storage Productivity Center uses this value when monitoring and configuring 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 have selected **Upload new key** in the **Select Key** list and have entered a user name and password in the **Admin Username** and **Admin Password** fields.
Important: If you upload a new key for the existing SAN Volume Controller user using Tivoli Storage Productivity Center, that new key overrides the existing key for the user. If you have another application using that SAN Volume Controller user (for example Tivoli Storage Productivity Center for Replication), the application cannot log in to SAN Volume Controller because the key associated with the user is overwritten.
If you have this situation, create a user ID specifically for Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication on the SAN Volume Controller cluster. Add the SAN Volume Controller to Tivoli Storage Productivity Center for Replication using the new user ID and password (or change the user ID and password while in the disconnected state).

Select User
Click this button to select from a list of existing users that are defined on the SAN Volume Controller.

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:

a. Start the puttygen utility.
b. Load the `.ppk` file. Click **File > Load private key**.
c. 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 SSH key pair. The passphrase protects your private SSH key file. If you do not have a passphrase, leave this field blank.

Click **Save**. The SAN Volume Controller is removed from the table after you have updated the credentials.

After a successful migration of System Storage DS8000, XIV system, or SAN Volume Controller credentials, the CIMOMs that manage these devices previously should be removed from the Tivoli Storage Productivity Center configuration. The migration process tries to determine and remove these CIMOMs during migration automatically, but if there are CIMOMs still remaining, they should be deleted

manually. For DS CIMOMs managing DS6000 or ESS devices and additional System Storage DS8000 devices, the System Storage DS8000 devices should be removed from the CIMOM if possible.

NetWare Tree Logins

You must specify a fully qualified login ID and password for each of the Novell Directory Services trees (NDS trees) discovered by licensed agents. Data Manager uses this login ID to log in to your NDS trees and gather information about the NetWare servers and volumes in those trees.

When you log in to the NDS trees, you see the licensed NDS trees discovered by agents installed in your environment. The login ID you assign to each NDS tree must have permission to enumerate the volumes within the NetWare servers on that tree.

Note: Only the Data agents support NDS trees.

To log in to the NDS trees, complete the following steps:

Expand **Administrative Services** > **Configuration** > **NetWare Tree Logins**.

Resource History Retention

You can specify how long to keep a history of the statistical elements collected by the system.

By specifying a number for days, weeks, or months for each element, 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.

You can set retention periods for:

- Directories
- Pings
- Disks
- Performance Monitors
- Filesystems
- Computer Uptime

Note:

- If you do not select a check box, the data related to that check box is retained permanently. This action might cause a large amount of data to accumulate in the database repository over time. If you select a check box and enter a value of 0, the data related to that check box is removed immediately from the database repository.
- When you set the retention period for Performance Monitors, remember the following restrictions:
  - The hourly and daily values for Performance Monitors can be set to 0 only if the per performance monitoring task is also set to 0.
  - The hourly and daily values for Performance Monitors cannot be set lower than the per performance monitoring task value.

To set the retention periods, complete the following steps:

1. Expand **Administrative Services** > **Configuration** > **Resource History Retention**.
2. The Retain History panel is displayed in the content pane. Enter the information for the retention periods.

   Note: Some default profiles have `Accumulate history` settings. If set, these settings dictate the number of days, weeks, or months to be retained. If a scan includes any of these profiles, then the history data is persisted and appears in the corresponding reports. The following default profiles have `Accumulate history` settings:
   - TPCUser.By Access
   - TPCUser.By Creation
   - TPCUser.By Mod Not Backed Up
   - TPCUser.Summary By File Type
   - TPCUser.Summary By Filesystem/Directory
   - TPCUser.Summary By Group
   - TPCUser.Summary By Owner
   - TPCUser.Temporary Files
   - TPCUser.Wasted Space

Removed Resource Retention

You can specify the number of days to keep information about directories, file systems, disks, and so on, that have been removed from the system and can no longer be found.

Information is kept in the database repository. To clear the history record for a resource and activate a new period for resource retention, perform a discovery or probe job on the resource.

You can set the retention period for the following:
- Directories
- File systems
- Disks
- Filers
- Subsystems
- Fabric
- Tape
- Hypervisors
- Virtual Machines
- IBM Tivoli Storage Productivity Center Servers: When you remove a data source (agent) from a subordinate server or a storage entity that it monitors, the master server continues displaying information about that agent or storage entity in roll up reports for a specified amount of time. You can specify how quickly that information is removed by changing this setting.

If the check box for a storage resource is not selected, the data related to the resource is retained permanently. This action might cause a large amount of data to accumulate in the database repository over time. If you select a check box and enter a value of 0, the data related to that check box is removed immediately from the database repository.

The retention time for cluster resource groups that are monitored by Tivoli Storage Productivity Center is determined by the values that you enter in the **Filers** and **Filesystems** fields. If the values in the **Filesystems** and **Filers** fields are different, the field with the higher value (longer retention period) is applied to cluster resource groups.

The following list shows the fields that determine the retention period for Storwize V7000 File Module storage resources:

- **Filers**: Storage systems, storage pools, and nodes.
- **Disks**: Network Shared Disks (NSDs).
- **Filesystems**: File systems, exports, file sets, and quotas.

To set the retention period, complete the following steps:

1. Expand **Administrative Services** > **Configuration** > **Removed Resource Retention**.
2. The Retain Removed panel is displayed in the content pane. Select the resource and enter the number of days to keep.

Resource History Retention for Databases

You can specify how long to keep a history of the database-related statistical elements collected by the system.

By specifying a number of days, weeks, or months for each element, you can control the amount of data that is retained and is available for historical analysis and charting. The longer you keep the data, the more informative your analysis.

You can retain the histories for:

- Databases-Tablespaces
- Tables

Note: If you do not select a check box, the data related to that check box is retained permanently. This action might cause a large amount of data to accumulate in the database repository over time. If you select a check box and enter a value of 0, the data related to that check box is removed immediately from the database repository.

To set the retention period, complete the following steps:

1. Expand **Administrative Services** > **Configuration** > **Resource History Retention for Databases**.
2. The Retain History panel is displayed in the content pane. Enter the retention period information.

Removed Resource Retention for Databases

You can specify how long to keep information in the enterprise repository that is related to a database entity that has been removed from the system and can no longer be found.

By specifying a number for days, you can indicate how long to keep information for table spaces and tables that have been removed from the system.

Note: If you do not select a check box, the data related to that check box is retained permanently. This action might cause a large amount of data to

accumulate in the database repository over time. If you select a check box and enter a value of 0, the data related to that check box is removed immediately from the database repository.

To specify a retention period, complete the following steps:

1. Expand **Administrative Services** > **Configuration** > **Removed Resource Retention for Databases**.
2. The Retain Removed panel is displayed in the content pane. Enter the information for the retention period.

Configuration History Settings

This page allows you to specify how often the system captures snapshots of your configuration and when to delete them. It also gives the number of snapshots in the database, indicates when the last one was taken (and allows you to refresh this setting), and enables you to create and enter a title for a snapshot on demand.

You must configure and save the settings on this page before you can use the Configuration History function.

To specify the configuration history settings, complete the following steps:

1. In the Create snapshot every field, type how often (in hours) you want the system to take snapshot views of the configuration.
2. To automatically delete snapshots, click the check box to place a check mark before Delete snapshots older than. In the box that follows that field, type how long you want to keep the snapshots (in days) before they are automatically deleted.
3. To optionally refresh the date and time of when the latest snapshot was created, click **Update**.
4. To optionally create a snapshot on demand, click **Create Snapshot now**. If you choose, you might want to type a name for the snapshot in the box **Title this snapshot** (optional).
5. To change your settings to the default, click **Reset** to defaults. The defaults are create snapshots every 12 hours and delete snapshots older than 14 days.
6. Click **File > Save** to save the configuration history settings.

Configuration: Device server specific

This section provides configuration information that is specific to the Device server.

Configuring switches

This section provides configuration information for switches.

IBM Tivoli Storage Productivity Center is a storage area network (SAN) management application that discovers devices in the SAN and displays a topology of the SAN environment. Tivoli Storage Productivity Center is designed to operate using industry-based standards for communicating with fiber channel switches and other SAN devices. This communication can be done using the simple network management protocol (SNMP) interface for out-of-band agents, the FC-GS-3 interface for in-band Fabric and Storage Resource agents, Storage Management Initiative (SMI) agents, or a combination of these agent types. FC-GS-3 refers to the Fibre Channel Generic Services 3 standard. To gather and display the information as expected, the switches must be configured correctly. The configuration varies by vendor and the type or types of agents that are used. One

of the common sources of customer problems is incorrect configuration of the switches that are being managed. This incorrect configuration leads to missing information and misconceptions that the Tivoli Storage Productivity Center product does not work with certain switches. The switch vendors supported are Brocade, Cisco, McData, and QLogic. Other vendors such as IBM, often sell these switches under their own labels.

Determining the agent type or types to used with a switch

For Brocade and McData fabrics, the preferred type of agent is the SMI agent. The SMI agent provides most fabric functions and the other agent types can be added for redundancy. However, Storage Resource agents or in-band Fabric agents are required to gather host bus adapter (HBA) information.

For QLogic and Cisco fabrics, a combination of different agent types is required to enable all functions.

For information about information that is gathered by each agent type, refer to the topic about discovering storage resources in the IBM Tivoli Storage Productivity Center User's Guide.

Using in-band Fabric or Storage Resource agents

With Tivoli Storage Productivity Center in-band Fabric and Storage Resource agent discovery, the Tivoli Storage Productivity Center for agent software is installed on SAN-attached hosts. The in-band Fabric and Storage Resource agents collect information about the fabric across the fiber channel network by querying the switch and the attached devices through the host bus adapter (HBA) in the system. For the switches to successfully receive and respond to the queries, there are some basic requirements:

- The switch must support the FC-GS-3 standard interface for discovery.
  - Name server
  - Configuration server
  - Unzoned name server
- For zone control functions, the fabric zone server must be supported, except in the case of Brocade switches.

For in-band Fabric agent discovery, Fabric events are automatically sent from the agent to Tivoli Storage Productivity Center There is no need for configuration.

For Storage Resource agent discovery, you must configure SNMP traps to be sent from the switches in your fabric to the Tivoli Storage Productivity Center Device server.

Using out-of-band SNMP agents

Out-of-band SNMP agent discovery collects some of the same information that can be obtained by in-band Fabric and Storage Resource agents, but out-of-band agent discovery is performed differently. In out-of-band discovery, Tivoli Storage Productivity Center queries the switch directly rather than going through an in-band Fabric or Storage Resource agent and the fiber channel network.Tivoli Storage Productivity Center uses the SNMP protocol to send queries across the IP network to management information bases (MIBs) supported on the switch. Tivoli Storage Productivity Center uses the FC Management MIB (sometimes referred to as the FA MIB) and the FE MIB. The queries are sent only to switches that have

been added to Tivoli Storage Productivity Center for use as SNMP agents. SNMP information is collected for a single switch. The out-of-band discovery registers each switch. The following list shows the basic requirements for the switch to successfully receive and respond to the query:

- The FC Management MIB and FE MIB must be enabled on the switch.
- The switch must be configured to receive SNMPv1 queries and respond in SNMPv1. Some switches are configured to use SNMPv2 or SNMPv3 by default.
- The community string configured in Tivoli Storage Productivity Center must match one of the community strings configured on the switch with read access. Cisco switches must additionally have a community string match for write access. The default community strings in Tivoli Storage Productivity Center are "public" for read access and "private" for write access. Additional community strings can be defined on the switches, but are not used.
- SNMP access control lists must include the Tivoli Storage Productivity Center system. Some lists automatically include all hosts while others exclude all by default.

Another aspect of the SNMP configuration includes trap notification. SNMP traps are generated by the switch and directed to Tivoli Storage Productivity Center as an indication that something in the fabric has changed and that a discovery must occur to identify changes. The default configuration for handling switch traps is to send them from the switch to port 162 on the Tivoli Storage Productivity Center system. To successfully generate and receive traps, there are some configuration requirements:

- The trap destination parameter on the switch must be set. This parameter is the host that receives the trap and sends it to Tivoli Storage Productivity Center. The parameter is set on the switch.
- The destination port parameter on the switch must be set. Tivoli Storage Productivity Center listens on port 162 by default. The parameter is set on the host.
- The traps must be sent as SNMPv1. This parameter is set on the switch.
- The trap severity level must be set to generate traps for change conditions. This level typically means to send error level traps and anything more severe. This parameter is set in Tivoli Storage Productivity Center.

Using SMI agents

You must install or enable an SMI agent for to perform the following tasks:

- Gather fabric performance data.
- Work with a set of McData inactive zone sets if you are using Brocade Enterprise Fabric Connectivity Management (EFCM) (The zone sets are synchronized with the EFCM inactive zone sets for the server if the SMI agent is configured to run with a proxy connection to the EFCM Server.)
- Collect and configuring Brocade zone aliases.

For information about installing or enabling a SMI agent for the switch, contact your switch vendor.

Configuring the inactive zone sets for a fabric data source

Some fabrics have multiple locations to retrieve a database for inactive zone sets. In some scenarios, and for some switch vendors, these definitions for inactive zone sets are not synchronized.

For heterogeneous fabrics composed of McDATA and Brocade switches, IBM Tivoli Storage Productivity Center allows you to choose which CIM agent to use to retrieve inactive zone sets and when configuring zoning. To set the inactive zone sets for **Data Source**, go to **Fabric Manager > Fabrics**, select the fabric, and click **Select Inactive Zone Sets Data Source**.

Excluding devices from Fabric agent scans

If you want to run the Fabric agent on a system that can access removable media devices such as tapes, you can remove those devices from Fabric agent attribute scans. Some removable media devices cannot handle command queuing, causing a long tape read or write command to fail. To avoid this problem, follow these steps:

1. Stop the Fabric agent.
2. Create a file named `ExcludeList` in the following directory:

   agent_install_directory`/conf`
3. In the `ExcludeList` file, enter the worldwide names of the devices that you want to be excluded from Fabric agent scans, each on a new line. Save the file. The file name needs to be the same as the one provided.

   Here is an example of an `ExcludeList` file to exclude the four WWNs 10000000c920d02a, 10000000c920ccf9, 10000000c93f51ca, 10000000c926342a from receiving **scsi** commands from the Attribute scanner:

   ```
   10000000c920d02a
   10000000c920ccf9
   10000000c93f51ca
   10000000c926342a
   ```
4. Start the Fabric agent.

Managing a SAN without agents

In the following situations, there might not be any agents on the SAN:

- The hosts do not currently have a Storage Resource agent or Fabric agent installed.
- The host operating system is not supported by the Storage Resource agent or Fabric agent.
- The customer requirements do not require the deployment of a Storage Resource agent or Fabric agent (topology map only).

In these cases, it is recommended that an agent is installed on the Device server itself. This action allows the Device server to use advanced features like Remote Node Identification (RNID) which requires an agent.

Normally the Device server does not have a Fibre Channel host bus adapter. In this configuration, the following steps are taken:

1. A Fibre Channel host bus adapter is added to the manager.
2. An agent is installed on the Device server(the Device server is installed first).
3. All storage devices are verified to ensure that they use LUN masking techniques. The LUN masking techniques prevent the Device server from accessing the disks used by the host systems.
4. The Fibre Channel host bus adapter is attached to the SAN to be managed. This host is added to each zone that is intended to be managed by the Device server.

Setting timeout values for the Device server

This topic describes how to set 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 the native interfaces (DS8000, the XIV system, SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000), you see this error message:

```
HWN020103E The external process exceeded the timeout limit and was cancelled.
```

To set timeout values for the Device server, follow these steps:

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 timeout_value
   ```

   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 **5990**.
   - *timeout_property* is one of the following strings:
     - http.timeout
     - CIMClientWrapper.Timeout
     - Probe.Timeout.Array
     - Probe.Timeout.InBand
     - Probe.Timeout.LMM
     - Discovery.Timeout
     - CIMOMManager.TestConnectionTimeout

     For the storage systems that use the native interface, the *timeout_property* strings are:
     - NAPI.Timeout.TestConnection
     - NAPI.Timeout.Probe
     - NAPI.Timeout.EventPoll
   - *Timeout_value* is the value for the timeout.

   .
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.

Service Location Protocol (SLP) configuration considerations

This topic provides some configuration considerations for enabling Tivoli Storage Productivity Center to discover a larger set of storage devices through SLP. These recommendations cover some of the more common SLP configuration problems. This topic discusses router configuration, SLP directory agent configuration, and environment configuration.

For additional information about SLP, see the **Service Location Protocol Request for Comments** website at http://www.ietf.org/rfc/rfc2165.txt.

Note: The storage systems that use the native interfaces (DS8000, the XIV system, SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000) do not use SLP discovery.

Router configuration

Configure the routers in the network to enable general multicasting or to allow multicasting for the SLP multicast address and port, 239.255.255.253, port 427. The routers of interest are the ones associated with subnets that contain one or more storage devices that are to be discovered and managed by Tivoli Storage Productivity Center.

To configure your router hardware and software, refer to your router and configuration documentation.

SLP directory agent configuration

This topic provides information about suggestions for configuring the SLP directory agent.

Configure the SLP directory agents (DAs) to circumvent the multicast limitations. With statically configured DAs, all service requests are unicast by the user agent. Therefore, it is possible to configure one DA for each subnet that contains storage devices that are to be discovered by Tivoli Storage Productivity Center. One DA is sufficient for each of the subnets. Each of these DAs can discover all services within its own subnet, but no other services outside its own subnet. To allow Tivoli Storage Productivity Center to discover all the devices, it needs to be statically configured with the addresses of each of these DAs. This operation can be accomplished by using the Tivoli Storage Productivity Center Discovery Preference panel.

You can use this panel to enter a list of DA addresses. Tivoli Storage Productivity Center sends unicast service requests to each of these statically configured DAs, and sends multicast service requests on the local subnet on which Tivoli Storage Productivity Center is installed. Configure an SLP DA by changing the configuration of the SLP service agent (SA) that is included as part of an existing CIM Agent installation. This action causes the program that normally runs as an SLP SA to run as an SLP DA.

Note: The change from SA to DA does not affect the CIMOM service of the subject CIM Agent, which continues to function normally, sending registration and deregistration commands to the DA directly.

Environment configuration

This section provides information about the configuration of your environment.

It might be advantageous to configure SLP DAs in the following environments:

- In environments where there are other non-Disk Manager SLP UAs that frequently perform discovery on the available services, an SLP DA must be configured. This action ensures that the existing SAs are not overwhelmed by too many service requests.
- In environments where there are many SLP SAs, a DA helps decrease network traffic that is generated by the multitude of service replies. It also ensures that all registered services can be discovered by a given UA. The configuration of an SLP DA is recommended when there are more than 60 SAs that need to respond to any given multicast service request.

SLP registration and slptool

IBM Tivoli Storage Productivity Center uses Service Location Protocol (SLP) discovery, which requires that all the CIMOMs that Disk Manager discovers are registered using the SLP.

In a non-multicast network environment, SLP can only discover CIMOMs that are registered in its IP subnet. For CIMOMs outside of the IP subnet, you need to use an SLP DA and register the CIMOM using **slptool**. Ensure that the **CIM_InteropSchemaNamespace** and **Namespace** attributes are specified.

For example, enter the following command:

```
slptool register service:wbem:https://myhost.com:port
```

Where *myhost.com* is the name of the server hosting the CIMOM, and *port* is the port number of the service, for example 5989.

Note: slptool is installed with a CIMOM. Run the command from the machine hosting the CIMOM.

SLP discovery

A common problem with SLP discovery is due to IP multicasting being disabled on the network router. Communication between the SLP SA and UA is done with IP multicasting. Follow these recovery procedures when there are SLP discovery problems and IP multicasting is disabled on the network router.

Note: The storage systems that use native interfaces (DS8000, XIV system, SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000), do not use SLP discovery.

There are two recovery procedures when there are SLP discovery problems and IP multicasting is disabled on the network router:

1. Configure one DA for each subnet within the environment.
2. Enable IP multicasting on the router which is disabled by default. Here is a list of common router configurations for multicasting:
   - Internet Group Management Protocol (IGMP) is used to register individual hosts in particular multicast groups and to query group membership on particular subnets.
   - Distance Vector Multicast Routing Protocol (DVMRP) is a set of routing algorithms that use a technique called reverse path forwarding. These algorithms provide the best solution for how multicast packets are to be routed in the network.
   - Protocol-Independent Multicast (PIM) comes in two varieties: dense mode (PIM-DM) and sparse mode (PIM-SM). The dense mode and sparse mode

routines are optimized for networks where either a large percentage of nodes requires multicast traffic (dense) or a small percentage of nodes requires the sparse traffic.

- Multicast Open Shortest Path First (MOSPF) is an extension of OSPF. It is a link-state unicast routing protocol that attempts to find the shortest path between any two networks or subnets to provide the most optimal packet routing.

To properly configure the routers for multicasting, see the reference and configuration documentation from the router manufacturer.

Configuring IBM Tivoli Storage Productivity Center for a Windows domain

Use these topics to configure IBM Tivoli Storage Productivity Center for a Windows domain.

Configuring the Data server for a Windows domain

By default, the IBM Tivoli Storage Productivity Center Data server is installed to operate under a local account created at installation time. However, if you are functioning in a Windows domain environment, and want to use domain accounts for authentication and authorization with Tivoli Storage Productivity Center, then you must reconfigure the service.

Note: These configuration steps are for a Windows system that is a member of a Windows domain and not for the Windows Domain Primary Domain Controller.

To configure the Data server, complete the following steps:

1. Open the Windows Services panel and verify that the Windows service **IBM Tivoli Storage Productivity Center - Data Server** is configured to log on as a domain user in order to perform Network Management Functions.
2. If the Data server is not configured as indicated in step 1, perform the following steps:
   a. In the Windows Services panel, highlight **IBM Tivoli Storage Productivity Center - Data Server**.
   b. Right-click the service menu.
   c. Select **Properties**.
   d. Select the **Log on** tab.
   e. Specify a domain account under **This account** using the *domain_name**domain_user* format.
   f. Enter and confirm the password for the domain account.
   g. Click **Apply** for the setting, and then click **OK**.
   h. Restart the Data server.
3. Log on to the Tivoli Storage Productivity Center GUI as a Tivoli Storage Productivity Center superuser.
4. Click **Administrative Services > Configuration > Role-to-Group mappings**.
5. Delete any instances of previous domain group mappings.
6. Save the changes.
7. Create new Role-to-Group mappings by specifying only the domain group name.
8. Save the changes.

9. You can now log on to the Tivoli Storage Productivity Center GUI with a domain account belonging to the domain group by using the *domain_name**domain_user* format.

Granting local administrative privileges to a domain account

For Windows users, the user account which the Storage Resource agent runs under requires local administrative rights. Because these rights are not necessarily guaranteed for domain users in a Windows domain environment, this topic provides information about how to grant local administrative rights for 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 the Organizational Unit (OU) upon which you want to apply the Group Policy. Click **Properties**.
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**. 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\tapeadmin", "TPC\tapegroup", 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 the user right for 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.

Adding Tivoli Storage Productivity Center to the Windows domain

Use this information to add Tivoli Storage Productivity Center to the Windows domain for OS authentication. This procedure requires two steps, one for configuring Windows and one for configuring Tivoli Storage Productivity Center.

Configuring Windows

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.

Add the DNS IP as an IP of the domain controller, follow these steps:

1. Click **Start** > **Settings** > **Control Panel** > **Network Connections**.
2. Right-click **Local Area Connection**. Click **Properties**.
3. The Local Area Connection Properties window is displayed. Select **Internet Protocol Version 4 (TCP/IPv4)**. Click **Properties**.

   Note: This procedure uses Windows 2008. If you have Windows 2003, select **Internet Protocol** .
4. The Internet Protocol Version 4 window is displayed. On the General tab, enter the IP address of the domain controller in the **Preferred DNS server** field. Click **Advanced**.
5. The Advanced TCP/IP Settings window is displayed. On the DNS tab, enter the domain controller information in the **DNS Suffix for this connection** field. Select the check box for **Register this connection's addresses in DNS**. Click **OK**.
6. Right-click **My Computer** > **Properties**. Click **Change settings**.
7. The System Properties window is displayed. On the Computer Name tab, click **Change**.
8. The Computer Name/Domain Changes window is displayed. Click **More**.
9. The DNS Suffix and NetBIOS Computer Name window is displayed. Enter the domain controller information in the **Primary DNS suffix of this computer** field. Click **OK**.
10. The Computer Name/Domain Changes window is displayed. Click **Domain**.
11. Provide the authentication credentials to the domain controller.
12. Restart the computer for the changes to take effect.

Configure Tivoli Storage Productivity Center with domain controller credentials

Follow these steps:

1. On the desktop, click **My Computer** > **Manage**.
2. Click **Configuration** > **Local Users and Groups** > **Groups**.
3. In the Groups pane on the right, double-click **Administrators**.
4. The Administrators window is displayed. Click **Add**.
5. The Select Users, Computers, or Groups window is displayed. Enter the domain controller user in the object names to select field. Click **Check Names**. Click **OK** twice to exit each window.
6. Click **Configuration** > **Services**.
7. Perform these steps for each of the services in the content pane:

- IBM Tivoli Storage Productivity Center - Data Server
- IBM WebSphere Application Server V6.1 - CSM
- IBM WebSphere Application Server V6.1 - DeviceServer
- Tivoli Integrated Portal - TIPProfile_Port_16310

a. Double-click the service.
b. Click **Properties**.
c. On the Properties window, click the **Log On** tab.
d. On the Log On tab, select **This account**. Enter the domain controller password in the Password and Confirm password fields. Click **OK**.
e. You see the message:

   ```
   The account domain_controller\user_ID has been granted the Log On
   As a Service right.
   ```

   Click **OK**.
f. Stop and restart each service.

8. You can now log in to Tivoli Storage Productivity Center (as well as Tivoli Storage Productivity Center for Replication) using the domain controller and user ID.

Configuring IP addressing

This section provides information about configuring IP addressing.

Configuring Tivoli Storage Productivity Center with multiple IP addresses

If the system where Tivoli Storage Productivity Center is to be installed has multiple IP addresses, then a configuration value must be set manually as a post-installation task using the **tpctool setdscfg** command. The value to be set is for the local IP address which must be used for subscription for CIM Indications for CIM agents.

Note: This task does not apply to storage systems that use the native interfaces (DS8000, XIV system, SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000).

If you are using IPv6 machines, see "Planning for Internet Protocol Version 6" on page 51.

For multiple IPv6 addresses, the IPv6 address to use for CIM indication subscription by Tivoli Storage Productivity Center can be specified by setting the property `System.LocalIPv6Address` as described.

With dual stack IPv4 and IPv6 Tivoli Storage Productivity Center servers, two IP addresses are required to subscribe to IPv4 CIMOMs and IPv6 CIMOMs. The configuration property `System.LocalIPv6Address` is used for IPv6 CIMOMs and the property `System.LocalIPAddress` is used for IPv4 CIMOMs.

To change the IP address, follow these steps:

1. Open a command prompt window on the server system.
2. Change to the following directory:

   ```
   cd TPC_installation_directory\cli
   ```

3. Enter the following command:

```
tpctool setdscfg -user user_ID -pwd password -url
host:port -property System.LocalIPv6Address value
```

Where:

user_ID
is the user ID.

password
is the password for the user.

host is either the host name or IP address of the system running Tivoli Storage Productivity Center.

port is a valid port number for the HTTP service of the Device server (the default is 9550).

value is the local IP address which must be used for subscription for CIM Indications for CIM agents.

4. Verify that the command was successful by entering this command:

```
tpctool getdscfg -user user_ID -pwd password
-url host:port -property System.LocalIPv6Address
```

Changing the HOSTS file

When you install Windows on your computers, you must follow these steps to avoid addressing problems with IBM Tivoli Storage Productivity Center. The problem is caused by the address resolution protocol which returns the short name (not-fully qualified host name). This problem can be avoided by changing the entries in the corresponding host tables on the DNS server and on the local computer. The fully qualified host name must be listed before the short name.

To change the HOSTS file, follow these steps:

1. You can find the **HOSTS** file in the `%SystemRoot%\system32\drivers\etc\` directory.
2. You have to edit the **HOSTS** file. An example of a **HOSTS** file is shown.

```
# Copyright (c) 1993-1995 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows NT.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name.
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be
inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#      102.54.94.97     rhino.acme.com    # source server
#       38.25.63.10     x.acme.com        # x client host
192.168.123.146     jason     jason.groupa.mycompany.com
```

3. Enter the fully qualified host name as the first line to be searched in the table. (Lines preceded with a # sign are comment lines.) For example, add the line in bold highlight to the file. The lines must be entered in the order shown.

```
# For example:
#
#      102.54.94.97     rhino.acme.com      # source server
```

```
#       38.25.63.10     x.acme.com          # x client host
192.168.123.146     jason.groupa.mycompany.com       jason
192.168.123.146     jason     jason.groupa.mycompany.com
```

Note: Host names are case sensitive. This requirement is a WebSphere limitation. Check your host name. For example, if your computer shows the name as JASON (uppercase), then you must enter JASON in the HOSTS file.

Starting the Tivoli Storage Productivity Center GUI using a web browser

This topic describes how to start the IBM Tivoli Storage Productivity Center GUI using a web browser.

Overview

The following information pertains to starting Tivoli Storage Productivity Center using Java Web Start:

- The IBM WebSphere Application Server is automatically installed on the same computer as the Tivoli Storage Productivity Center Device server and Data server.
- The Java Web Start application enables a system to communicate with Tivoli Storage Productivity Center that is running on AIX, Linux, UNIX, and Windows systems. Users launch Tivoli Storage Productivity Center using an Internet Explorer or Firefox web browser. The following list shows the supported web browser versions by operating system:
  - AIX: Firefox 2.0 or 3.0
  - Linux: Firefox 2.0 or 3.0
  - Windows: Internet Explorer 7 or 8, or Firefox 2.0 or 3.0
- The remote system must have IBM Java Runtime Environment (JRE) 6 installed. From the remote system, open a web browser window and enter the web address of the target server. The URL is in the format `http://`*`device_server_location`*`:`*`device_server_port`*`/ITSRM/app/welcome.html`.

  Note: The default Device server port is 9550; however, this number might have been changed during installation. Check with your system administrator to see if this port has changed.
- The Java Web Start web browser displays messages based on your system. If your language is not supported, English displays in your web browser.

You can start the Tivoli Storage Productivity Center GUI in one of the following ways:

- Starting the Tivoli Storage Productivity Center GUI without IBM Tivoli Integrated Portal
- Starting the Tivoli Storage Productivity Center GUI through IBM Tivoli Integrated Portal

Note:

- Each time that you start Tivoli Storage Productivity Center GUI from the remote system, Java Web Start compares the version running on the remote system to determine if there is an updated version on the server. If the server version is more current, Java Web Start opens the more current version of the GUI.

- If you are running Java Web Start on a computer that is on the other side of a firewall from the Tivoli Storage Productivity Center server, you must enable some ports (in both directions: in and out). These ports are 9550 (port for the Device server) and 9549 (port for the Data server).
- If you are using Internet Explorer 8, the first time you click **Start Tivoli Storage Productivity Center**, the browser responds with the following message:

  ```
  Do you want to view ONLY the web page content delivered securely?
  ```

  For this message, click **No**.

  Internet Explorer responds with this message:

  ```
  To help protect your security, IE blocked this site from downloading files to
  your computer, Click here for options...
  ```

  Click this message and select **Download file**. You can then launch Java Web Start.

Starting the GUI without IBM Tivoli Integrated Portal

To configure Java Web Start to start Tivoli Storage Productivity Center from a web browser, complete the following steps:

1. Start the Device server on the server where you installed Tivoli Storage Productivity Center.
2. From the remote system, open a web browser window and enter the web address of the target server. The URL is in the format `http://` *`device_server_location`*`:`*`device_server_port`*`/ITSRM/app/welcome.html`.

   Note: The default Device server port is 9550; however, this number might have been changed during installation. Check with your system administrator to see if this port has changed.
3. Tivoli Storage Productivity Center requires that IBM JRE 1.6 is installed on your remote system. If the JRE is not installed, a message displays with a link that enables you to download the JRE.
4. If necessary, download and install IBM JRE 1.6.
5. After the JRE has been installed, Java Web Start on the server determines if Java Web Start is installed on the remote system. If Java Web Start is not installed, a message displays with a link that enables you to download Java Web Start.

After you have completed these steps, you can start the Tivoli Storage Productivity Center from a remote system using the following mechanisms:

- From an Internet Explorer or Firefox browser window, either by typing in the appropriate web address or by clicking a saved bookmark.
- From the Java Web Start Cache Viewer.

Note:

- For Internet Explorer, the browser uses Java Web Start to open the `launch.jnlp` file and continue with the download and launch of Tivoli Storage Productivity Center.

  For Firefox, you see a window asking what you would like to do with the `launch.jnlp` file. Select the **Open with** radio button. Ensure that the program being used to open the file is Java Web Start. Check the box for **Do this automatically for files like this from now on**. Click **OK**.
- For example, if you are using Firefox on AIX or Linux, follow these steps:

1. Download a MIME type editor extension for Mozilla Firefox 2.0 (for example, Mime Edit 0.60) on the AIX or Linux system.
2. Start the Firefox browser.
3. Click **File** > **Open File**.
4. Go to the `mime_edit-0.60-fx+tb.xpi` file and click **Open**. Firefox guides you through the installation of the extension and then tells you to restart Firefox.
5. When you have restarted Firefox, click **Tools** > **Mimetypes**.
6. If there is an entry for "application/x-java-jnlp-file" in the File types list, select it and click **Edit**. In the section When a file of this type is encountered, select **Open it with**. Enter the path to `javaws` on your AIX or Linux system. The most common path is `/usr/java6/jre/javaws/javaws`. Click **OK** to close the dialog used for editing the entry. Click **OK** to close the `Mimetypes` interface.
7. If there is no entry for "application/x-java-jnlp-file" in the File types list, click **New Type**. Enter the following values:
   - MIME Type = application/x-java-jnlp-file
   - Description = Java Web Start
   - Extension = jnlp

   In the section When a file of this type is encountered, select **Open it with**. Enter the path to `javaws` on your AIX or Linux system. The most common path is `/usr/java6/jre/javaws/javaws`. Click **OK** to close the dialog used for editing the entry. Click **OK** to close the `Mimetypes` interface.

Using Internet Explorer or Firefox on Windows (with IBM Tivoli Integrated Portal)

If you are using Internet Explorer or Firefox on Windows, follow these steps:

1. Download and install IBM JRE 1.6 (if you do not have it).
2. Start an Internet Explorer or Firefox web browser, and type the following information in the address bar:

   `http://`*`hostname`*`:`*`port`*

   Where *`hostname`* defines the server that is running Tivoli Integrated Portal such as the server name or IP address and *port* defines the port number for Tivoli Integrated Portal. If the default port was accepted during the installation of Tivoli Integrated Portal, the port number is 16310. Contact your Tivoli Storage Productivity Center administrator if you need to verify the host name and port number.
3. On the IBM Tivoli Integrated Portal logon page, type your Tivoli Storage Productivity Center ID and password in the **User ID** and **Password** boxes and click **Log in**.

   Note: To have administrator permissions in IBM Tivoli Integrated Portal, you must log on as a member of the Tivoli Storage Productivity Center superuser role.
4. In the IBM Tivoli Integrated Portal navigation tree, click Tivoli Storage Productivity Center.
5. On the Tivoli Storage Productivity Center portlets page, click **Start Storage Productivity Center**.
6. For Internet Explorer, the browser uses Java Web Start to open the `launch.jnlp` file and continue with the download and launch of Tivoli Storage Productivity Center.

For Firefox, you see a window asking what you would like to do with the `launch.jnlp` file. Select **Open with**. Ensure that the program being used to open the file is Java Web Start. Check the box for **Do this automatically for files like this from now on**. Click **OK**.

7. Java Web Start continues with the download and start of the Tivoli Storage Productivity Center GUI.

Using Firefox on Linux (with IBM Tivoli Integrated Portal)

If you are using Firefox 3.0 on Linux, follow these steps:

1. Download and install IBM JRE 1.6 (if you do not have it).
2. Start the Firefox web browser, and type the following information in the address bar:

   `http://hostname:port`

   Where *hostname* defines the server that is running Tivoli Integrated Portal such as the server name or IP address and *port* defines the port number for Tivoli Integrated Portal. If the default port was accepted during the installation of Tivoli Integrated Portal, the port number is 16310. Contact your Tivoli Storage Productivity Center administrator if you need to verify the host name and port number.
3. On the IBM Tivoli Integrated Portal logon page, type your Tivoli Storage Productivity Center ID and password in the **User ID** and **Password** boxes and click **Log in**.

   Note: To have administrator permissions in IBM Tivoli Integrated Portal, you must log on as a member of the Tivoli Storage Productivity Center superuser role.
4. In the IBM Tivoli Integrated Portal navigation tree, click Tivoli Storage Productivity Center.
5. On the Tivoli Storage Productivity Center portlets page, click **Start Storage Productivity Center**.
6. For Firefox, you see a window asking what you would like to do with the `launch.jnlp` file. Select **Open with** and enter the path to `javaws` on your Linux system. The most common path is `/opt/ibm/java-i386-60/jre/bin/javaws`. Check the box for **Do this automatically for files like this from now on**. Click **OK**.
7. Java Web Start continues with the download and launch of the Tivoli Storage Productivity Center GUI.

Using Firefox on AIX (with IBM Tivoli Integrated Portal)

If you are using Firefox 2.0 on AIX follow these steps:

1. Download and install IBM JRE 1.6 (if you do not have it).
2. Download a MIME type editor extension for Mozilla Firefox 2.0 (for example, Mime Edit 0.60) on the AIX system.
3. Start the Firefox browser.
4. Click **File > Open File**.
5. Go to the `mime_edit-0.60-fx+tb.xpi` file and click **Open**. Firefox guides you through the installation of the extension and then tells you to restart Firefox.
6. When you have restarted Firefox, click **Tools** > **Mimetypes**.

7. If there is an entry for "application/x-java-jnlp-file" in the File types list, select it and click **Edit**. In the section When a file of this type is encountered, select **Open with**. Enter the path to `javaws` on your AIX system. The most common path is `/usr/java6/jre/javaws/javaws`. Click **OK** to close the dialog used for editing the entry. Click **OK** to close the `Mimetypes` interface.
8. If there is no entry for "application/x-java-jnlp-file" in the File types list, click **New Type**. Enter the following values:
   - MIME Type = application/x-java-jnlp-file
   - Description = Java Web Start
   - Extension = jnlp

   In the section "When a file of this type is encountered" select **Open with**. Enter the path to `javaws` on your AIX system. The most common path is `/usr/java6/jre/javaws/javaws`. Click **OK** to close the dialog used for editing the entry. Click **OK** to close the `Mimetypes` interface.
9. Start the Firefox web browser, and type the following information in the address bar:

   `http://`*`hostname`*`:`*`port`*

   Where *`hostname`* defines the server that is running Tivoli Integrated Portal such as the server name or IP address and *port* defines the port number for Tivoli Integrated Portal. If the default port was accepted during the installation of Tivoli Integrated Portal, the port number is 16310. Contact your Tivoli Storage Productivity Center administrator if you must verify the host name and port number.
10. On the IBM Tivoli Integrated Portal logon page, type your IBM Tivoli Storage Productivity Center ID and password in the **User ID** and **Password** boxes and click **Log in**.

    Note: To have administrator permissions in IBM Tivoli Integrated Portal, you must log on as a member of the Tivoli Storage Productivity Center superuser role.
11. In the IBM Tivoli Integrated Portal navigation tree, click Tivoli Storage Productivity Center.
12. On the Tivoli Storage Productivity Center portlets page, click **Start Storage Productivity Center**.

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 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.

Configuring the DB2 JDBC driver

This topic describes how to configure the DB2 JDBC driver for the Tivoli Common Reporting component.

To configure the DB2 JDBC driver, follow these steps:

1. Stop the IBM Tivoli Integrated Portal.
2. Copy the following files:
   - db2jcc.jar
   - db2jcc_license_cu.jar

   Copy the files from this directory:

   Windows
   `C:\Program Files\IBM\SQLLIB\java`

   UNIX or Linux
   usr_or_opt`/IBM/SQLLIB/java`

   To this directory:

```
C:\Program Files\IBM\tivoli\tip\products\tcr\
lib\birt-runtime-2_2.1\ReportEngine\plugins\
org.eclipse.birt.repot.data.oda.jdbc_2.2.1_r22x_v20070919\
dirvers
```

3. Start the IBM Tivoli Integrated Portal.

Configuration guidelines for 500 or more agents

This topic provides information if you have 500 or more agents in IBM Tivoli Storage Productivity Center.

If you have 500 or more agents for the Data Server, follow these steps.

Note: Probes can overwhelm Data Server resources, because by default, they can return information all at once.

1. Probe the machines at least once a day or more, depending on when you want to test for alert conditions. This action applies to alert conditions other than directory alerts, quotas, or constraints.
2. If you use anything but the "ALL" groups (ALL file systems, ALL computers), you need to manually populate the groups.
3. Always run a probe before a scan. Set the following parameters in the `server.config` file:

   MaxConnections=1200
   The default is 500. Agents can have multiple connections to the server.

   routerThreads=3 (max)
   Incoming connections need to be routed to the correct Data Manager "service" queue and can stack up behind this thread. You can watch this behavior by watching the connections in the "server service" and the "agent service" in the GUI. The server service runs the router and the agent service is where the connections are queued once routed and saved by any of three threads here to the repository.

4. Set the following parameter in the `Scheduler.config` file:

   MaxSubmitthreads=8
   Tells how many threads are used to tell the agents to start a job. Agent connections can queue up the scheduler service. After a job is run, the agent makes a connection to communicate with this thread to give it the job status.

Using a master image to set up Storage Resource agent servers

If you use a master operating system image to deploy new servers in your environment, you can include the Storage Resource agent on that master image. The master image enables the agents to start and register with the Tivoli Storage Productivity Center server automatically upon deployment. This support applies only to Storage Resource agents running in daemon mode.

The default agent directory is:

- For Windows: `C:\Program Files\IBM\TPC\agent`
- For UNIX: `/opt/IBM/TPC/agent`
- For AIX: `/usr/IBM/TPC/agent`

Follow these instructions to prepare the master image for the Tivoli Storage Productivity Center agents:

1. Install the Storage Resource agent in daemon mode.
2. Stop the Storage Resource agent.

   For the Windows system: Click **Start** > **Settings** > **Control Panel** > **Administrative Tools** > **Services**. Stop the following service: **IBM Tivoli Storage Resource Agent -** ***directory***. *directory* is where the Storage Resource agent is installed. The default directory is `TPC_installation_directory\agent`.

   For the UNIX or Linux system, run the following commands:

   ```
   cd /opt/IBM/TPC/agent/bin/
   ./agent.sh stop
   ```
3. Create two files in the root directory for the agent.

   REGISTERSRA
   Uppercase is important here with no file extension. This file causes the agent to run a probe, the register with the server.

   REGISTERSRA_REGENGUID
   Uppercase is important here with no file extension. This file causes the agent to regenerate a Globally Unique Identifier (GUID), run a probe, then register with the server.

   The default agent directories are:

   For the Windows system: `C:\Program Files\IBM\TPC\agent`.

   For the UNIX or Linux system: `/opt/IBM/TPC/agent`.
4. Delete the contents of the *`agent_installation_directory`*`/logs`. This action clears existing messages so that you can view new messages.
5. Make the master image copies from this Tivoli Storage Productivity Center agent system.
6. When a new system is preinstalled with this image and started, the REGISTERSRA or REGISTERSRA_REGENGUID file is run. Check for the new agent in the GUI by clicking **Administrative Services** > **Data/Storage Resource Agents**. You must refresh the views.

Configuring LUN provisioning for Oracle Solaris

IBM Tivoli Storage Productivity Center for Data provides a file system extension feature that can be used to automatically increase file system capacity for managed hosts when utilization reaches a specified level. This function allows for the automatic provisioning of (TotalStorage Enterprise Storage Server, DS6000, DS8000) LUNs when there is not enough space available in a volume group to extend a file system. This section provides information about LUN provisioning for Solaris.

LUNs can be provisioned for file system hosts running Solaris, but the hosts must be configured to avoid a reboot after provisioning. Before you install the Data agent, run the following steps:

1. Assign TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUNs to Solaris Host Bus Adapters (HBAs).
2. Modify the HBA configuration file to include persistent name binding.
3. Modify the SCSI Disk configuration file to allow the maximum number of LUNs.
4. If you are using multipathing, ensure that TotalStorage Enterprise Storage Server, DS6000, or DS8000 multipaths are detected by the Veritas Dynamic Multipathing (VxDMP) utility.

This section provides basic instructions for performing these configuration steps. For detailed information, see the HBA and VxDMP documentation.

Assigning TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUNs to Oracle Solaris HBAs

This section provides information about assigning TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUNs to Solaris HBAs.

You must assign at least one TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUN to each HBA on the Solaris host.

If you are using multipathing, there are different ways to configure either the host and TotalStorage Enterprise Storage Server, DS6000, or DS8000. For example:

- For an TotalStorage Enterprise Storage Server, DS6000, or DS8000 without internal multipath configuration, assign the same LUNs to the World Wide Port Node (WWPN) of each HBA.
- For an TotalStorage Enterprise Storage Server, DS6000, or DS8000 with internal multipath configuration, assign the LUNs to the WWPN of one HBA or assign the same LUNs to the WWPNs of two or more HBAs.

Modifying the HBA configuration file

The HBA configuration file must be modified to include Persistent Name Binding on HBAs and targets so that both the controller and target numbers remain the same across system reboots. This section provides information about what to modify in the configuration file.

The HBA configuration file (for example, `qla2200.conf`) must be modified to include Persistent Name Binding on HBAs and targets so that both the controller and target numbers remain the same across system reboots. You must reboot the system with the new configuration for the changes to take effect.

QLogic QLA2200 and QLA2300 HBAs have been tested for use with IBM Tivoli Storage Productivity Center. You can use the QLogic SANblade Control FX (**scfx**) application to modify the configuration file for these HBAs. The **scfx** application is included as part of the device driver installation package. The **scfx** application is installed in the `/opt/QLogic_Corporation/SANblade_Control_FX` directory.

Setting Persistent Name Binding for QLA2200 and QLA2300 HBAs by using the scfx command

This section describes how to set Persistent Name Binding in the HBAs by using the **scfx** command for LUN provisioning under Oracle Solaris.

Follow these steps:

1. Install the QLogic HBA Driver, Common API Library, and QLogic SANblade Control FX (**scfx**) application if you have not already done so. For installation instructions, see the *SANblade 2200 Series User's Guide* or *SANblade 2300 Series User's Guide*. After these packages are installed successfully, restart and reconfigure the system by using the **reboot -- -rv** command.
2. After the system is rebooted, use **scfx** to configure Persistent Bind on HBAs and Targets in the `/kernel/drv/qla2xxx.conf` file.
   a. Start the **scfx** application. For example:

      ```
      # /opt/QLogic_Corporation/SANblade_Control_FX/scfx
      ```

      The main window of the **scfx** application consists of three sections:

      Menu Bar
      The menu bar provides three options: **File**, **Tools**, and **Help**.

HBA Tree

The HBA Tree displays the host with its connected adapters (HBAs), devices and LUNs. The HBAs are displayed with a model name and instance number. For example, `Adapter 2200 (Instance #0)`. If a device is connected to an HBA, it has a plus sign (+) to the left of the HBA, which can be expanded to view the list of attached devices. The devices are listed with their World Wide Port Names (WWPN).

Click the plus sign next to a device to expand the tree to show all the LUNs in that device. For a RAID device, such as an TotalStorage Enterprise Storage Server, DS6000, or DS8000, there are multiple LUNs per device.

Note: Expand all the devices to search the TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUNs assigned to the system and note the WWPN of the target device. This information is required to identify the SCSI Target ID assigned or specified for the Persistent Bind Targets Setting.

Tabbed Pages

The contents of the Tabbed Pages changes depending on what is currently selected in the HBA Tree.

b. Select an HBA.

Select an adapter in the HBA Tree. The Tabbed Pages show the **HBA Information**, **HBA Options**, **Target Settings**, **Boot Device**, **Diagnostics**, and **Utilities** tabs.

c. Select the Persistent Bind HBA Setting.

Click the **HBA Options** tab. In the **Select Parameter Section** drop-down list, select **Advanced Host Parameters**. Select the check box for **Persistent Bind HBA**. Click **Save**.

d. Select the Persistent Bind Target Setting.

Click the **Target Settings** tab. Select the check box for each target in the **Bind** column. If the check boxes are already checked and disabled, proceed to the next step. In the **Target ID** column, you can either accept the pre-selected SCSI Target ID or change to a different value. Each SCSI target ID must be unique and range from 0 to 255.

Note: Write down the selected Target ID for each TotalStorage Enterprise Storage Server subsystem device.
Click **Save**.

e. Repeat Steps b through d for the next HBA.

f. Exit the **scfx** application.

From the **Menu Bar**, select **File | Exit**. A Reboot Reminder dialog is displayed. Click **OK** to exit.

3. Restart and reconfigure the system by using the **reboot -- -rv** command.

Modifying the SCSI disk configuration file

You must configure the SCSI disk configuration file for the maximum number of LUNs per target for LUN provisioning for Oracle Solaris.

You must configure the SCSI disk (`sd.conf`) configuration file for the maximum number of LUNs (256) per target. The system must then be rebooted with the new configuration for the changes to take effect. Follow these steps:

1. Identify the SCSI Target ID assigned to the TotalStorage Enterprise Storage Server.
2. Edit the `/kernel/drv/sd.conf` file to include all the possible target and LUN mappings for the RAID device. For example, assume the SCSI Target ID assigned for an TotalStorage Enterprise Storage Server is 2. You can allow up to 256 LUNs (0 - 255) for this target:

```
name="sd" class="scsi" target=0 lun=0;
name="sd" class="scsi" target=1 lun=0;
name="sd" class="scsi" target=2 lun=0;
name="sd" class="scsi" target=2 lun=1;
name="sd" class="scsi" target=2 lun=2;
 :
 :
name="sd" class="scsi" target=2 lun=253;
name="sd" class="scsi" target=2 lun=254;
name="sd" class="scsi" target=2 lun=255;
name="sd" class="scsi" target=3 lun=0;
name="sd" class="scsi" target=4 lun=0;
 :
 :
name="sd" class="scsi" target=253 lun=0;
name="sd" class="scsi" target=254 lun=0;
name="sd" class="scsi" target=255 lun=0;
```

   In this example, the system can detect up to 256 targets with 1 LUN (for example, multiple RAID devices with a total of 256 LUNs) and up to 256 LUNs for target 2 (for example, a RAID device with a total of 256 LUNs).
3. Restart and reconfigure the system by using the **reboot -- -rv** command.

Checking for TotalStorage Enterprise Storage Server, DS6000, or DS8000 multipaths in VxDMP

If you are using IBM TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUNs with multipaths, you must ensure that all the paths are detected by Veritas Dynamic Multipathing (VxDMP) utility. This section provides information about how to check for multipathing in the VxDMP utility.

The VxDMP utility is an administrative interface to the Veritas Volume Manager (VxVM) Dynamic Multipathing (DMP) facility. It lists the paths under a DMP device, gets the DMP device corresponding to a path, lists all the disk controllers on the system, lists all the paths through a host disk controller, lists all the DMP nodes through a disk array, and enables or disables a host disk controller on the system. For more information, and detailed instructions, see the VxDMP documentation.

To list all disk controllers on the system, enter the following command:

```
# vxdmpadm listctlr all
```

The following sample output shows that controllers c3 and c4 are connected to the IBM TotalStorage Enterprise Storage Server with an Enclosure Type of IBM_SHARK and an Enclosure Name of IBM_SHARK0.

```
CTLR-NAME       ENCLR-TYPE      STATE      ENCLR-NAME
=====================================================
c1              Disk            ENABLED      Disk
c3              IBM_SHARK       ENABLED      IBM_SHARK0
c4              IBM_SHARK       ENABLED      IBM_SHARK0
```

To list all subpaths for controller c3, enter the following command:

```
# vxdmpadm getsubpaths ctlr=c3
```

The following sample output shows that the **DMPNODENAME** is the same as the device name for each TotalStorage Enterprise Storage Server LUN:

```
NAME         STATE         PATH-TYPE  DMPNODENAME  ENCLR-TYPE   ENCLR-NAME
=============================================================================
c3t4d0s2     ENABLED        -        c3t4d0s2     IBM_SHARK    IBM_SHARK0
c3t4d1s2     ENABLED        -        c3t4d1s2     IBM_SHARK    IBM_SHARK0
```

To list all subpaths for controller c4, enter the following command:

```
# vxdmpadm getsubpaths ctlr=c4
```

The following sample output shows that the **DMPNODENAME** for each TotalStorage Enterprise Storage Server LUN is from controller c3. This means that VxDMP refers to the TotalStorage Enterprise Storage Server, DS6000 or DS8000 LUNs as devices from controller c3 and mask devices on controller c4 from VxVM:

```
NAME         STATE         PATH-TYPE  DMPNODENAME  ENCLR-TYPE   ENCLR-NAME
=============================================================================
c4t4d0s2     ENABLED        -        c3t4d0s2     IBM_SHARK    IBM_SHARK0
c4t4d1s2     ENABLED        -        c3t4d1s2     IBM_SHARK    IBM_SHARK0
```

How to import authentication information for a Storage Resource agent

The Storage Resource agent is installed as a non-daemon or daemon process. IBM Tivoli Storage Productivity Center stores the authentication information to connect to the host on which the Storage Resource agent has installed for the non-daemon agent. This authentication information can be changed depending on the environment.

To change a Storage Resource agent's authentication information (for non-daemon service), follow these steps:

1. Export the authentication information for a Storage Resource agent. For more information, go to the information center. Search for **Exporting a Storage Resource agent list**.
2. The data file exported contains information such as the host name, user ID, password, certificate location, and passphrase for every agent selected. The information is separated by the pipe character (|). For example,

   agent_host|*user*|*password*|*certificate*|*passphrase*

   You can update the password or passphrase in encrypted format or plain text format. If you want to update the password or passphrase in encrypted format, then you can use the **tpctool**. For example go to this directory and run the **tpctool**:

   ```
   cd TPC_install_directory/cli
   tpctool encrypt string_to_be_encrypted
   ```

   This generates an encrypted string. Place this string in the data file to be imported and add @ENC@ to the end of the encrypted string. For example,

   agent_host|usera|*encrypted_password*@ENC@|certificate|
   encrypted_passphrase@ENC@

   encrypted_password is the encrypted string for the password and *encrypted_passphrase* is the encrypted string for the passphrase.

3. Import the data file. For more information, go to the information center. Search for **Updating a Storage Resource agent list**.

Configuring Tivoli Storage Productivity Center for DS8000 LDAP authentication

You must configure IBM Tivoli Storage Productivity Center to use LDAP for single sign-on support for the DS8000 R4.2.

Overview

Configuring Tivoli Storage Productivity Center and DS8000 for single sign-on involves these general steps:

For Tivoli Storage Productivity Center

1. Extract the certificate. This certificate is used for securing communication between the Authentication Client on the HMC and the Authentication Service (server component) on Tivoli Storage Productivity Center.
2. Create a truststore which includes the certificate from step 1.
3. You need to know the URL for the Authentication Service.

For DS8000

1. Create a Storage Authentication Service (SAS) policy with information collected from Tivoli Storage Productivity Center and the LDAP server.
2. Test the Storage Authentication Service policy using a valid LDAP user mapped to a DS8000 user role in the policy.
3. Activate the Storage Authentication Service policy using a valid LDAP user mapped to the DS8000 administrative user role in the policy.

Procedure for configuring Tivoli Storage Productivity Center

This procedure assumes that Tivoli Storage Productivity Center is set up with the LDAP repository.

To configure Tivoli Storage Productivity Center, complete these steps:

1. You need to know the URL for the Authentication Service.

   Note: An example of the Authentication Service URL is:

   `https://TIP_server_host:16311/TokenService/services/Trust`

   Here is an example:

   `https://tpcserver1.storage.mycompany.com:16311/TokenService/services/Trust`

   The port for the Authentication Service (16311) is one plus the default Tivoli Integrated Portal port (16310). If you change the default port, for example, 17522, then the port number to use for the Authentication Service is one plus the Tivoli Integrated Portal port. In this example, the port number would be:

   `https://tpcserver1.storage.mycompany.com:17523/TokenService/services/Trust`
2. Start Internet Explorer and log in to IBM Tivoli Integrated Portal by entering the following text in the address field:

   https://hostname:port/ibm/console/logon.jsp

where *hostname* defines the server that is running Tivoli Integrated Portal and port defines the port number for Tivoli Integrated Portal. The default port is 16316. Contact your Tivoli Storage Productivity Center administrator to verify the host name and port number. You can log in to Tivoli Integrated Portal with your LDAP user ID and password.

3. Create the truststore in Tivoli Integrated Portal. When logged into IBM Tivoli Integrated Portal, go to the Personal Certificates page for the Default keystore. Click **Security > SSL certificate and key management > Key stores and certificates > NodeDefaultKeyStore > Personal Certificates**. On this page, select the **default** certificate and click **Extract**. On the next page, enter the following information:

   Certificate file name
   Enter a file name for the extracted certificate. This file automatically gets created in `TIP_installation_directory/profiles/TIPProfile/etc/`.

   For Windows, the default directory is `C:\TIP_installation_directory\profiles\TIPProfile\etc\`.

   Accept and select the default data type. Click **OK**.

4. Create the truststore file and import the certificate into the truststore file using the **ikeyMan** tool.
   a. Launch the **ikeyMan** tool.

      For example on Windows:

      ```
      c:\Program Files\IBM\tivoli\tip\bin\ikeyman.bat
      ```

   b. Click **Key Database File > New**. On the New panel, enter the following information and click **OK**:

      Key database type
      Enter or leave the default JKS.

      File Name
      Enter a file name. For example, enter tpc_ess.jks.

      Note: The default location is:

      ```
      c:\Program Files\IBM\tivoli\tip\bin\
      ```

      Location
      Enter a location. For example, enter `c:\tpc\`. Click **OK**.

   c. The next panel prompts you to specify a password for this truststore. Specify a password that you can remember. Click **OK**.

   d. On the next panel, click **Add**. This action opens the Add CA certificate from a file panel. Click **Browse** and select the certificate file you created in step 3. Click **Done**, then click **OK**.

      Note: Look for the certificate file, change the Files of Type to **All files**. Click **Open**.

   e. You see a prompt to specify a label. Provide any label. An example of a label is: ESS_Cert. Click **OK**.

   f. The ESS_Cert is now one of the certificates listed.

   g. Exit the **ikeyman** tool and locate the truststore file (for this example, `tpc_ess.jks`). You need this truststore file and the password for configuring the LDAP-based policy on DS8000.

   h. You are now finished with Tivoli Integrated Portal and the truststore setup.

5. Find the user ID and password that is used in LDAP to use for the DS8000 Storage Authentication Service policy configuration page.

   This user ID is used for authenticating with the Authentication Service. It can be any user ID in LDAP, or a user ID that is also used by Tivoli Storage Productivity Center. This user ID is used as the "Application Client User ID" for a Storage Authentication Service policy on the DS8000.
6. Find the name of a group in LDAP with which you can log in to Tivoli Storage Productivity Center and the DS8000. You would use this LDAP group on the DS8000 also, for mapping to DS8000 roles.

   You can go to the Tivoli Storage Productivity Center **Role-to-Group Mapping** node to find out which LDAP group is mapped to the role in Tivoli Storage Productivity Center.

   To find the LDAP group name, open the Tivoli Storage Productivity Center GUI and click **Tivoli Storage Productivity Center > Configuration > Role-to-Group Mapping**.

   The information gathered in steps 1, 3, 4, 5, and 6 is used on the DS8000 Storage Authentication Service policy creation page.
7. Configure DS8000 R 4.2.

Configuring DS8000 for LDAP authentication

Follow these steps:

1. Add the IP address of the DS8000 Hardware Management Console to the Internet Explorer list of trusted sites using the following steps:
   a. Open the Internet Explorer by clicking the Internet Explorer icon located on the Quick Launch toolbar.
   b. From the Internet Explorer toolbar, click **Tools** > **Internet options**.
   c. Click the **Security** tab, click the **Trusted sites** icon and then click **Sites**.
   d. In the **Add this web site to the zone** field, type the IP address of the DS8000 Hardware Management Console (HMC). Click **Add** and the IP address is added to the **Websites** field.
   e. Click **Close** and then click **OK** to exit the Internet Options window, and then close the Internet Explorer.
2. To access the DS8000 GUI, complete the following steps:
   a. In the Tivoli Storage Productivity Center GUI, expand the **Element Manager** tree.
   b. Click DS8000.
   c. Complete one of the following steps:
      - In the content pane, select a DS8000 system.
      - Click **DS8000 Element Manager > Select Action > Add Element Manager**.
   d. In **Select Action**, select **Launch Default Element Manager**. This action opens the DS8000 Storage Manager GUI so that you can administer DS8000. Enter the user name and password and click **OK**.
3. On the DS8000 Storage Manager Welcome page, click **Real-time manager > Monitor System > User Administration**.
4. On the User and Authentication Policy Administration Summary page, select a Complex Name. Under the Select action menu, select **Create Storage Authentication Service Policy**.

5. The Authentication Service Configuration page is displayed. Enter the following information:
   - Policy Name
   - Authentication Service URL (primary)
   - Authentication Service Client User ID
   - Authentication Service Client Password
   - Confirm Authentication Service Client Password

   Click **Next**.

   Note: An example of the authentication URL is:
   `https://TIP_server_host:16311/TokenService/services/Trust`

   Here is an example:
   `https://tpcserver1.storage.mycompany.com:16311/TokenService/services/Trust`

   The port for the Authentication Service (16311) is one plus the default Tivoli Integrated Portal port (16310). If you change the default port, for example, 17522, then the port number to use for the Authentication Service is one plus the Tivoli Integrated Portal port. In this example, the port number would be:
   `https://tpcserver1.storage.mycompany.com:17523/TokenService/services/Trust`
6. The Truststore file Information page is displayed. Enter the following information:
   - Truststore File Location
   - Truststore File Password
   - Confirm Truststore File Password

   Click **Next**.
7. The Map External Users and User Groups to DS8000 User Roles page is displayed. Enter the following information:
   - External Entity Name
   - External Entity Type
   - DS8000 User Role

   Click **Add**. The entry is entered in the table at the bottom of this page. Select the entry you created and click **Next**.
8. The Verification page is displayed. Verify the information and click **Next**.
9. The Summary page is displayed. Click **Activate the Policy** if you want to activate the policy immediately. If you want to test the policy before activating it, do not select **Activate the Policy** and click **Finish** to create the policy. This scenario assumes that you want to test the policy before activating it. You see a message dialog indicating whether the policy was successfully created or not. If the policy was successfully created, close the message dialog.
10. The Manage Authorization Policy page is displayed. Select a policy. Under the Select action menu, click **Test Authentication Policy**.
11. The Test Storage Authentication Service Policy page is displayed. Enter the following information:
    - External User Name
    - External User Password

    Provide an LDAP user ID and password for External User Name and External User Password. The user ID must already be mapped to a valid DS8000 user

role in the Storage Authentication Service policy. This user ID does not have to be in the Administrator group. Click **OK**.

12. The Manage Authentication Policy page is displayed. Select the policy you want. Under the Select action menu, click **Activate Authentication Policy**.
13. The Activate Storage Authentication Service Policy page is displayed. Enter the following information:
    - External User Name
    - External User Password

    Provide an LDAP user ID and password for External User Name and External User Password. The user ID must already be mapped to a valid DS8000 user role in the Storage Authentication Service policy. This user ID must be in the Administrator group. Click **OK**. The policy is now activated. Close the page.

Configuring multiple Tivoli Storage Productivity Center servers with one DS8000 R4.2

You can configure multiple Tivoli Storage Productivity Center servers to use LDAP for single sign-on support for the DS8000 R4.2.

Procedure for configuring multiple Tivoli Storage Productivity Center servers

Follow these steps:

1. Configure one server as described in "Configuring Tivoli Storage Productivity Center for DS8000 LDAP authentication" on page 359. This server is called `TPC_server1`.
2. Install a second Tivoli Storage Productivity Center server with the same LDAP information as the first server. The second server is called `TPC_server2`.
3. Open a command prompt window. Go to the following directory:

   TIP_installation_directory`/bin`
4. Run the **wsadmin** command to export LTPA keys from `TPC_server1` into a file on `TPC_server2`.

   ```
   wsadmin -user TIP_admin_ID -password TIP_admin_password -lang jython
   -port TIP_SOAP_port -host TPC_server1_hostname_or_IP_address
   -f "TPC_install_dir_on_TPC_server2/tip/scripts/exportLTPAKeys.py"
   "LTPA_keys_file_name" LTPA_keys_password
   ```

   An example is:

   ```
   c:\Program Files\IBM\Tivoli\TIP\bin>wsadmin -user tpcsuperuser
   -password tpcsuperuser -lang jython
   -port 16313 -host 9.56.98.41
   -f "c:/program files/ibm/tpc/tip/scripts/exportLTPAKeys.py"
   "c:/share/ltpaKeys_serv1" ltpa123
   ```

   This creates a file named `ltpaKeys_serv1` containing the LTPA keys of `TPC_server1`. The LTPA keys are imported into `TPC_server2`.

   Note: Use forward slashes.
5. In the same command window, run the following command to import the LTPA keys into IBM Tivoli Integrated Portal and then into the Device server.

   ```
   wsadmin -user TIP_admin_ID -password TIP_admin_password -lang jython
   -f "TPC_install_dir_on_TPC_server2/tip/scripts/importLTPAKeys.py"
   "LTPA_keys_file_name" LTPA_keys_password
   ```

An example is:

```
c:\Program Files\IBM\Tivoli\TIP\bin>wsadmin -user tpcsuperuser
-password tpcsuperuser -lang jython
-f "c:/program files/ibm/tpc/tip/scripts/importLTPAKeys.py"
"c:/share/ltpaKeys_serv1" ltpa123
```

Note: Use forward slashes.

6. Change to the Device server WebSphere bin folder and run the same command there.

```
c:\Program Files\IBM\TPC\device\apps\was\bin>wsadmin -user tpcsuperuser
-password tpcsuperuser -lang jython
-f "c:/program files/ibm/tpc/tip/scripts/importLTPAKeys.py"
"c:/share/ltpaKeys_serv1" ltpa123
```

Note: Use forward slashes.

7. The LTPA keys in TPC_server1 and TPC_server2 are now synchronized. You can complete a successful single sign-on launch from TPC_server2 to the DS8000 R4.2. The DS8000 uses the same policy that was set up when you set up TPC_server1.

   The same steps can be used to start the same DS8000 from any number of IBM Tivoli Storage Productivity Center servers.

   Note: This is not a high-availability setup because the policy in DS8000 is still pointing to only one Embedded Security Service, which is that of TPC_server1.

Setting up dual Tivoli Storage Productivity Center servers for high availability

This section describes how to set up dual Tivoli Storage Productivity Center servers for high availability.

Procedure

Follow these steps:

1. Configure one Tivoli Storage Productivity Center server as described in "Configuring Tivoli Storage Productivity Center for DS8000 LDAP authentication" on page 359. In this example, this server is called TPC_server1.
2. Install a second Tivoli Storage Productivity Center server with the same LDAP information as the first server. In this example, this server is called TPC_server2.
3. Open a command prompt window. Go to the following directory:

   TIP_install_directory/bin
4. Run the following WebSphere command to export the LTPA keys from TPC_server1 into a file on TPC_server2.

```
wsadmin -user TIP_admin_ID -password TIP_admin_password
-lang jython
-port TIP_SOAP_port
-host TPC_server1_hostname_or_IP_address
-f "TPC_install_directory_on_TPC_server2/TIP/scripts/exportLTPAKeys.py"
"LTPA_keys_file_name" LTPA_keys_password
```

   Here is an example:

```
C:\Program Files\IBM\Tivoli\TIP\bin> wsadmin -user tpcsuperuser
-password tpcsuperuserpassword
-lang jython
```

```
-port 16313
-host 9.54.91.40
-f "c:/Program Files/IBM/TPC/TIP/scripts/exportLTPAKeys.py"
"c:/share/ltpakeys_serv1" ltpa123
```

This creates a file named ltpakeys_serv1 which contains the LTPA keys of TPC_server1. The LTPA keys are imported into TPC_server2.

Note: Use forward slashes with the **-f** parameter.

5. In the same command window, run the following WebSphere command to import the LTPA keys into Tivoli Integrated Portal and then into the Device server.

```
wsadmin -user TIP_admin_ID -password TIP_admin_password
-lang jython
-f "TPC_install_directory_on_TPC_server2/tip/scripts/importTPAKeys.py"
"LTPA_keys_file_name" LTPA_keys_password
```

Here is an example:

```
C:\Program Files\IBM\Tivoli\TIP\bin>wsadmin -user tpcsuperuser
-password tpcsuperuserpassword
-lang jython
-f "c:/program files/ibm/tpc/tip/scripts/importLTPAKeys.py"
"c:/share/ltpakeys_serv1" ltpa123
```

Note: Use forward slashes with the **-f** parameter.

6. Change to the Device server WebSphere bin folder and run the same command there.

```
c:\Program Files\IBM\TPC\device\apps\was\bin>wsadmin -user tpcsuperuser
-password tpcsuperuserpassword
-lang jython
-f "c:/program files/ibm/tpc/tip/scripts/importLTPAKeys.py"
"c:/share/ltpakeys_serv1" ltpa123
```

Note: Use forward slashes with the **-f** parameter.

7. You need to synchronize the second authentication service (on TPC_server2) with the correct LTPA keys:

```
C:\Program Files\IBM\Tivoli\TIP\bin>wsadmin -user tpcsuperuser
-password tpcsuperuserpassword
-lang jython
-c "AdminTask.importESSLTPAKeys
('[-pathname c:/share/ltpakeys_serv1 -password ltpa123]')"
```

Restart the second Tivoli Integrated Portal server (on TPC_server2).

8. The LTPA keys in TPC_server1 and TPC_server2 are now synchronized.
9. If you are using the Java client. add the SSL certificates for all servers to the truststore file. Follow these steps:
   a. Log in to Tivoli Integrated Portal on TPC_server1 and extract the certificate. In this example, the certificate is named cert1.cer.
   b. Log in to Tivoli Integrated Portal on TPC_server2 and extract the certificate. In this example, the certificate is named cert2.cer.
   c. On TPC_server1, take the cert1.cer and cert2.cer certificates at one location and use the Java keytool command to create a truststore.
   d. Go to c:\Program Files\IBM\Tivoli\tip\java\bin and add these two certificates:

```
keytool -import -alias TPCServer1
  -file c:\cert1.cer
  -keystore c:\ess.truststore.jks
  -storetype jks
  -storepass password
keytool -import -alias TPCServer2
  -file c:\cert2.cer
  -keystore c:\ess.truststore.jks
  -storetype jks
  -storepass password
```

e. Verify that the two certificates exist in the keystore by running this command:

```
keytool -list
-keystore c:\ess.trustore.jks
-storepass password
```

This command must list the two aliases (TPCServer1 and TPCServer2) in the keystore.

f. Copy the ess.trustore.jks keystore to c:\Program Files\IBM\TPC\device\conf on TPC_server1 and TPC_server2.

Configuring Tivoli Storage Productivity Center and SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 for single sign-on

To use one user ID and password to access multiple applications, configure Tivoli Storage Productivity Center and SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 for single sign-on.

Overview

Configuring Tivoli Storage Productivity Center and SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 for single sign-on involves these general steps:

For Tivoli Storage Productivity Center

1. You must know the web address for the Authentication Service.
2. You must know the HTTP basic authentication user name and password.

For SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000

Configure the remote authentication service of the cluster.

Configuring Tivoli Storage Productivity Center for single sign-on

This procedure assumes that Tivoli Storage Productivity Center is set up with the LDAP repository. Follow these steps:

1. Extract the Secure Sockets Layer (SSL) certificate in one of the following two ways:

   Note: The format of the certificate must be PEM (Privacy Enhanced Mail).

   - Open the Tivoli Integrated Portal console. The default address for the console is https://*TIP_server_host*:16316/ibm/console/logon.jsp where *TIP_server_host* is the IP address or host name of the Tivoli Integrated Portal server.

a. On the Personal Certificates page for the default keystore, click **Security** > **SSL certificate and key management** > **Keystores and certificates** > **NodeDefaultKeyStore** > **Personal Certificates**.
b. Select the default certificate and click **Extract**.
c. On the resulting page, enter the following information:

Certificate file name
Enter a file name for the extracted certificate. This file is automatically created in *TIP_install_directory*/profiles/TIPProfile/etc/. In Microsoft Windows, the default directory is C:*TIP_install_directory*\profiles\TIPProfile\etc\. Click **OK**.

Data Type
Select Base64-encoded ASCII data.

- From any Linux or UNIX system, run the following Shell command:

```
echo | openssl s_client -connect TIP_server_ip:16311 2>&1 \
| sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p'
```

where *TIP_server_ip* is the IP address of the Tivoli Integrated Portal server.

2. Determine the web address for the authentication service.

Note:

- SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 does not support Domain Name System (DNS) host names. IP addresses must be used.
- The default port for the unencrypted (HTTP) Tivoli Integrated Portal Authentication Service is 16310.
- The default port for the encrypted (HTTPS) Tivoli Integrated Portal Authentication Service is 16311.
- Web address for HTTP:

```
http://TIP_server_ip:port/TokenService/services/Trust
```

where *TIP_server_ip* is the IP address of the Tivoli Integrated Portal server.

- Web address for HTTPS:

```
https://TIP_server_ip:port/TokenService/services/Trust
```

where *TIP_server_ip* is the IP address of the Tivoli Integrated Portal server.

Examples of authentication service address for secure communication:

- IPv4:

```
https://10.1.2.100:16311/TokenService/services/Trust
```

- IPv6:

```
https://[fec0::ffff:214:5eff:fe29:a02d]:16311/TokenService/services/Trust
```

3. Find the authentication service user ID and authentication service password that is used on the cluster for remote authentication. It can be any user ID in LDAP or a user ID that is also used by Tivoli Storage Productivity Center. Use this user ID and password when you configure SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000.

Configuring SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 for remote authentication service

You can use either the SAN Volume Controller, Storwize V7000 Unified or Storwize V7000 Management GUI or the command-line interface (CLI) to configure the remote authentication of a cluster. To configure remote authentication service on SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000, you must have the Security Administrator role.

To use the SAN Volume Controller Console Version 5.1 or later to configure remote authentication of a cluster, follow these steps:

1. Log on to the SAN Volume Controller Management GUI and select **Launch the SAN Volume Controller Console** for the particular cluster.
2. From the main menu, select **Manage Authentication** > **Remote Authentication**.
3. Check **Enabled** to enable remote authentication.
4. Enter the appropriate attributes:
   - Remote service user name
   - Remote service password
   - Protocol HTTP or HTTPS
   - For HTTPS connections, the SSL certificate to upload (the SSL certificate that you extracted from Tivoli Storage Productivity Center)
   - Remote service web address (IPv4 web address or IPv6 address)

To use the SAN Volume Controller Management GUI Version 6.1 or later, Storwize V7000 Unified, or Storwize V7000 Management GUI to configure remote authentication of a cluster, follow these steps:

1. In a browser window, enter the management IP address for the system:
   `https://cluster_domain_name or IP_address`
2. Log onto the management GUI with your username and password.
3. Click **User Management** > **Users** > **Global Actions** > **Configure Remote Authentication**.
4. Select **Enable Remote Authentication** and enter the following information about the remote authentication services:
   - Remote service user name
   - Remote service password
   - Web address for the remote authentication service
5. Click **OK**. If you are using the encrypted authentication service (HTTPS), you will see a message that the SSL certificate is automatically retrieved from the web address. Click **OK** again.

Disabling unlicensed or unused Tivoli Storage Productivity Center components

If your Tivoli Storage Productivity Center server is installed on a system with less than 8 GB of RAM, you might experience performance or stability issues if you enable all Tivoli Storage Productivity Center components (Disk, Data, and Replication).

If this is the case, and you are not licensed for some components, or do not intend to use some components, you can disable components after installation.

Because Tivoli Storage Productivity Center for Replication is the newest integrated component of Tivoli Storage Productivity Center, it might also be the most common component to initially disable when upgrading an existing Tivoli Storage Productivity Center server to version 4. To disable Tivoli Storage Productivity Center for Replication, follow these steps:

On Windows

To stop the Tivoli Storage Productivity Center for Replication server, go to **Start > Settings > Control Panel > Administrative Tools > Services**. Stop the following service:

```
IBM WebSphere Application Server V6.1 - CSM
```

If you want to set your system so that Tivoli Storage Productivity Center for Replication does not start up on system reboot, you must disable this service.

On Linux and AIX

To stop the Tivoli Storage Productivity Center for Replication server, run this command from a command prompt window:

```
/opt/IBM/replication/eWAS/profiles/CSM/bin/stopServer.sh server1
```

If you want to set your system so that Tivoli Storage Productivity Center for Replication does not start up on system reboot, you must go into `/etc/inittab` and remove the line that starts up Tivoli Storage Productivity Center for Replication.

If you are upgrading a stand-alone Tivoli Storage Productivity Center for Replication server to version 4, you will notice you now have access to many new integrated capabilities in Tivoli Storage Productivity Center. If you choose to not take advantage of those capabilities, you can disable them by following these steps:

On Windows

Go to **Start > Settings > Control Panel > Administrative Tools > Services** and stop the following services:

IBM WebSphere Application Server V6.1 - DeviceServer

IBM Tivoli Storage Productivity Center - Data Server

IBM Tivoli Common Agent - *ca_directory*

ca_directory is where the common agent is installed. The default is *TPC_install_directory*`\ca`

IBM Tivoli Storage Resource Agent - *sra_directory*

sra_directory is where the agent is installed. The default is *TPC_install_directory*`\agent`

Tivoli Integrated Portal - TIPProfile_Port_xxxxx_

xxxxx indicates the port specified during installation. The default port is 16310.

IBM ADE Service

The Tivoli Integrated Portal registry.

Note:

- Stop Tivoli Integrated Portal and IBM ADE Service only if no other applications are using these services.
- If you want to set your system so that Tivoli Storage Productivity Center does not start up on system reboot, you must disable these services.

On Linux

To stop the Tivoli Storage Productivity Center services, run these commands in a command prompt window:

```
Data server: /usr_or_opt/IBM/TPC/data/server/tpcdsrv1 stop
Device server: /usr_or_opt/IBM/TPC/device/bin/linux/stopTPCF.sh
Common agent: /usr_or_opt/IBM/TPC/ca/endpoint.sh stop
Storage Resource agent: /usr_or_opt/IBM/TPC/agent/bin/agent.sh stop
```

On AIX

To stop the Tivoli Storage Productivity Center services, run these commands in a command prompt window:

```
Data server: stopsrc -s TSRMsrv1
Device server: /usr_or_opt/IBM/TPC/device/bin/aix/stopTPCF.sh
Common agent: /usr_or_opt/IBM/TPC/ca/endpoint.sh stop
Storage Resource agent: /usr_or_opt/IBM/TPC/agent/bin/agent.sh stop
```

Stop Tivoli Integrated Portal on AIX and Linux

To stop Tivoli Integrated Portal, run this command in a command prompt window:

```
install_directory/tip/profiles/TIPProfile/bin/stopServer.sh server1
```

Wait for the server to complete the operation.

To stop the IBM ADE Service, run this command in a command prompt window:

```
Source the environment:
      . /var/ibm/common/acsi/setenv.sh
   Run this command:
      /usr/ibm/common/acsi/bin/acsisrv.sh -stop
         (wait for operation to complete)
```

Note: Stop Tivoli Integrated Portal and IBM ADE Service only if no other applications are using these services.

Configuring and controlling the Tivoli Storage Productivity Center Monitoring Agent

You might want to change the parameters or start and stop the Tivoli Storage Productivity Center Monitoring Agent. This topic describes how to configure the Tivoli Storage Productivity Center Monitoring Agent to change the Tivoli Enterprise Monitoring Services connection, log path, or authentication parameters. This topic also describes how to start and stop the Tivoli Storage Productivity Center Monitoring Agent.

Configuring the Tivoli Storage Productivity Center Monitoring Agent on Windows

This topic describes how to configure the Tivoli Storage Productivity Center Monitoring agent on Windows.

To configure the Tivoli Storage Productivity Center Monitoring agent on Windows, complete the following steps:

1. Log on to the system with administrator authority on Windows.
2. Open the Tivoli Enterprise Monitoring Services Console. Click **Start** > **All Programs** > **IBM Tivoli Monitoring** > **Manage Tivoli Monitoring Services**.

Note: Depending on how you have your Tivoli Monitoring Services components distributed across your enterprise, you will see different types of Tivoli Enterprise Monitoring Services components displayed.

3. On the Manage Tivoli Enterprise Monitoring Services window, select **Monitoring Agent for TPC**. Right-click **Monitoring Agent for TPC** and click **Reconfigure**.
4. For information about the different parameters to change, see "Installing and configuring the Tivoli Storage Productivity Center Monitoring Agent on Windows" on page 236.

Configuring the Tivoli Storage Productivity Center Monitoring Agent on AIX or Linux

This topic describes how to configure the Tivoli Storage Productivity Center Monitoring Agent on AIX or Linux.

To configure the Tivoli Storage Productivity Center Monitoring agent on AIX or Linux, complete the following steps:

1. Log on to the system with the root user ID or a user ID that has root access rights.
2. Open the Tivoli Enterprise Monitoring Services Console. Open a terminal session window and go to the Tivoli Monitoring Services bin installation directory. The default directory is:

   ```
   /opt/IBM/ITM/bin
   ```

3. Run the following command:

   ```
   ./itmcmd manage
   ```

   Note: Depending on how you have your Tivoli Monitoring Services components distributed across your enterprise, you will see different types of Tivoli Enterprise Monitoring Services components displayed.
4. On the Manage Tivoli Enterprise Monitoring Services window, select **Monitoring Agent for TPC**. Right-click **Monitoring Agent for TPC** and click **Configure**.
5. For information about the different parameters to change, see "Installing and configuring the Tivoli Storage Productivity Center Monitoring Agent on AIX or Linux" on page 240.

Starting and stopping the Tivoli Storage Productivity Center Monitoring Agent

This topic describes how to start and stop the Tivoli Storage Productivity Center Monitoring Agent.

To start or stop the Tivoli Storage Productivity Center Monitoring agent on Windows, complete the following steps:

1. Open the Tivoli Enterprise Monitoring Services Console. Click **Start** > **All Programs** > **IBM Tivoli Monitoring** > **Manage Tivoli Monitoring Services**.
2. On the Manage Tivoli Enterprise Monitoring Services window, select **Monitoring Agent for TPC**. Right-click **Monitoring Agent for TPC**. Click **Start** to start the agent, **Stop** to stop the agent, or **Recycle** to recycle the agent.
3. Wait for the agent to start, stop, or recycle.

To start or stop the Tivoli Storage Productivity Center Monitoring agent on AIX or Linux, complete the following steps:

1. Open the Tivoli Enterprise Monitoring Services Console. Open a terminal session window and go to the Tivoli Monitoring Services bin installation directory. The default directory is:

   ```
   /opt/IBM/ITM/bin
   ```

2. Run the following command:

   ```
   ./itmcmd manage
   ```

3. On the Manage Tivoli Enterprise Monitoring Services window, select **Monitoring Agent for TPC**. Right-click **Monitoring Agent for TPC**. Click **Start Service** to start the agent, **Stop Service** to stop the agent, or **Recycle Service** to recycle the agent.
4. Wait for the agent to start, stop, or recycle.

Installing and configuring the Tivoli Storage Productivity Center server with multiple NIC cards

If your Tivoli Storage Productivity Center server has multiple network interface cards (NIC), install the Tivoli Storage Productivity Center server using an IP address. After you install the Tivoli Storage Productivity Center server, all incoming and outgoing communication are successfully handled.

Installing Tivoli Storage Productivity Center for a multiple network configuration

If the Tivoli Storage Productivity Center server you are installing has multiple NIC, and is configured to use multiple network addresses, ensure that you use the appropriate IP address during installation. If you set up the DNS to resolve the host names, you can use fully qualified host names instead of IP addresses during installation.

- On the "User ID and password, and server information" page of the Tivoli Storage Productivity Center installation wizard, enter the server IP address or the fully qualified host name. The following example shows the typical installation panel where you can specify the Tivoli Storage Productivity Center server IP address for **Server name**.

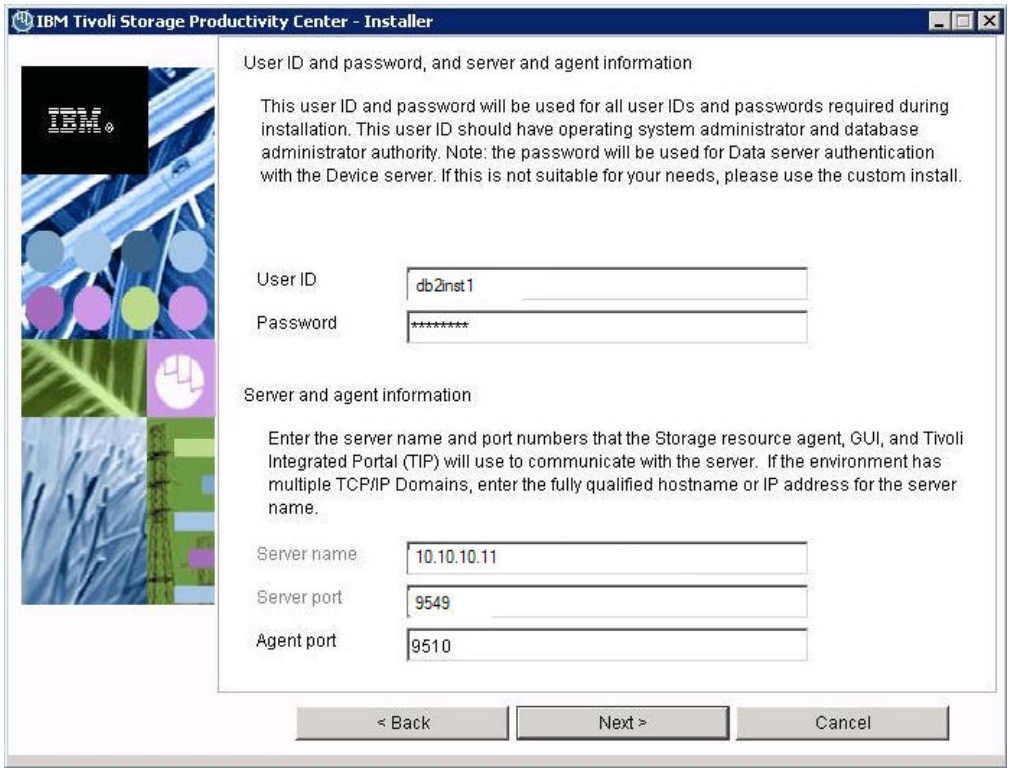


Figure 43. User ID and password panel

- On the "Data server, Device server, and Storage Resource agent information" page of the Tivoli Storage Productivity Center installation wizard, enter the appropriate IP address for both the Data server and Device server. The following example shows where you can define the server name during custom installation.

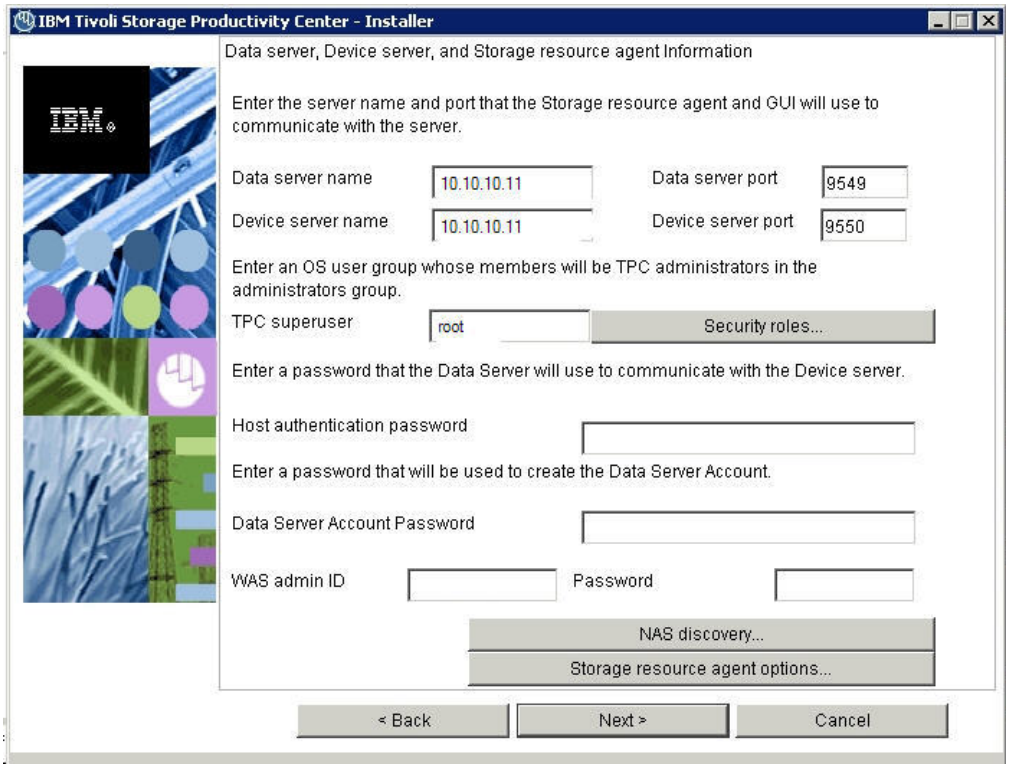


Figure 44. Data server, Device server, and Storage Resource Agent information panel

Outgoing communication initiated by the Tivoli Storage Productivity Center server

All the outgoing communication that is initiated by the Tivoli Storage Productivity Center server is not affected if the server is configured for a multiple network environment.

For example, if you have a Tivoli Storage Productivity Center server with two IP addresses: 10.10.10.11 and 9.9.9.10, and 10.10.10.11 is used during installation, all outgoing transmissions can be sent to the devices and agents in both networks.

The following list includes examples of outgoing communication that is initiated by the Tivoli Storage Productivity Center server:

Storage systems using native interfaces
Run probe, performance management, and provisioning jobs, and collect data events (SAN Volume Controller, Storwize 7000, and XIV Systems)

Switches (SNMP)
Run SNMP discovery and probe jobs

CIM agents
Run discovery, probe, and performance management jobs, and provisioning jobs

VMware vSphere or vCenter
Run discovery and probe jobs

Agents (Storage Resource agents, Data agents, and Fabric agents)
Deploy agents, run probe, discovery, scan, and batch report jobs, and run scripts

Tivoli Storage Productivity Center servers
Run probe jobs

Incoming communication that is initiated by the devices, agents, and GUI

Incoming communication that is initiated by the devices or agents can work with only the IP address that is specified during the installation with the exception of DS8000 events.

For DS8000 events, the Tivoli Storage Productivity Center server must initiate and establish a socket connection directly with the Hardware Management Console (HMC) to receive events. The DS8000 HMC uses that socket connection to send events. As long as the Tivoli Storage Productivity Center server can initiate the communication to the HMC, DS8000 events can be received.

Tivoli Storage Productivity Center informs devices and agents to initiate communication to the IP address provided during the installation. This example uses the IP address 10.10.10.11. However, depending on the communication, you might be able to change the IP address. For example, Tivoli Storage Productivity Center does not configure SAN switches to send SNMP traps to Tivoli Storage Productivity Center, so you can use either 9.9.9.10 or 10.10.10.11.

The following list includes examples of incoming communication that are initiated by the devices, agents, and the GUI:

DS8000 events
Events sent by the HMC to the Tivoli Storage Productivity Center server

SNMP trap notifications
SNMP traps sent from the switches and other devices

CIM indications
Indications sent by the CIM agents

Servers (agents)
Job results and registration

Tivoli Storage Productivity Center GUI
Any request.

CIM indications

A CIM indication is an event that occurs on a managed object, for example, the completion or failure of an operation. The CIM indications are managed by the CIM object manager. Tivoli Storage Productivity Center uses the CIM agents for the managed objects to gather information about the device.

Manually customize CIM indications on a Tivoli Storage Productivity Center system that has multiple IP addresses. To configure Tivoli Storage Productivity Center to receive CIM indications in an IPv4, IPv6, and dual stack (IPv4 and IPv6) environment, see “Configuring Tivoli Storage Productivity Center with multiple IP addresses” on page 345.

Note: The manual customization task does not apply to storage devices that use the native interfaces.

Creating an SSH certificate for the root user ID

You can create a Secure Shell (SSH) certificate for authentication for the Virtual I/O Server. Follow the certificate-generation instructions. However, if you want to use Telnet to connect to the Virtual I/O Server using the padmin user ID, you must follow this procedure.

Procedure

To create an SSH certificate using the padmin user ID, follow these steps:

1. Telnet to the remote system using the padmin user ID.
2. Set up the AIX environment. Run the following command:

   ```
   oem_setup_env
   ```
3. Change to the following `/.ssh` directory.
4. Enter **ssh-keygen**. Accept the default names (for example, id_rsa).
5. Enter the passphrase. Two new files are created:

   id_rsa This is the private key.

   id_rsa.pub
   This is the public key.
6. Create an `authorized_key` file in the same location as the `id_rsa.pub` file. Enter the following command:

   ```
   cat >> id_rsa.pub >> authorized_keys
   ```

   The following example shows the command input and output (the commands are in bold):

   ```
   # ssh-keygen

   Generating public/private rsa key pair.
   Enter file in which to save the key (//.ssh/id_rsa):
   Enter passphrase (empty for no passphrase):
   Enter same passphrase again:
   You identification has been save in //.ssh/id_rsa.
   Your public key has been save in //.ssh/id_rsa.pub.
   The key fingerprint is:
   xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx root@<server>

   # cat id_rsa >> authorized_keys
   # ls -l

   -rw-r--r-- 1 root system 1743 Oct 15 09:40 authorized_keys
   -rw--- 1 root system 1743 Oct 15 09:39 id_rsa
   -rw-r--r-- 1 root system 399 Oct 15 09:39 id_rsa.pub
   ```
7. Copy the `id_rsa` (private key) to your server machine.

   Note: You must copy the file in binary mode.
8. To connect to the remote system using the private key, enter the following information in the Remote Agent Machines window when you install the Storage Resource agent:
   - User
   - Certificate Location (`c:\keys\id_rsa`)
   - Passphrase

Configuring DB2, AIX, and Linux for IPv6-only environment

Use this information to configure DB2, AIX, and Linux for an IPv6-only environment.

Configuring the AIX system for IPv6 only

Note: For IPv6 support, the AIX operating system must have level TL 5300–06 installed.

Follow these steps:

1. Obtain the most recent versions of **openssh** and **openssl** packages for AIX and install them. Some older version of **openssh** does not work in an IPv6-only environment.
2. Change **sshd** (Secure Shell Daemon) on AIX system to accept IPv6 connections.
   a. In the `/etc/ssh/sshd_config` file, uncomment the line "ListenAddress:".
   b. Restart **sshd** with the following commands:
      ```
      stopsrc -g ssh
      startsrc -g ssh
      ```
   c. From another IPv6 system, verify that you contact AIX over IPv6 (using ssh).
3. In SMIT, set the IPv4 address to 0.0.0.0 for all interfaces. Save the file.
4. Edit the `/etc/resolv.conf` file to use IPv6 DNS server or servers.

Configuring DB2 on AIX for IPv6 systems

To get DB2 on AIX to work on IPv6 systems, follow these steps:

1. Identify the host name used by DB2 in the `db2nodes.cfg` file:
   ```
   # cat ~db2inst1/sqllib/db2nodes.cfg
   0 myhost 0
   #
   ```
2. Edit the `/etc/hosts` file and make sure that the host name found in the `db2nodes.cfg` file resolves to an IPv6 address. Use the **vi** editor to verify that the host name is not on any line with an IPv4 address. In particular, ensure that the host name is not listed as an alias for the IPv4 loopback address 127.0.0.1.
   ```
   # vi /etc/hosts
   127.0.0.1 loopback localhost
   ::1 localhost
   2001:db8:0:0:209:6bff:fe09:63fa myhost.mydomain myhost
   ```
3. Stop DB2 and set DB2 to use IPv6 addressing. Restart DB2.
   a. Source the DB2 profile:
      ```
      . ~db2inst1/sqllib/db2profile
      ```
   b. Stop DB2:
      ```
      db2stop
      ```
   c. Configure DB2 to use IPv6.
      ```
      db2set
      ```
      An example of the output is: DB2FCMCOMM=TCPIP6.
   d. Start DB2.
      ```
      db2start
      ```

In some installations, the AIX server does not have a graphical console attached to the server. In this situation, you can select another system with an X11 server to display the IBM Tivoli Storage Productivity Center installation and Tivoli Storage Productivity Center application. The X11 server must have IPv6 configured and an SSH client installed. Open an SSH connection from a shell on the X11 server desktop with the -X option to permit forwarding of X11 applications from the remote AIX server. Start the Tivoli Storage Productivity Center installation program or application from the SSH shell.

```
ssh -X my_IPv6_host
/opt/IBM/TPC/gui/TPCD.sh
```

Configuring DB2 on Linux for IPv6-only systems

To get DB2 on Red Hat Enterprise Linux Version 4 systems to work in an IPv6-only environment, follow these steps:

1. Install DB2 in dual-stack configuration.
2. Stop DB2 and set DB2 to use IPv6 addressing:
   a. As the root user from the UNIX command line, run this command:
      ```
      su - db2inst1
      ```
   b. Stop DB2 by running this command:
      ```
      db2stop
      ```
   c. Configure DB2 to use IPv6 by running this command:
      ```
      db2set
      ```

      An example of the output is: DB2FCMCOMM=TCPIP6.

      Note: The host name in the `db2nodes.cfg` file resolves to an IPv6 address. This action can require you to change the domain or search directive in the `/etc/resolv.conf` file to specify a domain in which the host name can resolve to IPv6. You can also edit the `/etc/hosts` file so that the host name resolves to an IPv6 address.
   d. Start DB2 by running this command:
      ```
      db2start
      ```

Chapter 4. Upgrading and migrating the Tivoli Storage Productivity Center family

You can upgrade previous TotalStorage Productivity Center 3.3.2 or later releases to Tivoli Storage Productivity Center version 4.2.2.1. You can upgrade previous IBM TotalStorage Productivity Center for Replication version 3.4 or later to Tivoli Storage Productivity Center for Replication version 4.2.2.1. This section provides information about upgrading and migrating.

For information about upgrading releases before Tivoli Storage Productivity Center 3.3.2 and Tivoli Storage Productivity Center for Replication 3.4, see the Tivoli Storage Productivity Center 4.2.2 information center.

When you upgrade Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication, if a component fails to upgrade, then just the component is not upgraded.

If Tivoli Storage Productivity Center exists but Tivoli Storage Productivity Center for Replication does not, then this upgrade is a fresh installation of Tivoli Storage Productivity Center for Replication and an upgrade of Tivoli Storage Productivity Center. If a failure occurs, an error message is displayed but there is no rollback.

If Tivoli Storage Productivity Center for Replication exists and Tivoli Storage Productivity Center does not, then this upgrade is a fresh installation of Tivoli Storage Productivity Center and an upgrade of Tivoli Storage Productivity Center for Replication. If there is a failure in the Tivoli Storage Productivity Center installation, a rollback occurs of Tivoli Storage Productivity Center (Tivoli Storage Productivity Center for Replication remains).

Tivoli Storage Productivity Center for Replication is no longer a stand-alone application. Tivoli Storage Productivity Center Version 4.2.2.1 installs Tivoli Integrated Portal and Tivoli Storage Productivity Center for Replication.

With version 4.2 or later, Tivoli Storage Productivity Center for Replication no longer supports DB2 as the datastore for its operational data. Tivoli Storage Productivity Center for Replication uses an embedded repository for its operational data. The Tivoli Storage Productivity Center for Replication installation program automatically migrates any Tivoli Storage Productivity Center for Replication operational data in an existing Tivoli Storage Productivity Center for Replication DB2 database to the Tivoli Storage Productivity Center for Replication embedded repository as part of upgrading to Tivoli Storage Productivity Center for Replication. New Tivoli Storage Productivity Center for Replication installations use the embedded repository by default.

General information for upgrading and migrating Tivoli Storage Productivity Center components

This section describes general information and tools for upgrading and migrating Tivoli Storage Productivity Center components.

© Copyright IBM Corp. 2001, 2012

Upgrading Tivoli Storage Productivity Center for Replication

With Tivoli Storage Productivity Center for Replication version 4.2 or later, an embedded database is used for the database repository. DB2 is no longer supported as the database repository for Tivoli Storage Productivity Center for Replication.

As a part of the upgrade process, an embedded repository is created and any data from an existing DB2 Tivoli Storage Productivity Center for Replication database is copied to the new database repository. This process is automatic and does not require any input or action from the user.

For a typical upgrade procedure, follow this general procedure:

1. Issue the takeover command to the standby server and then upgrade the standby server.

   Note: Be aware that this action makes both the Tivoli Storage Productivity Center for Replication servers active.
2. Wait for the standby server to complete installation and start up.
3. While upgrading the active server, avoid making any configuration changes to the sessions.
4. Upgrade the active server.
5. If no changes have been made to the configuration while the active server is being upgraded, reestablish the High Availability function from the active server to the standby server. If configuration changes were made to the standby server, synchronize the High Availability function from the standby server to the active server. Then perform a takeover operation and reestablish the High Availability function from the active server to the standby server.

During the initial synchronization, the current information in the database is saved and held until the synchronization is complete. If an error occurs during this process, the server database is restored to its original state before the synchronization process began. If an error occurs during the synchronization process which causes the status to be in the disconnected or inconsistent state, you can reconnect to a synchronized state.

Fix packs and patches

All fix packs and patches use the upgrade procedure to upgrade IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication.

You must have a valid Tivoli Storage Productivity Center license to use the upgrade procedure. For information about fix packs and patches, go to http://www-01.ibm.com/support/docview.wss?&uid=swg21320822.

Performance database migration tool (Tivoli Storage Productivity Center only)

After you upgrade Tivoli Storage Productivity Center, migrate the Tivoli Storage Productivity Center database using the performance database migration tool (**partitiontables.bat** or **partitiontables.sh**).

You should run the performance database migration tool after these Tivoli Storage Productivity Center migration scenarios:

- version 3.x to version 4.1 or later

- version 4.1.x to version 4.2 or later.

Both of these migration scenarios are used in preparation for migrating an earlier version of Tivoli Storage Productivity Center to version 5.1 or later.

Tivoli Storage Productivity Center functions normally without the performance database migration, but with decreased performance as compared to the migrated database. The Tivoli Storage Productivity Center version 4.2 or later database has been changed to improve performance by taking advantage of the table partitioning and multidimensional clustering features of DB2. These changes are automatically included when you install Tivoli Storage Productivity Center. Because the performance database migration tool can take a long time to run (depending on the size of the database to be migrated), you can run the migration tool at a time that is convenient for you. You only have to run the performance database migration tool once. It is a good practice to run the performance database migration tool before you apply any fix packs or PTF for Tivoli Storage Productivity Center.

When you run the performance database migration tool, you can check the progress of the database migration in the `migrateTable.log` file in the *`TPC_install_directory`*`\data\server\tools` directory. You see warning messages if the migration cannot be performed on the database or if the migration has been previously completed. The performance database migration tool prints out messages indicating which table is currently being migrated and which subsystem ID or switch ID is currently being migrated. (The migration tool works with both Disk Manager and Fabric Manager data.) The performance database migration tool can be run more than once if an error occurred during execution. Tables that were successfully migrated during previous attempts are not migrated in subsequent runs.

Note: The performance database migration tools partitiontables.bat and partitiontables.sh must be run from disk 1 of the Tivoli Storage Productivity Center installation image or DVD. For example, `<disk1>/data/scripts/partitiontables.sh`.

Storage Subsystem Credential Migration tool (Tivoli Storage Productivity Center only)

If you have DS8000, XIV, or SAN Volume Controller storage systems, you must migrate your storage system credentials for Tivoli Storage Productivity Center. These storage systems now use the native storage system interface instead of the CIM agents.

You can migrate your storage system credentials in one of the following ways:

- Use the migration tool from the Tivoli Storage Productivity Center DVD. Use the DVD or download the migration tool: `MigrateUserInfo.bat` (for Windows) or `MigrateUserInfo.sh` (for UNIX or Linux). Run the migration tool from the directory `UserMigrationTool`.
- Use the Tivoli Storage Productivity Center GUI to migrate the credentials (see "Update Storage Subsystem Credentials" on page 328)
- Run the Tivoli Storage Productivity Center upgrade installation program and specify that you want to use the migration tool. Run the migration tool during the Tivoli Storage Productivity Center upgrade procedure.

Note: For Windows, the credential migration tool has a dependency on MSVCR90.dll if you run the tool from a command window. If you get an error

exception when running this tool, check to see that your system has this DLL file. If you use the tool from the installation wizard, the wizard automatically loads the required DLL files.

If Tivoli Storage Productivity Center can access the DS8000 using the existing credentials, you can continue to use the existing credentials. For XIV and SAN Volume Controller storage systems, you must manually update the credentials.

The storage system credential migration applies to all DS8000 systems, XIV systems, and SAN Volume Controller systems.

If you have run a CIMOM discovery job for a storage system but have not run a probe job for that system, then upgrade Tivoli Storage Productivity Center, the IP address does not display in the GUI. You must manually enter the IP address for that storage system.

Note: After a successful migration of DS8000, XIV, or SAN Volume Controller credentials, the CIMOMs that manage these devices previously should be removed from the Tivoli Storage Productivity Center configuration. The migration process tries to determine and remove these CIMOMs during migration automatically, but if there are CIMOMs still remaining, they should be deleted manually. For DS CIMOMs managing DS6000 or ESS devices and additional DS8000 devices, the DS8000 devices should be removed from the CIMOM if possible.

Upgrading TotalStorage Productivity Center 3.3.2 or later to Tivoli Storage Productivity Center 4.2.2.1 or later

If you have TotalStorage Productivity Center 3.3.2 or later installed and want to upgrade to Tivoli Storage Productivity Center 4.2.2.1 or later, follow these general steps.

Follow these general steps:

If you have DB2 version 8.1 installed

Follow these steps:

1. Upgrade DB2 8.1 to DB2 version 9.7. See "Upgrading DB2 8.2 to DB2 9.7" on page 386.
2. Upgrade TotalStorage Productivity Center 3.3.2 or later to Tivoli Storage Productivity Center 4.2.2.1 or later. See "Upgrading Tivoli Storage Productivity Center" on page 398.

If you have DB2 version 9.1 installed

Follow these steps:

1. If you have DB2 version 9.1 with fix pack 2 installed, apply fix pack 5.
2. Upgrade TotalStorage Productivity Center 3.3.2 or later to Tivoli Storage Productivity Center 4.2.2.1 or later. See "Upgrading Tivoli Storage Productivity Center" on page 398.

If you have DB2 version 9.1 with fix pack 5 installed

Follow these steps:

1. If you want to stay with DB2 version 9.1 with fix pack 5, then you can upgrade TotalStorage Productivity Center 3.3.2 or later to Tivoli Storage Productivity Center 4.2.2.1 or later. See "Upgrading Tivoli Storage Productivity Center" on page 398.

2. If you want to upgrade to DB2 version 9.7, then upgrade to DB2 version 9.7, then upgrade TotalStorage Productivity Center 3.3.2 or later to Tivoli Storage Productivity Center 4.2.2.1 or later. See "Upgrading Tivoli Storage Productivity Center" on page 398.

If you have DB2 version 9.5 with fix pack 3a installed

Follow these steps:

1. If you have DB2 version 9.5 with fix pack 3a installed, you do not need to upgrade DB2.
2. Upgrade TotalStorage Productivity Center 3.3.2 or later to Tivoli Storage Productivity Center 4.2.2.1 or later. See "Upgrading Tivoli Storage Productivity Center" on page 398.

Upgrading Tivoli Storage Productivity Center for Replication from V3.4 or later to V4.2.2.1 or later

You can upgrade Tivoli Storage Productivity Center for Replication version 3.4 or later to version 4.2.2.1 or later.

Tivoli Storage Productivity Center for Replication is no longer a stand-alone application. Tivoli Storage Productivity Center Standard Edition Version 4.2.2.1 or later installs Tivoli Integrated Portal and Tivoli Storage Productivity Center for Replication. Before you can upgrade, verify that you have the minimum prerequisites required for Tivoli Storage Productivity Center and IBM Tivoli Integrated Portal. For more information about hardware and software requirements, see "Hardware requirements" on page 2 and "Software requirements" on page 5.

You are not required to uninstall the previous version of Tivoli Storage Productivity Center for Replication to upgrade from one version to the next. Version 4.2.2.1 or later can be installed on an existing version 3.4 installation if you have met the hardware and software requirements to support Tivoli Storage Productivity Center and IBM Tivoli Integrated Portal.

Note: Tivoli Storage Productivity Center for Replication automatically migrates the database to an embedded database repository as part of the upgrade procedure.

After you have completed the necessary migration tasks, use the Tivoli Storage Productivity Center installation wizard to complete a typical or custom installation of Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication. For information about installing, see Chapter 2, "Installing the Tivoli Storage Productivity Center family," on page 153.

Upgrading DB2 from DB2 9.1 or DB2 9.5 to DB2 9.7

If you are upgrading DB2 from DB2 9.1 or DB2 9.5 to DB2 9.7, follow this general procedure.

For information about upgrading to DB2 9.7, go to this Web site: http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp.

For the general procedure to upgrade the DB2 database repository to DB2 9.7, follow these steps:

1. Stop the Tivoli Storage Productivity Center services and Agent Manager. For information about stopping the services, go to the Tivoli Storage Productivity Center information center. Search for **Starting and stopping the Tivoli Storage Productivity Center services**.
2. Check to see if your databases are ready for migration:
   a. On DB2 9.1 or DB2 9.5, disconnect all users, stop the back-end process, and stop your existing instances by running the following commands:
      ```
      db2stop force (disconnects all users and stops the instance)
      db2 terminate (terminates back-end process)
      ```
   b. Verify that the local databases that are owned by the current instance are ready to be upgraded and generate a log file. Run the **db2ckupgrade** command.

      Note: For Windows, click **Start** > **Programs** > **IBM DB2** > **Command Line Tools** > **Command Window**.
      Go to the directory for the **db2ckupgrade** command:
      ```
      cd DB2_9.7_installer_kit\ESE\image\db2\Windows\utilities
      ```

      Run this command:
      ```
      db2ckupgrade database_name -l db2ckupgrade.log -u user_name
      -p password
      ```

      where *database_name* is the name of the database and db2ckupgrade.log is the name of the log file created in the current directory. The log file includes error and warning messages. Each time you run this command, it overwrites the existing log file. You can rename the log file to avoid losing the information from a previous run. You must correct these errors before you upgrade.
   c. Check the log file. For a successful run, the output provides this message:
      ```
      db2ckupgrade was successful. Database(s) can be upgraded.
      ```
3. Back up the database. Run the following DB2 command from a command prompt window:
   ```
   DB2 BACKUP DATABASE TPCDB USER user_name USING password
   TO path_to_backup_folder
   ```

   where *user_name* is the user ID, *password* is the password, and *path_to_backup_folder* is the directory path to the backup folder. Test the image integrity by running this command:
   ```
   DB2CKBKP *.001
   ```
4. Run the DB2 extended security command:
   ```
   db2set DB2_EXTSECURITY=NO
   ```
5. Install DB2 9.7. For information about installing DB2 9.7, see “Installing DB2” on page 155. Select the **Install New** option on the Install a Product panel. Install all DB2 add-on products that you installed on your previous DB2 version.
6. Upgrade the DB2 instances.
   a. On DB2 9.1 or DB2 9.5, disconnect all users, stop the back-end process, and stop your existing instances by running the following commands:
      ```
      db2stop force (disconnects all users and stops the instance)
      db2 terminate (terminates back-end process)
      ```
   b. Log on to the DB2 9.1 or DB2 9.5 server with root authority on UNIX or local administrator authority on Windows.

c. Upgrade your existing instances by running the **db2iupgrade** command from the target DB2 version 9.7 copy location.

   For Windows, run this command:

   ```
   "%DB2PATH%"\bin\db2iupgrade InstName /u:user,password
   ```

   Where DB2PATH is set to the location you specified during DB2 Version 9.7 installation, *user,password* are the user name and password under which the DB2 service runs, and *InstName* is the name of the instance.

   For UNIX or Linux, run this command:

   ```
   $DB2DIR/instance/db2iupgrade [ -u fencedID ] InstName
   ```

   Where DB2DIR is set to the location you specified during DB2 Version 9.7 installation, *fencedID* is the user name under which the fenced user-defined functions (UDFs) and stored procedures run, and *InstName* is the login name of the instance owner. This example upgrades the instance to the highest level for DB2 database product that you installed.

d. Check the following configuration file for Windows (2008):

   ```
   C:\ProgramData\IBM\DB2\DB2COPY1\DB2\db2node.cfg
   ```

   Remove the extra line. The line should look like this:

   ```
   0 beesspc BEESPC 0
   ```

   Where beesspc and BEESSPC is the name of the Tivoli Storage Productivity Center server.

e. Log on to the DB2 9.7 server as a user with sufficient authority to start your instance.

f. Restart your instance by running the **db2start** command.

g. Verify that your instance is running on DB2 9.7 by running the **db2level** command. The output provides a string like:

   ```
   DB2 v9.7.x.x
   ```

   Where x is a digit number.

7. Upgrade the DB2 Administration Server.

   a. Log on to the DB2 9.7 server with root authority on UNIX or Linux or local administrator authority on Windows.

   b. Upgrade your existing DB2 Administration Server by running the **dasmigr** command.

      For Windows, run this command:

      ```
      %DB2PATH%\bin\dasmigr
      ```

      Where DB2PATH indicates the location that you specified during DB2 Version 9.7 installation.

      For UNIX or Linux, run this command:

      ```
      $DB2DIR/instance/dasmigr
      ```

      Where DB2DIR indicates the location that you specified during DB2 Version 9.7 installation.

8. Upgrade the databases.

   a. Log on to the DB2 9.7 server as the instance owner or a user with SYSADM authority.

b. Migrate the TPCDB database. Run this command from the DB2 command prompt window:

```
db2 migrate db TPCDB
```

c. From a DB2 command prompt window, run this command:

```
db2 update dbm cfg using JDK_PATH 'C:\Program Files\IBM\SQLLIB\java\jdk'
```

d. Verify that your database upgrade is successful. Connect to the upgraded databases and run a small query:

```
db2 connect to database_name

     Database Connection Information

    Database server        = DB2/INT 9.7.0
    SQL authorization ID   = ADMINISTRATOR
    Local database alias   = database_name

   db2 "select * from syscat.dbauth"
```

e. For SSPC only: Open the DB2 - DB2COPY1 - DB2-0 Properties window. Click the **Log On** tab. Check to see if the "Log On As" service for DB2 - DB2COPY1 - DB2-0 is .\administrator.

Set this DB2 service to Automatic. Open the Services window and find the DB2 - DB2COPY1 - DB2-0 service. Right-click this service and click **Automatic**.

9. Start the Tivoli Storage Productivity Center services and Agent Manager (if the Agent Manager is installed). For information about starting the services, go to the Tivoli Storage Productivity Center. Search for **Starting and stopping the Tivoli Storage Productivity Center services**.
10. Upgrade Tivoli Storage Productivity Center. For information about upgrading Tivoli Storage Productivity Center, see “Upgrading Tivoli Storage Productivity Center” on page 398.

Upgrading DB2 8.2 to DB2 9.7

Use this information to learn about the general upgrade steps required to migrate the database repository from DB2 8.2 with any fix pack to DB2 9.7.

Use this procedure only after you have upgraded from DB2 8.2 with any fix pack to DB2 9.7 and have upgraded or installed Tivoli Storage Productivity Center 4.2.2.1 or later (including Agent Manager, is installed).

The default database name for Tivoli Storage Productivity Center is **TPCDB**. The default database name for Agent Manager is **IBMCDB**. The same procedure is used for both Tivoli Storage Productivity Center and Agent Manager.

To migrate the DB2 database, follow these general steps:

1. Stop the Tivoli Storage Productivity Center services and Agent Manager (if you have Agent Manager installed).
2. Pre-check the database for migration.
3. Back up the DB2 8.1 database.
4. Migrate the DB2 instance.
5. Migrate the database.
6. Verify the migration.
7. Start the IBM Tivoli Storage Productivity Center services and Agent Manager (if you have Agent Manager installed).

Note: In this topic, DB2 8.1 is synonymous with DB2 8.2.

Migrating the database repository for Windows

This topic provides general migration steps to migrate the database repository from DB2 8.1 with any fix pack to DB2 9.7 on Windows.

Note: This procedure also applies to the Agent Manager if you have the Agent Manager installed.

This scenario assumes that you have Agent Manager installed on the same machine. If you do not have Agent Manager installed, skip the instructions for Agent Manager.

Follow these general steps for migrating the database repository:

1. Upgrade DB2 to 9.7. See the following steps.
2. Upgrade IBM Tivoli Storage Productivity Center. See Chapter 4, "Upgrading and migrating the Tivoli Storage Productivity Center family," on page 379.

To migrate the DB2 database on Windows, follow these steps:

1. Stop the IBM Tivoli Storage Productivity Center services.
2. Verify that your databases are ready for migration.
   a. Log on to the DB2 server as the DB2 instance owner that you want to migrate.
   b. Stop the instance by running the **db2stop** command:
      ```
      db2stop force
      ```
   c. From the DB2 command prompt window, go to the following directory:

      For UNIX or Linux
      Change to the $DB2DIR/bin directory where DB2DIR is the location that you specified during the DB2 Version 9.7 installation.

      For Windows
      Insert the DB2 Version 9.7 product CD into the CD-ROM drive and change to the \db2\Windows\utilities directory.
   d. Run the **db2ckupgrade** command to verify that the databases that are owned by the current instance are ready to be migrated and to generate a log file. Here is the syntax of the command:
      ```
      db2ckupgrade database -l db2ckupgrade.log -u admin_user password
      ```
      Here is an example of the command:
      ```
      db2ckupgrade TPCDB -l db2ckupgrade.log -u db2admin -p password
      ```

      If the **db2ckupgrade** command is successful, the databases can be migrated. The db2ckupgrade.log file is created in the current directory and includes information about errors and warnings.

      Each time you run this command, it overwrites the existing log file. You can rename the log file to avoid losing error information. You must correct these errors before you migrate the database.

      When the **db2migr** command runs the **db2ckupgrade** command, the log file specified is the migration.log file in the instance home directory for Linux and UNIX or in the current directory for Windows.

      Ensure that the log file for the **db2ckupgrade** command the following text:
      ```
      Version of DB2CKUPGRADE being run: VERSION 9.7.
      ```

This text confirms that you are running the correct level of the **db2ckupgrade** command.

e. Start the instance by running the **db2start** command.

3. Revoke the EXECUTE privilege on migrated routines from PUBLIC. To revoke the EXECUTE privilege on migrated routines from PUBLIC, follow these steps:

   a. Run the **db2undgp** command. In the following example, the option **-o** creates a file that contains all the REVOKE statements needed to remove the EXECUTE privilege from PUBLIC:

   ```
   db2undgp -d database -o revoke.db2
   ```

   You can review or edit this file to remove any specific statements when you want to keep the EXECUTE privilege granted to PUBLIC for any routine.

   b. Grant the EXECUTE privilege to specific users on all your routines. The following example shows how to grant this privilege on all functions under a specific schema:

   ```
   CONNECT TO TPCDB
   db2 GRANT EXECUTE ON FUNCTION  schema_name.* to USERID
   ```

   c. Run all your routines as a user granted with the EXECUTE privilege to ensure that the routines run successfully.

4. Back up your databases before migration.

   To perform an offline full backup for each of your local databases, follow these steps:

   a. Disconnect all applications and users from the database. To get a list of all database connections for the current instance, run the **list applications** command:

   ```
   db2 list applications
   ```

   If all the applications are disconnected, you see this message:

   ```
   SQL1611W No data was returned by the Database System Monitor.
        SQLSTATE=00000
   ```

   b. To disconnect all applications and users, use the **force application** command:

   ```
   db2 force application all
   ```

   c. Back up your database using the **backup database** command.

   Syntax command:

   ```
   db2 BACKUP DATABASE database USER user_ID USING password
   TO backup_directory
   ```

   This is an example of the command:

   ```
   db2 BACKUP DATABASE TPCDB USER db2admin USING password TO C:\DB2BK
   ```

   In partitioned database environments, back up all the database partitions.

   If you performed a full offline database backup recently and you cannot perform another one before migration, you can perform an incremental offline database backup instead. For information about how to perform an incremental offline database backup, see step 5 on page 389.

   d. Optional step. You can test the integrity of the backup image to ensure that the image can be restored using the **db2ckbkp** command. Here is the syntax of the command:

   ```
   db2ckbkp database
   ```

Here is an example of the command:

```
cd backup_directory
db2ckbkp SAMPLE.0.arada.NODE0000.CATN0000.20051014114322.001
```

You see this output:

```
[1] Buffers processed:  #######
Image Verification Complete - successful.
```

5. Migrate the DB2 server. To migrate the DB2 server from DB2 Version 9.1 or DB2 UDB Version 8 to DB2 Version 9.7, follow these steps:
   a. log on to the DB2 server as a user with Local Administrator authority.
   b. Install DB2 Version 9.7 by running the setup.exe command to launch the DB2 Setup wizard. You have three choices:

      Work with Existing option
      This option appears on the Install a Product panel. In the Work with Existing panel, select the DB2 copy name with the migrate action. All your instances running on the selected DB2 copy and your DAS are automatically migrated to the DB2 Version 9.7 copy. The selected DB2 copy and add-on products are uninstalled.

      You get a warning message that recommends that you run the **db2ckupgrade** command if you have local databases. If you completed the pre-migration tasks, ignore this warning message and continue with the migration. Otherwise, verify that your databases are ready for DB2 migration before continuing with the installation.

      Install New option
      This option appears on the Install a Product panel. This option creates a copy of DB2 Version 9.7, and you must migrate your instances after installation.

      Work with Existing option, create response file
      This option appears on the Install a Product panel. In the Work with Existing panel, select the DB2 copy name with the migrate action. In the Select the installation panel, response file creation, or both window, select **Save my installation setting in a response file** option to create a response file for a response file installation. The response file has the required MIGRATE_PRIOR_VERSIONS keyword, the DB2 copy name, and the installation path.

      The response file installation is the same as the first choice, all your instances running on the selected DB2 copy and your DAS are automatically migrated to the DB2 Version 9.7 copy.

   c. If you want your applications to access the DB2 Version 9.7 copy through the default interface or if you migrated your existing DB2 UDB Version 8 copy, set the DB2 Version 9.7 copy as the DB2 default copy. You must define a default copy if you migrated from DB2 UDB Version 8 because there is no default copy defined in your DB2 server.
   d. Optional step. When you select to install a new copy, migrate the DAS if you want to keep your existing DAS configuration and use new functions available with DB2 Version 9.7. If your DAS is running on DB2 UDB Version 8, you need to migrate it to use the Control Center to administer your DB2 Version 9.7 instances.
6. Migrate your databases. To migrate a DB2 database, follow these steps:
   a. Log on to the DB2 server as the instance owner or a user with SYSADM authority.

b. Optional step. Rename or delete the db2diag.log file so that a new file is created. Also, remove or move any existing dump files, trap files, and alert log files from the directory indicated by the **diagpath** parameter. Move these files to another directory. By doing this action, the files only contain information about the migration process and helps you isolate and understand any problems that might occur during database migration.

c. Migrate the database using the **migrate database** command. Here is the syntax for the command:

```
db2 MIGRATE DATABASE database_alias USER user_ID USING password
```

Where *database_alias* is the name or the alias of the database you want to migrate. *user_ID* is the user ID with SYSADM authority. Here is an example of the command:

```
db2 MIGRATE DATABASE TPCDB USER db2admin USING password
```

d. If the database migration fails and returns the error message SQL1704N with a reason code that describes the cause of the failure, find this SQL error code and determine the list of the possible solutions for each reason code. One of the most common causes of migration failure is that the log file space is not large enough, in which case the following error is returned:

```
SQL1704N  Database migration failed.  Reason code "3".
```

You must increase the log file size and run the **migrate database** command again. Once the database migration is complete, reset the value of **logfilsiz**, **logprimarylogsecond** database configuration parameters. There are additional error codes that are returned by the **migrate database** command for specific cases that are not supported by the database migration. These cases are described in the migration restrictions for DB2.

e. If the database migration returns the warning message SQL1243W, you need to drop or rename the SYSTOOLS.DB2LOOK_INFO table. Otherwise, the ALTER TABLE and COPY SCHEMA statements fail to run. Check if the SYSTOOLS.DB2LOOK_INFO table exists by running the following command:

```
db2 "SELECT tabname, tabschema, definer FROM syscat.tables
          WHERE tabschema = 'SYSTOOLS' AND tabname = 'DB2LOOK_INFO'"
```

If you created this table, simply rename it by running the RENAME statement:

```
db2 RENAME SYSTOOLS.DB2LOOK_INFO TO new_table_name
```

If you did not create this table, simply remove it by running the DROP command:

```
db2 DROP TABLE SYSTOOLS.DB2LOOK_INFO
```

7. Verify that your database migration is successful. Connect to the migrated databases and run a small query. Here is an example of the **db2 connect** command:

```
db2 connect to TPCDB
```

8. After you have migrated a database from DB2 8.1 to DB2 9.7, there are several parameters that need to be set manually.

Self-tuning memory

The self-tuning memory can be enabled for migrated databases by setting the **self_tuning_mem** configuration parameter to ON and setting these parameters:

self_tuning_mem

Run the following command in Windows and Linux for IBM Tivoli Storage Productivity Center:

```
DB2 UPDATE DB CFG FOR TPCDB USING SELF_TUNING_MEM ON
```

For Agent Manager, run this command:

```
DB2 UPDATE DB CFG FOR IBMCDB USING SELF_TUNING_MEM ON
```

pckcachesz

Run the following command in Windows and Linux for IBM Tivoli Storage Productivity Center:

```
DB2 UPDATE DB CFG FOR TPCDB USING PCKCACHESZ AUTOMATIC
```

For Agent Manager, run this command:

```
DB2 UPDATE DB CFG FOR IBMCDB USING PCKCACHESZ AUTOMATIC
```

database_memory

Run the following command for IBM Tivoli Storage Productivity Center:

```
For Windows
DB2 UPDATE DB CFG FOR TPCDB USING DATABASE_MEMORY AUTOMATIC
```

Run the following command for Agent Manager:

```
For Windows
DB2 UPDATE DB CFG FOR IBMCDB USING DATABASE_MEMORY AUTOMATIC
```

avg_appls

This parameter is used by the query optimizer to estimate how much buffer pool is available at run-time for the access plan chosen. Run this command for Windows and Linux for IBM Tivoli Storage Productivity Center:

```
DB2 UPDATE DB CFG FOR TPCDB USING AVG_APPLS AUTOMATIC
```

Run this command for Agent Manager:

```
DB2 UPDATE DB CFG FOR IBMCDB USING AVG_APPLS AUTOMATIC
```

9. Restart the IBM Tivoli Storage Productivity Center services.
10. Upgrade IBM Tivoli Storage Productivity Center to 4.1. See Chapter 4, "Upgrading and migrating the Tivoli Storage Productivity Center family," on page 379.

Migrating the database repository for AIX

This topic provides general migration steps to migrate the database repository from DB2 8.1 with any fix pack to DB2 9.7 on AIX.

When installing or upgrading Tivoli Storage Productivity Center or DB2 on AIX, you need a Telnet/SSH session and a graphical display.

Follow these general steps for migrating the database repository:

1. Upgrade DB2 to 9.7. See the following steps.
2. Upgrade Tivoli Storage Productivity Center. See Chapter 4, "Upgrading and migrating the Tivoli Storage Productivity Center family," on page 379.

To migrate the DB2 database on AIX, follow these steps:

1. Verify that Tivoli Storage Productivity Center is up and running:

a. Open the Tivoli Storage Productivity Center GUI.
b. Check to see if the following Tivoli Storage Productivity Center services are running:

```
Administrative Services > Services > Data Server
Administrative Services > Services > Device Server
```

2. Stop the following Tivoli Storage Productivity Center services:
   a. Close the Tivoli Storage Productivity Center GUI.
   b. Stop the Tivoli Storage Productivity Center Data server and Device server by running these commands:

```
Data Server:
stopsrc -s TSRMsrv1

Device Server:
/opt/IBM/TPC/device/bin/aix/stopTPCF.sh
```

   c. Stop the Data agent and Fabric agent by stopping the common agent:

```
/opt/IBM/TPC/ca/subagents/TPC/Data/tpcdagt2 stop
```

3. Make sure there are no DB2 connections to the instance used for the Tivoli Storage Productivity Center database:
   a. Log in with the instance owner ID. (You can also log in using the telnet/putty session.) The command is:

```
su - db2_8.1_instance_owner
```

      For example:

```
su - db2inst1
```

   b. Stop all DB2 connections by running this command:

```
db2 force applications all
```

   c. List the DB2 applications by running this command:

```
db2 list applications
```

      This is the output expected:

```
SQL1611W No data was returned by Database System Monitor.
SQLSTATE=00000
```

4. Make a database backup by running the following command and then exit:

```
db2 backup db database_name
exit
```

   For example:

```
db2 backup db TPCDB
exit
```

5. Install DB2 9.7 on your system.

   Note:

   - Open and use a new graphical display that is not loaded with the parameters for DB2 8.1. Otherwise, you get an error for the DB2 9.7 installation.
   - You do **NOT** need to create an instance while installing DB2 9.7. Migrate the instance used in DB2 8.1 to DB2 9.7 after installation of DB2 9.7.

6. Migrate the DB2 8.1 instance to DB2 9.7 instance:
   a. Go to the display loaded with the DB2 8.1 parameters.
   b. Verify that the instance can be migrated from DB2 8.1 to DB2 9.7:
      1) Switch to the instance owner ID:

```
su - db2_8.1_instance_owner
```

For example:

```
su - db2inst1
```

2) List the instances on DB2 8.1:

```
db2ilist -a
```

For example, the output is displayed as follows for the Tivoli Storage Productivity Center database:

```
db2inst1 32/opt/IBM/db2/V8.1
```

3) Stop the DB2 instance:

```
db2stop
```

4) Exit from the login session:

```
exit
```

c. Migrate the database from DB2 8.1 to DB2 9.7:

1) Open a new telnet/putty/SSH session with the root ID.

2) Change the directory to the DB2 9.7 installation location and instance. For example:

```
cd /opt/ibm/db2/V9.7/instance
```

3) Migrate the DB2 instance:

```
./db2imigr -d -u db2_fence_id db2_8.1_instance_name
```

For example:

```
./db2imigr -d -u db2fenc1 db2inst1
```

This command returns "exit 0."

4) Switch to the instance owner ID:

```
su db2_8.1_instance_owner
```

For example:

```
su - db2inst1
```

5) Check the DB2 level:

```
db2level
```

The output displays the version information for DB2 v9.7.0.2.

6) Migrate the Tivoli Storage Productivity Center database:

```
db2start
db2 migrate database TPC_database
```

For example:

```
db2start
db2 migrate database TPCDB
```

7) Exit from the login session.

```
exit
```

7. Configure DB2 V9.7:

a. Switch to the instance owner ID:

```
su - db2_9.7_instance_owner
```

For example:

```
su - db2inst1
```

b. Set the self-tuning memory option:

```
db2 update db cfg for TPC_database using self_tuning_mem on
```

For example:

```
db2 update db cfg for TPCDB using self_tuning_mem on
```

c. Set the package cache size option:

```
db2 update db cfg for TPC_database using pckcachesz automatic
```

For example:

```
db2 update db cfg for TPCDB using pckcachesz automatic
```

d. Set the database memory option:

```
db2 update db cfg for TPC_database using database_memory computed
```

For example:

```
db2 update db cfg for TPCDB using database_memory computed
```

e. Exit from the login session:

```
exit
```

8. Start the Tivoli Storage Productivity Center services.
9. Upgrade Tivoli Storage Productivity Center to 4.1. See Chapter 4, "Upgrading and migrating the Tivoli Storage Productivity Center family," on page 379.

Migrating the database repository for Linux

This topic provides general migration steps to migrate the database repository from DB2 8.1 with any fix pack to DB2 9.7 on Linux.

When installing or upgrading Tivoli Storage Productivity Center or DB2 on Linux, you need a Telnet/SSH session and a graphical display interface.

Follow these general steps for migrating the database repository:

1. Upgrade DB2 to 9.7. See the following steps.
2. Upgrade Tivoli Storage Productivity Center. See Chapter 4, "Upgrading and migrating the Tivoli Storage Productivity Center family," on page 379.

To migrate the DB2 database on Linux, follow these steps:

1. Verify that Tivoli Storage Productivity Center is up and running:
   a. Open the Tivoli Storage Productivity Center GUI.
   b. Check to see if the following Tivoli Storage Productivity Center services are running:

```
Administrative Services > Services > Data Server
Administrative Services > Services > Device Server
```

2. Stop the following Tivoli Storage Productivity Center services:
   a. Close the Tivoli Storage Productivity Center GUI.
   b. Stop the Tivoli Storage Productivity Center Data server and Device server:

```
Data Server:
/opt/IBM/TPC/data/server/tpcdsrv1 stop

Device Server:
/opt/IBM/TPC/device/bin/linux/stopTPCF.sh
```

   c. Stop the Data agent and Fabric agent by stopping the common agent:

```
/opt/IBM/TPC/ca/subagents/TPC/Data/tpcdagt2 stop
```

3. Make sure there are no DB2 connections to the instance used for the Tivoli Storage Productivity Center database:
   a. Log in with the instance owner ID. (You can also log in using the telnet/putty session.) The command is:
      ```
      su - db2_8.1_instance_owner
      ```
      For example:
      ```
      su - db2inst1
      ```
   b. Stop all DB2 connections by running this command:
      ```
      db2 force applications all
      ```
   c. List the DB2 applications by running this command:
      ```
      db2 list applications
      ```
      This is the expected output:
      ```
      SQL1611W No data was returned by Database System Monitor.
      SQLSTATE=00000
      ```
4. Make a database backup by running the following command and then exit:
   ```
   db2 backup db database_name
   exit
   ```
   For example:
   ```
   db2 backup db TPCDB
   exit
   ```
5. Install DB2 9.7 on your system.

   Note: You do **NOT** need to create an instance while installing DB2 9.7. Migrate the instance used in DB2 8.1 to DB2 9.7 after installation of DB2 9.7.
6. Migrate the DB2 8.1 instance to DB2 9.7 instance:
   a. Verify that the instance can be migrated from DB2 8.1 to DB2 9.7:
      1) Switch to the instance owner ID:
         ```
         su - db2_8.1_instance_owner
         ```
         For example:
         ```
         su - db2inst1
         ```
      2) List the instances on DB2 8.1:
         ```
         db2ilist -a
         ```
         For example, the output displays as follows for the Tivoli Storage Productivity Center database:
         ```
         db2inst1 32/opt/IBM/db2/V8.1
         ```
      3) Stop the DB2 instance:
         ```
         db2stop
         ```
      4) Exit from the login session:
         ```
         exit
         ```
   b. Migrate the database from DB2 8.1 to DB2 9.7:
      1) Open a new telnet/putty/ssh session with the root ID.
      2) Change the directory to the DB2 9.7 installation location and instance. For example:
         ```
         cd /opt/ibm/db2/V9.7/instance
         ```
      3) Migrate the DB2 instance:

```
./db2imigr -d -u db2_fence_id db2_8.1_instance_name
```

For example:

```
./db2imigr -d -u db2fenc1 db2inst1
```

This command returns "exit 0."

4) Switch to the instance owner ID:

```
su - db2_8.1_instance_owner
```

For example:

```
su - db2inst1
```

5) Check the DB2 level:

```
db2level
```

The output displays the version information for DB2 v9.7.0.2.

6) Migrate the Tivoli Storage Productivity Center database:

```
db2start
db2 migrate database TPC_database
```

For example:

```
db2start
db2 migrate database TPCDB
```

7) Exit from the login session.

```
exit
```

7. Configure DB2 V9.7:
   a. Switch to the instance owner ID:

```
su - db2_9.7_instance_owner
```

For example:

```
su - db2inst1
```

   b. Set the self-tuning memory option:

```
db2 update db cfg for TPC_database using self_tuning_mem on
```

For example:

```
db2 update db cfg for TPCDB using self_tuning_mem on
```

   c. Set the package cache size option:

```
db2 update db cfg for TPC_database using pckcachesz automatic
```

For example:

```
db2 update db cfg for TPCDB using pckcachesz automatic
```

   d. Set the database memory option:

```
db2 update db cfg for TPC_database using database_memory computed
```

For example:

```
db2 update db cfg for TPCDB using database_memory computed
```

   e. Exit from the login session:

```
exit
```

8. Start the Tivoli Storage Productivity Center services.
9. Upgrade Tivoli Storage Productivity Center to 4.1. See Chapter 4, "Upgrading and migrating the Tivoli Storage Productivity Center family," on page 379.

Preparing for an upgrade

When you upgrade Tivoli Storage Productivity Center, you are upgrading all components (if all components are installed). This topic describes the preparation steps required before you do an upgrade installation.

If you have your database on a separate machine from the server, you must also upgrade the remote database.

Before you upgrade Tivoli Storage Productivity Center, back up your environment.

Backing up IBM Tivoli Storage Productivity Center

This section describes how to back up your IBM Tivoli Storage Productivity Center system before performing an upgrade installation.

It is important to back up your Tivoli Storage Productivity Center environment in the event of a problem or disaster. Backups should begin as soon as possible after your Tivoli Storage Productivity Center server is installed.

The frequency of back ups depends on your company's data protection policies, Tivoli Storage Productivity Center work load, and size of the Tivoli Storage Productivity Center environment.

A comprehensive backup plan for your Tivoli Storage Productivity Center environment should provide coverage of all of the following items:

- Installation documentation including user IDs and passwords for DB2, Agent Manager (if installed), Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication.
- Full server backup if possible (so that you can recover the entire server if necessary).
- DB2 backup of the Tivoli Storage Productivity Center databases (TPCDB and IBMCDB). IBMCDB is the Agent Manager database and TPCDB is the Tivoli Storage Productivity Center database.
- For Tivoli Storage Productivity Center and Tivoli Integrated Portal single signon authentication configuration, back up the WebSphere configuration files. The configuration files are located in the following directories:

  TIP_installation_directory/profiles/TIPProfile/bin
  TPC_installation_directory/device/apps/was/profiles/deviceServer/bin

  The backup file is named:

  WebSphereConfig_yyyy_mm_dd.zip

  Where *yyyy* is the year, *mm* is the month, and *dd* is the day.

  Run the following commands on UNIX or Linux systems::

  /IBM/Tivoli/tip/profiles/TIPProfile/bin/backupConfig.sh
  /IBM/TPC/device/apps/was/profiles/deviceServer/bin/backupConfig.sh

  Run the following commands on Windows systems:

  \IBM\Tivoli\tip\profiles\TIPProfile\bin\backupConfig.bat
  \IBM\TPC\device\apps\was\profiles\deviceServer\bin\backupConfig.bat
- Backup the following registries:

  InstallShield registries
  Back up the following registries:

```
AIX: /usr/lib/objrepos/InstallShield/Universal/IBM-TPC/
UNIX: /root/InstallShield/Universal/IBM-TPC
Windows: C:\Program Files\Common Files\InstallShield\Universal\IBM-TPC
```

SRM legacy registry
Back up the following registries:

```
AIX: subsystem TSRMsrv# where # can be any number
UNIX: /etc/Tivoli/TSRM
```

Windows registry
Back up the Windows registry.

Common agent registry (if you have Data agents and Fabric agents installed)
Back up the following registries:

```
AIX or UNIX: /usr/tivoli/ep*, /opt/tivoli/ep*
Windoes: C:\Program Files\Tivoli\ep*
```

- Back up the Tivoli GUID setting. Go to `C:\Program Files\Tivoli\guid` (for Windows) or `/opt/tivoli/guid` (for AIX or UNIX). Run the following command:

  ```
  tivguid -show >tpc_tivguid.txt
  ```

- Back up the Agent Manager files and directories if you have Agent Manager installed.

  AM_installation_directory`/AppServer/agentmanager/config/cells/AgentManagerCell/security.xml`
  AM_installation_directory`/AppServer/agentmanager/installedApps/AgentManager.ear/AgentManager.war/WEB-INF/classes/resources/AgentManager.properties`
  AM_installation_directory`/os.guid`
  AM_installation_directory`/certs`

- Back up Tivoli Storage Productivity Center server files and directories.

  TPC_installation_directory`/config`
  TPC_installation_directory`/data/config`
  TPC_installation_directory`/device/config`

- Back up the Data agent and Fabric agent files and directories (if you have the Data agent and Fabric agent installed).

  TPC_installation_directory`/config`
  TPC_installation_directory`/ca/cert`
  TPC_installation_directory`/ca/config`
  TPC_installation_directory`/ca/*.sys`
  TPC_installation_directory`/ca/subagents/TPC/Data/config`
  TPC_installation_directory`/ca/subagents/TPC/Fabric/config`

- Back up any interim fixes or work around code provided by Tivoli Storage Productivity Center support.

Upgrading Tivoli Storage Productivity Center

This topic describes how to upgrade all the Tivoli Storage Productivity Center components.

To upgrade Tivoli Storage Productivity Center, use the same starting procedure as the installation program. Depending on what components you have installed on your system, there are some differences in the panels you see. If you are upgrading the agents on a different server, see “Upgrading Storage Resource agents” on page 406.

There are two installation images:

disk 1 Disk 1 has four parts:

- Disk 1, Part 1 -- Base IBM Tivoli Storage Productivity Center installation program
- Disk 1, Part 2 -- Base IBM Tivoli Integrated Portal and IBM Tivoli Storage Productivity Center for Replication installation programs
- Disk 1, Part 3 -- IBM Tivoli Integrated Portal PTF installation program
- Disk 1, Part 4 -- IBM Tivoli Storage Productivity Center Monitoring Agent installation program for IBM Tivoli Monitoring. This is an optional program and only required if you are using Tivoli Monitoring.

The **disk 1** images part 1, part 2, part 3, and part 4 must be downloaded to the same directory.

Storage Resource Agent

Contains the files to perform local Storage Resource agent installations or upgrades. Download the file for the platform you want the agent to reside on.

The **Storage Resource Agent** location and operating system file name is the same as the **disk 1** image.

For Storage Resource agents, the image is located in the following location:

DVD/data/sra/*operating_system_name*

The operating systems supported are listed in the following table.

Table 44. Operating system for Storage Resource agents

Operating system	Operating system name
AIX	aix_power
HP-UX	hp-ux_itanium
Linux x86	linux_ix86
Linux for Power Systems Servers	linux_power
Linux s390	linux_s390
Solaris	solaris_spark
Windows	windows

Note:

- If you install the Storage Resource agent using the disk 1 image, you must uninstall the agent using the uninstallation wizard (InstallShield). If you uninstall the agent using the Tivoli Storage Productivity Center user interface, InstallShield does not recognize the uninstallation. Therefore, you must uninstall the agent that is installed from the disk 1 image using InstallShield.
- If you are deploying Storage Resource agent on UNIX or Linux systems, you must use root as the user ID.
- You cannot install the Storage Resource agent from the disk 1 image on machines that do not have a Data server installed. If you are installing only a Storage Resource agent, use the Tivoli Storage Productivity Center user interface. You can install the Storage Resource agent if you are also installing the Data server.
- You can upgrade TotalStorage Productivity Center 3.3.x to Tivoli Storage Productivity Center version 4.2.1 or later in an IPv4 environment to an environment that supports both IPv6 and IPv4. You cannot upgrade the product from an IPv4-only environment directly to an IPv6-only environment. for more information, see “Planning for Internet Protocol Version 6” on page 51.

- After you upgrade Tivoli Storage Productivity Center, all trace settings are reset to the default values. If you changed the trace settings, you must reset the trace settings after an upgrade.

Important: This procedure is disruptive to the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication services. For Tivoli Storage Productivity Center for Replication, you can perform a takeover to a standby server to ensure that consistency can be maintained for any active servers.

If you have installed all the TotalStorage Productivity Center components on your system and are upgrading Tivoli Storage Productivity Center, complete the following steps:

1. Back up your Tivoli Storage Productivity Center environment. See "Backing up IBM Tivoli Storage Productivity Center" on page 397.
2. Stop all Tivoli Storage Productivity Center jobs.
   a. Stop all jobs including performance monitor jobs, system and fabric probe jobs, scan jobs, and other probe jobs.
   b. If possible, reboot the Tivoli Storage Productivity Center servers. This action will stop any remaining Tivoli Storage Productivity Center Java processes that might not stop in a timely manner. If a reboot is not possible, stop and restart the Tivoli Storage Productivity Center server service. It is important for the Tivoli Storage Productivity Center Device server to stop and restart cleanly. If this does not occur, a server reboot might be indicated.
3. After the Tivoli Storage Productivity Center server and services are started, start the Tivoli Storage Productivity Center installation program and apply the upgrade.

   Note: On AIX or Linux, before running the installation program, source the **db2profile** file. For example, if the DB2 instance is **db2inst1**, source the file by typing:

   ```
   . /home/db2inst1/sqllib/db2profile
   ```

   Please note that there is a space between **.** and **/home**.
4. The Select a language panel is displayed. Select a language and click **OK**.
5. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select I accept the terms of the license agreement. Click **Next**.

   Note: 8 GB of RAM is required. If you have 4-8 GB of RAM, you can still install Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication. However, you get a warning message. Run only Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication to avoid system overload. For information about how to disable Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication after installation, see "Disabling unlicensed or unused Tivoli Storage Productivity Center components" on page 368.
6. The Select the type of installation you want to run panel is displayed. Select **Typical installation**. Select **Servers**, **Storage Resource Agent** (if you want the agent installed), and **Clients**. Click **Next**.

   Note: If you have a Data agent or Fabric agent on the server system and select **Storage Resource Agent**, the Data agent or Fabric agent gets migrated to a Storage Resource agent on the server system.

7. The User ID and password, and server and agent information panel is displayed. Enter the User ID and password. Click **Next**.

   If you are upgrading Tivoli Storage Productivity Center from 3.3.x to 4.2.1 or later, go to step 9.

   If you are upgrading Tivoli Storage Productivity Center from 4.1.0 or 4.1.1 to 4.2.1 or later, go to step 15 on page 404.

   Note:

   - If you have advanced Brocade discovery configured but do not have a Brocade SMI-S agent for the fabric, then you receive a warning message. The warning message indicates that advanced Brocade discovery is removed from the product. You can install and configure a Brocade SMI-S agent to manage the fabric before upgrading Tivoli Storage Productivity Center. Otherwise, your zoning information for the fabric is removed during the upgrade.
   - If you are using custom installation, you see the Brocade warning message when you see the Existing database schema information panel.

8. If Tivoli Storage Productivity Center detects that you have a DS8000, XIV, or SAN Volume Controller storage system, the "Storage Subsystem Credential Migration Tool" panel is displayed. For information about the Storage Migration Tool, see “Update Storage Subsystem Credentials” on page 328. If you want to run the migration tool at upgrade time, check the box for **Run Storage Subsystem Credential Migration Tool**. Otherwise, clear the box for this option. Click **Next**.

   Note: If you do not run the migration tool at upgrade time, when you first log into the Tivoli Storage Productivity Center GUI, the Welcome panel will display a button for **Update Subsystems**. You can update the storage subsystems credentials at that time or at a later time.

   If you update the credentials at a later time, the Tivoli Storage Productivity Center GUI displays a node for **Update Storage Subsystem Credentials**. To update the credentials from the GUI, expand **Administrative Services** > **Configuration** > **Update Storage Subsystem Credentials**. After you update the credentials, this node is no longer displayed in the GUI.

9. The Tivoli Integrated Portal panel is displayed. Enter this information and click **Next**.

   Specify the location to install TIP
   Accept the default or enter a location to install Tivoli Integrated Portal.

   Port Enter the port number for Tivoli Integrated Portal or accept the default (16310). Tivoli Integrated Portal uses 10 port numbers. For example, if the default port 16310 is accepted, Tivoli Integrated Portal uses the following ports: 16310, 16311, 16312, 16313, 16315, 16316, 16318, 16320, 16322, and 16323. If you enter a port number other than the default number, ensure that you have the following ports available for Tivoli Integrated Portal:

   - *base port*
   - *base port* +1
   - *base port* +2
   - *base port* +3
   - *base port* +5

- *base port* +6
- *base port* +8
- *base port* +10
- *base port* +12
- *base port* +13

where *base port* is the port value entered in the **Port** field.

Reuse an existing install TIP install
If you select this option, you see a list of existing IBM Tivoli Integrated Portal installations. Select a IBM Tivoli Integrated Portal installation.

TIP Administrator ID
Enter the IBM Tivoli Integrated Portal administrator ID.

Password
Enter the password for the IBM Tivoli Integrated Portal administrator ID.

10. The authentication panel is displayed. Enter this information and click **Next**.

OS Authentication
Uses the operating system for user authentication. If you select this option, go to step 15 on page 404.

LDAP/Active Directory
If you select LDAP or Microsoft Active Directory for authentication, you must have an LDAP or Active Directory installed.

If you select this option, you also see other panels for LDAP authentication. Click **Next**.

11. The Lightweight Directory Access Protocol (LDAP server information) panel is displayed. Enter this information and click **Next**.

LDAP Server Hostname
The fully qualified host and domain name of the machine where your LDAP-based directory is running.

LDAP Port Number
The port number on your LDAP server where the LDAP process is listening for communications.

Support referrals to other LDAP servers
The options are **ignore** or **follow** an LDAP referral. The default is **ignore**.

Referral is the process by which an LDAP server, instead of returning a result to a client application, returns a reference (a referral) to another LDAP server which might contain further information related to the client application's request.

Bind Distinguished Name

The distinguished name used to bind to the LDAP-based directory when performing a user or group search. This name is optional because some LDAP-based directories allow anonymous binds and others require you to bind with a specific name and password.

Note: Depending on how the LDAP server is set up, the Bind DN and password might or might not be optional. If the LDAP server allows for anonymous binds, then the Bind DN and password are optional.

If you want the ability to create or modify LDAP users and groups from the IBM Tivoli Integrated Portal administrative panel, then the Bind DN and password are required.

Bind Password
The password associated with the Bind Distinguished Name.

12. The Lightweight Directory Access Protocol (LDAP user and group information) panel is displayed. Enter this information and click **Next**.

Relative Distinguished Name for usernames
Tells IBM Tivoli Storage Productivity Center where to start a search in the directory when performing user authentication.

Attribute to use for usernames
Tells IBM Tivoli Storage Productivity Center which attribute in a user directory entry contains the user name for authentication.

Relative Distinguished Name for groups
Tells IBM Tivoli Storage Productivity Center where in the directory to start a search when performing a group search for authorization.

Attribute to use for groups
Tells IBM Tivoli Storage Productivity Center which attribute in a directory entry that contains the name for the group for authorization.

13. The Lightweight Directory Access Protocol (LDAP TPC Administrator user and group) panel is displayed. Enter this information and click **Next**.

LDAP TPC Administrator username
The user ID that IBM Tivoli Storage Productivity Center configures as the IBM Tivoli Storage Productivity Center superuser during installation. This user name must exist in the directory before you start the IBM Tivoli Storage Productivity Center installation and must be a member of your designated LDAP IBM Tivoli Storage Productivity Center Administrator group in the directory.

Note: If you install IBM Tivoli Storage Productivity Center on a Windows system and have IBM Tivoli Storage Productivity Center authenticate users against an LDAP-based repository, then the LDAP TPC Administrator user name must **not** contain a space or blank character in it. See the WebSphere Application Server APAR PK77578.

Ensure that you use the correct case when entering the values for the LDAP TPC Administrator user name and LDAP TPC Administrator group. If you use the incorrect case when entering these values, then you cannot log in to Tivoli Storage Productivity Center for Replication with the user name you expected to use.

LDAP TPC Administrator password
Password for the administrator.

LDAP TPC Administrator group
The group that Tivoli Storage Productivity Center maps to the Tivoli Storage Productivity Center superuser role during installation. This group must exist in the directory before you start the Tivoli Storage Productivity Center installation and must have the designated LDAP IBM Tivoli Storage Productivity Center Administrator as a member.

Note:

Ensure that you use the correct case when entering the values for the LDAP TPC Administrator user name and LDAP TPC Administrator group. If you use the incorrect case when entering these values, then you cannot log in to Tivoli Storage Productivity Center for Replication with the user name you expected to use.

14. The IBM Tivoli Storage Productivity Center role to group mapping panel is displayed. For information about authorizing users, see "Role-to-Group Mappings" on page 291.
15. The Summary information panel is displayed. Review the components to be installed. Click **Install**.
16. If you have the Monitoring Agent for Windows OS installed on the system, you might see a message to shut down this service. To stop this service on Windows, go to the Windows Services window and click **stop** for the Monitoring Agent for Windows OS -Primary service.
17. You see the installing panel. When Tivoli Integrated Portal has finished installing, Tivoli Storage Productivity Center starts the Tivoli Storage Productivity Center for Replication installation.

    When you upgrade the server, the server is restarted and therefore, any High Availability connection between the active and standby servers is disconnected. You must resynchronize the active and standby servers manually. If the standby server is upgraded, the server becomes an active server.

    Note: When you are upgrading the system, you might see several windows prompting you with the text **Replace Existing File**. Reply **Yes to All** to these prompts. Sometimes this dialog window is hidden behind the main installation panel. Make sure you check behind the main installation panel to see if there are any hidden dialog panels.
18. If you elected to upgrade the agents on this server, you do not have to upgrade the agents separately. The upgrade procedure handles upgrading the agents on this server. If you need to upgrade the agents on different computers, see "Upgrading Storage Resource agents" on page 406.
19. The successfully installed panel is displayed. Click **Finish**.

    Note: If you see the panel indicating that the product has been 100% installed and receive the following message:

    ```
    /opt/IBM/TPC/service/service.sh exist on this system and is newer
    than the file being installed. Do you want to replace this file?
    ```

    Click **Yes to All**.
20. If you are upgrading from a previous release (for example, Tivoli Storage Productivity Center 3.3.2 to 4.2.1 or later or installing a fix pack), and used the Language Pack in the previous release, you need to install the Tivoli Storage Productivity Center 4.2.1 or later Language Pack.

    Note: For Windows, if you want the log files in a language other than English (if English is set on Windows), you must change the setting on Windows to another language.
21. If you are upgrading your system on Windows, you must reboot the system.
22. If upgrading from Tivoli Storage Productivity Center version 3.x, migrate the Tivoli Storage Productivity Center database. This step can be performed right after upgrading or at a later time. The performance database migration tool can run a long time depending on the size of your database. It is a good

practice to run the performance database migration tool before you install any Tivoli Storage Productivity Center fix packs. To run the database migration tool, follow these steps:

a. Back up your Tivoli Storage Productivity Center database using the DB2 backup command.
b. Stop the Data server and Device server.
c. Go to the following directory on the installation media:

   TPC_install_media_disk1\data\scripts

d. Run the following command:

   ```
   partitiontables.bat (for Windows)
   partitiontables.sh (for UNIX and AIX)
   ```

The database connection information is automatically obtained and the `migrateTable.log` file is created under the following directory:

TPC_install_directory\data\server\tools

If the database migration tool is interrupted for some reason, it is safe to run the migration tool again. The tool picks up where it left off the last time. If the database migration has been completed and the command is run again, this tool returns immediately.

Upgrading the remote database

If the Tivoli Storage Productivity Center database is on a remote system from the server, you must also upgrade the remote database. You must upgrade the remote database before you upgrade the Tivoli Storage Productivity Center server.

Note: If you are upgrading the database schema from TotalStorage Productivity Center 3.3.x to Tivoli Storage Productivity Center 4.2.1 or later on Windows Server 2003, there is an issue with the maximum size of environment variables like the PATH variable. The maximum size of the PATH variable is 2048. However, in some cases, the PATH variable is truncated to 1024 characters. You need to install a hot fix from Microsoft. For information about the hot fix, go to http://support.microsoft.com/kb/906469.

To upgrade the remote database, follow these steps:

1. Start the Tivoli Storage Productivity Center installation program.
2. The Select a language panel is displayed. Select a language and click **OK**.
3. The Software License Agreement panel is displayed. Read the terms of the license agreement. If you agree with the terms of the license agreement, select I accept the terms of the license agreement. Click **Next**.
4. The Select the type of installation you want to run panel is displayed. Select **Custom installation**. Click **Next**.
5. The Select one or more components to install on the local or remote computer panel is displayed. Select **Database schema**. Click **Next**.
6. Wait for installation to complete.
7. If you have upgraded the database from DB2 V8 to DB2 V9.5, you must migrate the database after upgrading it. See “Upgrading DB2 8.2 to DB2 9.7” on page 386.
8. Migrate the Tivoli Storage Productivity Center database. This step can be performed right after upgrading or at a later time. The database migration tool can run a long time depending on the size of your database. You must,

however, run the database migration tool before you install any Tivoli Storage Productivity Center patches or PTF. To run the database migration tool, go to the following directory:

disk1_image_directory\data\scripts

Run this command:

```
partitiontables.bat (for Windows)
partitiontables.sh (for UNIX and AIX)
```

The database connection information is automatically obtained and the `migrateTable.log` file is created under the following directory:

TPC_install_directory\data\server\tools

Upgrading Storage Resource agents

This section provides general steps for upgrading an existing or older version of Storage Resource agents to the current release.

If you are upgrading the Storage Resource agents, you can upgrade them using one of the following methods:

- Use the Tivoli Storage Productivity Center user interface. For information, go to the information center. Search for **Upgrading the Storage Resource agent**.
- Use the Tivoli Storage Productivity Center user interface by selecting **Enable Auto Upgrade** for a Storage Resource agent. When you run a probe for the Storage Resource agent, the agent is automatically upgraded. To enable automatic upgrade, click **Administrative Services** > **Data Sources** > **Data/Storage Resource Agents**. In the content pane, select the agent you want and from the Select Action menu, click **Enable Auto Upgrade**.
- Use a Storage Resource agent command. See "Upgrading Storage Resource agents using a command."
- Use the installation wizard when upgrading the server. You can specify to upgrade the Storage Resource agent at the same time.

Upgrading Storage Resource agents using a command

This section describes how to upgrade the Storage Resource agents manually.

To upgrade the Storage Resource agent agent, follow these steps:

1. Go to the DVD location of the installation program (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 Storage Resource agent installation image location.
2. Run the upgrade command:

   ```
   ►►—bin/Agent— -upgrade— -installLoc—agent_install_directory————————►
                        (1)
   ►— -commType Daemon———————————————————————————————————————————————►◄
   ```

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 using the default location:

```
bin/Agent -upgrade
-installLoc "/opt/IBM/TPC/"
-commType Daemon
```

Here is an example for a non-daemon service using a non-default location:

```
bin/Agent
-upgrade -installLoc "C:\Program Files\IBM\TPC_SRA\"
```

Note:

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 Information Center. Search for **Return codes used by Storage Resource agent**.

Migrating Data agents and Fabric agents to Storage Resource agents

The Storage Resource agents now perform the functions of the Data agents and Fabric agents. You can migrate the Data agents and Fabric agents to Storage Resource agents by using the Tivoli Storage Productivity Center user interface, upgrade installation wizard (for server system only), or using a command.

For information about migrating the Data agent or Fabric agent to a Storage Resource agent during upgrade installation, see “Upgrading Tivoli Storage Productivity Center” on page 398.

Migrating the Data agent or Fabric agent using the user interface

This topic describes how to migrate the Data agent or Fabric agent to a Storage Resource agent using the Tivoli Storage Productivity Center user interface.

To schedule a migration of the Data agent or Fabric agent to a Storage Resource agent through the user interface, complete the following steps:

1. Open the Tivoli Storage Productivity Center user interface.
2. In the navigation tree, expand **Administrative Services** > **Configuration** > **Migrate Data/Fabric Agent**.
3. Right-click the **Migrate Data/Fabric Agent** node.
4. Click **Create Data/Fabric Agent Migration**.
5. Specify the information for the migration job in the appropriate fields.
6. Click **File** > **Save** in the menu bar to save the migration definition. The agents you specify are migrated at the scheduled time.

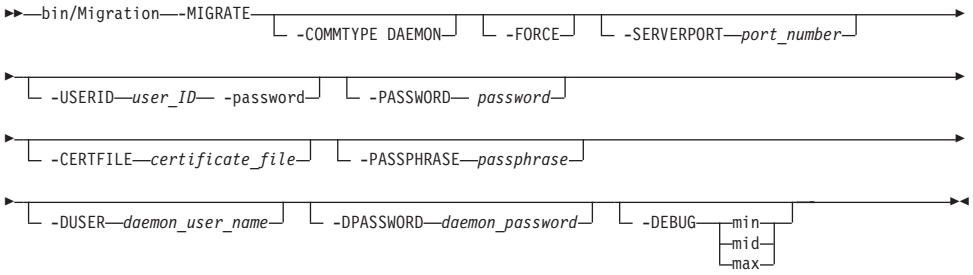
Migrating Data agents and Fabric agents using a command

This section describes how to migrate the Data agents and Fabric agents to Storage Resource agents using a command.

To migrate the Data agent or Fabric agent to the Storage Resource agent, follow these steps:

1. If the Data agent or Fabric agent is on a Virtual I/O system, you must follow this log in procedure:
   a. Log in with the **padmin** user ID.
   b. Run the following command to set up the AIX environment:
      `eom_setup_env`
2. Go to the DVD location of the installation program and go to the `bin` directory. The directory for disk 1:
   DVD_image_location`/data/sra/`*operating_system*`/bin`

   Where *DVD_image_location* is the Storage Resource agent installation image location and *operating_system* is the operating system type of the agent system.
3. Run the **Migration** command:



The parameters are:

-COMMTYPE DAEMON
Specify this parameter if you want the Storage Resource agent to run as a daemon service.

-FORCE
This forces the migration to continue even if the Storage Resource agent cannot perform all the functions of the Data agent. The **-FORCE** option applies to the following conditions:

- The Data agent is monitoring an unsupported RDBMS type.
- The Data agent is monitoring NetWare.
- The Data agent is responsible for a file system extension job.

-SERVERPORT *port_number*
Specify the port number of the Data server. This value is only required when the Data agent is not installed on the local system. The default port for the Data server is 9549.

-USERID *user_name*
For non-daemon service. The user ID defined on the agent system. This user ID is used by the server to connect to the agent system.

-PASSWORD *password*
For non-daemon service. Password for the user ID.

-CERTFILE *certificate_file*
For non-daemon service. The certificate used for SSH communication between the server and agent. This certificate needs to be stored on the server system.

-PASSPHRASE *passphrase*
For non-daemon service. The passphrase defined for the certificate used in SSH communication.

-DUSER *daemon_user_name* **(for Windows only)**
For daemon service. Login account used to start the Storage Resource agent service on Windows.

-DPASSWORD *daemon_password* **(for Windows only)**
For daemon service. Password for the Windows login account specified by the **-duser** parameter.

-DEBUG [min | mid | max]
Optional parameter for debugging purposes. You can specify the following:

min The minimum amount of debugging information.

mid A moderate amount of debugging information.

max The maximum amount of debugging information.

Here is an example for a daemon-based service:

```
bin/Migration -MIGRATE
-COMMTYPE DAEMON
-SERVERPORT 9549
-DEBUG max
```

Here is an example for a non-daemon service:

```
bin/Migration -MIGRATE
-SERVERPORT 9549
-USERID myuserid
-CERTFILE mycertfile
-PASSPHRASE mypassphrase
-DEBUG max
```

Note:

If you run the migration program outside of the *DVD_image_location* installation directory, then you must specify the full path.

Migrating Data agents and Fabric agents using the installation wizard

This section describes how to migrate the Data agents and Fabric agents to Storage Resource agents using the installation wizard.

You can migrate the Data agent or Fabric agent to the Storage Resource agent when you upgrade Tivoli Storage Productivity Center. For information about migrating the Data agent or Fabric agent, see "Upgrading Tivoli Storage Productivity Center" on page 398.

Upgrading CIM agents for storage systems

Before you upgrade a Common Information Model (CIM) agent, make sure that IBM Tivoli Storage Productivity Center supports the CIM agent that you want and that the agent is compatible with the firmware versions of your storage systems.

To ensure that a CIM agent is supported, 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. If you are uncertain about the CIM agent support, contact IBM customer support to help you with your upgrade plans.

Upgrading CIM agents

To upgrade a Common Information Model (CIM) agent, you must have the upgrade instructions that are supplied by the provider of the agent.

To upgrade the CIM agent, follow these steps:

1. Stop all Tivoli Storage Productivity Center server activity that might be dependent on the CIM agent (for example, performance monitor jobs, and discovery or probe jobs).
2. Stop the Tivoli Storage Productivity Center server services.
3. If any Tivoli Storage Productivity Center server processes remain, shut down the Tivoli Storage Productivity Center server for the duration of the CIM agent upgrade. If it is not possible to shut down the Tivoli Storage Productivity Center server, stop or kill the running processes.
4. Upgrade the CIM agent.
   a. Refer to the instructions supplied by the CIM agent to perform the upgrade.
   b. Verify that the CIM agent configuration is intact after the upgrade. Make sure that user accounts and passwords used with Tivoli Storage Productivity Center are still in place, and that all devices managed by the CIM agent are still listed.
5. Restart the Tivoli Storage Productivity Center server and Tivoli Storage Productivity Center server services.
6. Open the Tivoli Storage Productivity Center graphical user interface (GUI).
7. Run a CIM object manager (CIMOM) discovery job.
8. Run a probe job for each storage system that is managed by the CIM agent.

Upgrading System Storage N Series Gateway servers

Follow this procedure to upgrade Tivoli Storage Productivity Center before adding N Series Gateway servers by using the **Other NAS** node.

Deleting and unlicensing existing N Series Gateway servers

If you are upgrading Tivoli Storage Productivity Center from TotalStorage Productivity Center 3.3.x, there are specific steps to follow. You must first delete and unlicense existing N Series Gateway servers before adding the N Series Gateway servers.

Follow these steps:

1. Stop all Tivoli Storage Productivity Center services.
2. Make a backup copy of the Tivoli Storage Productivity Center database.
3. Open the Tivoli Storage Productivity Center GUI.
4. Manually delete all N Series Gateway servers.
   a. Expand **Administrative Services > Configuration > Manual NAS/Netware Server Entry**.

b. Select a NAS gateway server (for example, ratbert1) from the Manual NAS/Netware Server table.
c. Click **Delete**.
d. Click **Yes** to **Delete Manual NAS/Netware Server Entry** prompt.
e. Repeat steps b through d for other manually added NAS gateway servers.

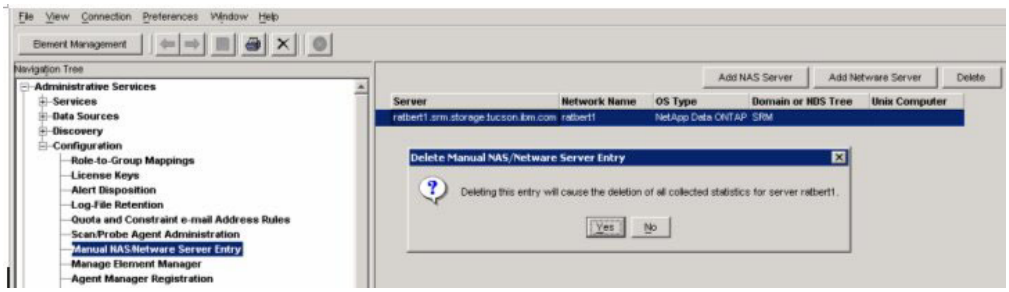


Figure 45. Delete Manual NAS/Netware Server Entry prompt

5. Unlicense all NAS gateway servers that were automatically discovered
   a. Click **Administrative Services > Configuration > License Keys**.
   b. Expand **IBM TPC for Data** under the Licensing tab.
   c. Clear the Licensed check box for all automatically discovered NAS gateway servers (for example, zinc).
   d. Click the disk icon to save the changes.
   e. Click **Yes** to the Save Confirmation prompt.

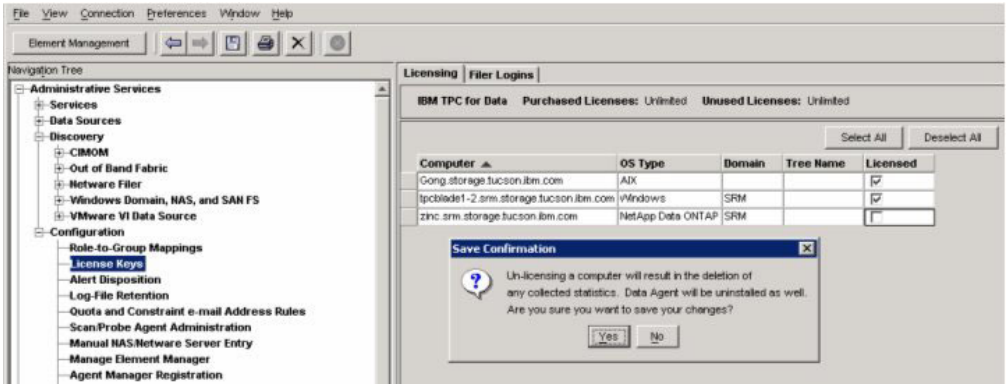


Figure 46. Save Confirmation prompt

Note: The OS Type and Licensed fields of existing NAS gateway servers are updated after it is manually added as **Other NAS**. See "Manually adding an N Series Gateway server " on page 412.

6. Exit the Tivoli Storage Productivity Center GUI.
7. Make another backup of the Tivoli Storage Productivity Center database after the N Series Gateway servers are deleted and unlicensed (make sure all Tivoli Storage Productivity Center services are stopped).
8. Restart all the Tivoli Storage Productivity Center services.
9. Go to "Manually adding an N Series Gateway server " on page 412.

Manually adding an N Series Gateway server

1. Start the Tivoli Storage Productivity Center GUI.
2. Manually add the N Series Gateway server as **Other NAS**.
   a. Click **Administrative Services > Configuration > Manual NAS/Netware Server Entry**.
   b. Click **Add NAS Server**.
   c. Enter information for the Network Name, Accessible from, SNMP Community, Login ID, and Password fields for an N Series Gateway server (for example, **ratbert1**).
   d. Check the **Add as Other NAS** check box.
   e. Enter the NAS Server Vendor Name or accept the Network Appliance as the default.
   f. Click **OK** to save the input. This action enables IBM Tivoli Storage Productivity Center to verify the NAS gateway server information and discover CIFS or NFS shares accessible to the selected Data agent or Storage Resource agent.

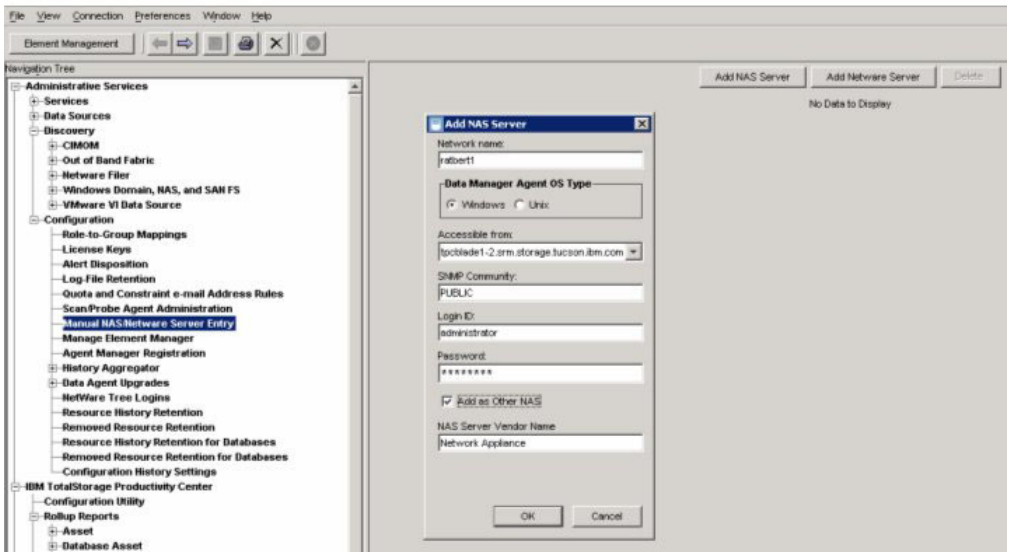


Figure 47. Add NAS Server panel

   g. Verify that the N Series Gateway server is added and listed with **Other NAS** as OS Type.

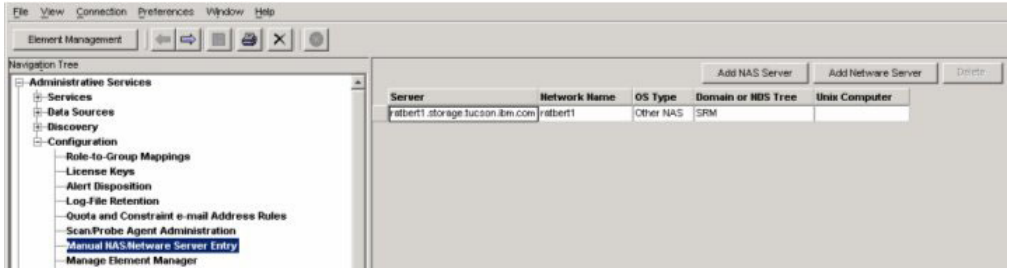


Figure 48. Verify that the N Series Gateway server is added

   h. Repeat steps b to g for other N Series Gateway servers.

3. Assign a Scan/Probe agent or agents to the exported CIFS or NFS shares accessible to the selected Data agent or Storage Resource agent.
   a. Click **Administrative Services > Configuration > Scan/Probe Agent Administration**.
   b. Select one or more file systems (for example, /vol/tpc1, /vol/tpc2, and /vol/vol0 for ratbert1) and assign the Data agent or Storage Resource agent by clicking **Set agent per row** or **Set agent** for all the selected rows.

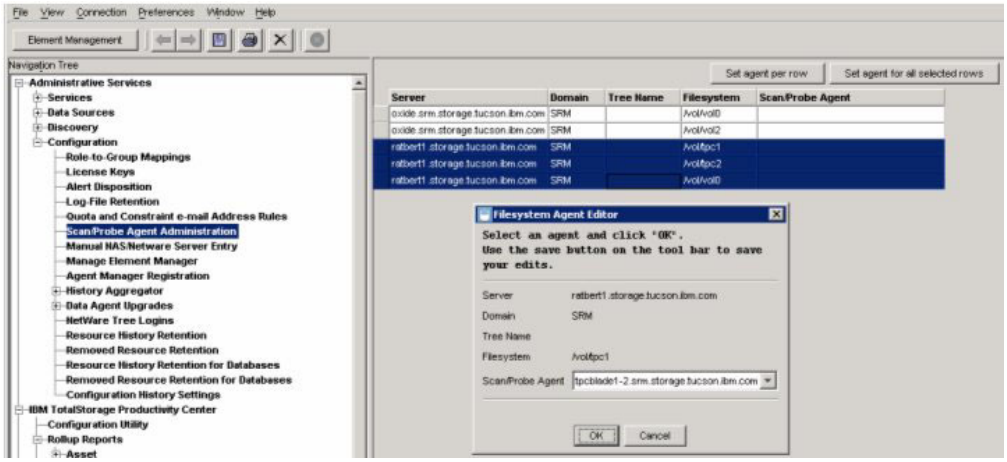


Figure 49. File system Agent Editor panel

   c. Click **OK** to save the assignment.

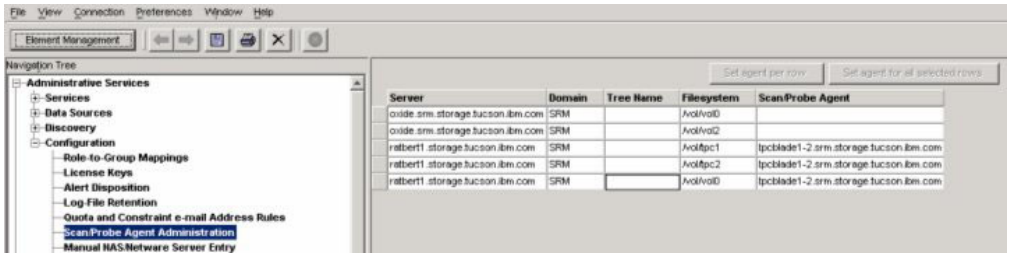


Figure 50. Scan/Probe Agent Administration panel

   d. Repeat steps b and c for other file systems (for example, /vol/vol0 and /vol/vol2 for oxide).
   e. Click the disk icon to save the **Scan/Probe Agent assignment** (so that the N Series Gateway servers and their file systems can be displayed when a probe or scan job is run).
4. Create a probe job to probe one or more N Series Gateway servers.
5. Create a scan job to scan one or more file systems from the N Series Gateway servers.
6. View asset report on a NAS gateway server and its file systems as **Other NAS**.
   a. Click **Administrative Services > Data Manager > Reporting > Asset > By OS Type > Other NAS**.
   b. Click an N Series Gateway server (for example, `ratbert1`) to view information for the server.

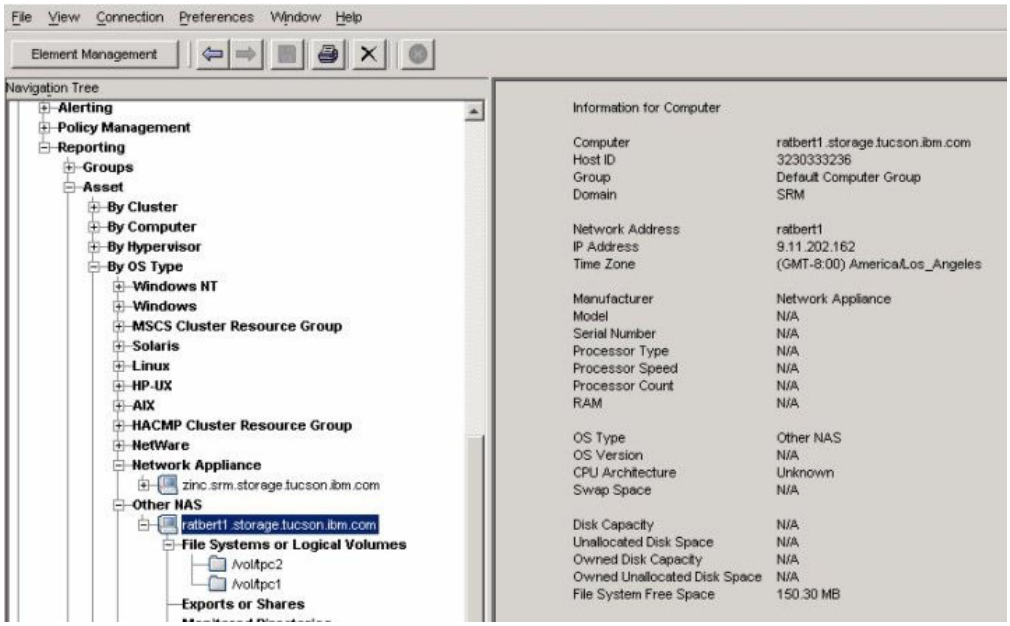


Figure 51. View information for N Series Gateway server

c. Click a file system (for example, `/vol/tpc2`) to view file system information.

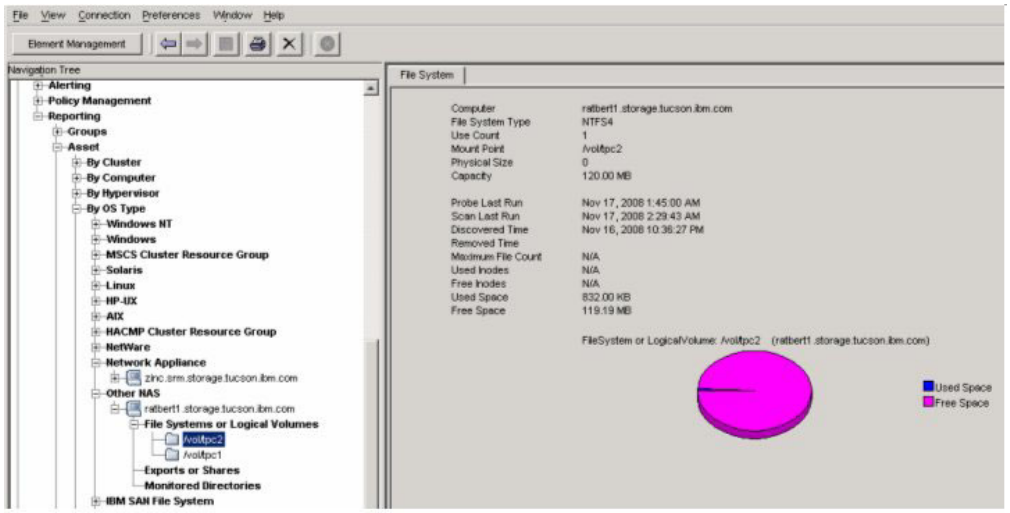


Figure 52. View information for a file system

7. You can see the following information and reports for **Other NAS**:
   a. Dashboard
   b. System reports
   c. Batch reports
   d. Asset reports
   e. Capacity reports
      - Capacity reports - Charting
   f. Usage reports
      - Usage reports - Filesystem Quota Violation
      - Capacity reports - Charting

8. You can monitor the N Series Gateway servers and create and monitor alerts for **Other NAS**:
   - Monitoring
     - Monitoring - Assign a Scan agent to a Filesystem
     - Monitoring - Grouping
     - Monitoring - Scan job
     - Monitoring - Filesystem Quotas
   - Alerting
     - Filesystem Alert definition
     - Alerting - Directory Alerts
     - Alerting - Alert Log
     - Alerting - Alerts shown on the topology

Chapter 5. Uninstalling the Tivoli Storage Productivity Center family

This section provides procedures for uninstalling the IBM Tivoli Storage Productivity Center family. The components must be uninstalled in a specific order.

Uninstall the components and related software in this order:

1. IBM Tivoli Storage Productivity Center:
   - Storage Resource agents, GUI, and CLI (these components are uninstalled using the Tivoli Storage Productivity Center uninstallation wizard)
   - Data Server or Device server (these components are uninstalled using the Tivoli Storage Productivity Center uninstallation wizard). When you are uninstalling the Device server and Data Server, you must also uninstall the DB2 database (remote or local) that was used by the Device server and Data Server.
   - DB schema (uninstalled using the Tivoli Storage Productivity Center uninstallation wizard)
2. IBM Tivoli Storage Productivity Center for Replication
3. DB2 (optional)

Note:

- On the Windows operating system, if you have installed the Monitoring Agent service, you must stop the **Monitoring Agent for Windows OS - Watchdog** and **Monitoring Agent for Windows OS - Primary** services before you uninstall Tivoli Storage Productivity Center.
- When uninstalling Tivoli Storage Productivity Center on AIX and Linux systems, you must have some free space on the hard disk drive to start the uninstallation wizard. This space is not checked by the uninstallation wizard. You must have at least 50 KB of free space in `/tmp` and 10 KB of free space in `/etc`.
- When you uninstall Tivoli Storage Productivity Center on Windows, there might be some directories or files that remain even if the uninstallation indicates that the operation was successful. If you do not delete all Tivoli Storage Productivity Center entries, you cannot reinstall Tivoli Storage Productivity Center. After uninstalling Tivoli Storage Productivity Center, check the `C:\Program Files\IBM` directory for any TPC entries. Delete any TPC entries and then reinstall Tivoli Storage Productivity Center.
- The uninstallation wizard uninstalls all components:
  - Tivoli Storage Productivity Center
  - IBM Tivoli Integrated Portal
  - IBM Tivoli Storage Productivity Center for Replication
- If you install the Storage Resource agent using the disk 1 image, you must uninstall the agent using the uninstallation wizard (InstallShield). If you uninstall the agent using the Tivoli Storage Productivity Center user interface, InstallShield does not recognize the uninstallation. Therefore, you must uninstall the agent that is installed from the disk 1 image using InstallShield.
- If you deploy the Storage Resource agent using the GUI, you must uninstall the agent using the GUI. You cannot uninstall the agent using the uninstallation wizard.

© Copyright IBM Corp. 2001, 2012

- If you receive an error message with a return code when uninstalling the Storage Resource agent, see the Tivoli Storage Productivity Center Information Center. Search for **Return codes used by the Storage Resource agent**.

Uninstalling Tivoli Storage Productivity Center in a single-server environment by using the wizard

In a single-server environment, you can uninstall Tivoli Storage Productivity Center by using a wizard. With the wizard, the uninstallation process requires minimal user interaction.

Source the DB2 profile `db2profile` for the instance owner of the DB2 database.

The following example sources the DB2 profile:

```
. /home/db2inst1/sqllib/db2profile
```

To uninstall the software in a single-server environment:

1. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
2. Start the Tivoli Storage Productivity Center uninstallation wizard in one of the following operating systems:
   - For the Windows operating system:
     a. Click **Start** > **Control Panel** > **Programs** > **Programs and Features**.
     b. Highlight **Tivoli Storage Productivity Center** and click **Uninstall/Change**.
   - For the AIX or Linux operating systems:

     Run the following command from the root directory:

     TPC_installation_directory`/_uninst/uninstall`

     where *TPC_installation_directory* is where Tivoli Storage Productivity Center is installed.
3. Review the message that indicates that all installed Tivoli Storage Productivity Center components is uninstalled.
4. Click **Uninstall** and click **Next**. You cannot cancel the uninstallation process after it starts.
5. Restart the system if needed.
   - For the AIX or Linux operating system, you do not have to restart the system.
   - For the Windows operating system, you can restart the system now or later. You must restart the system to reinstall Tivoli Storage Productivity Center.
6. If an error occurred during the installation process, review the installation log files.

Related tasks:

"Planning for IBM Tivoli Storage Productivity Center authorization" on page 33
When a user ID is authenticated to IBM Tivoli Storage Productivity Center through the GUI, CLI, or APIs, group membership (either an LDAP-compliant repository group or an operating system group) determines the authorization level of the user in IBM Tivoli Storage Productivity Center. This topic describes the user roles and authorization levels.

Uninstalling IBM Tivoli Storage Productivity Center for Replication with InstallShield

InstallShield is the recommended method for uninstalling IBM Tivoli Storage Productivity Center for Replication.

Tivoli Storage Productivity Center has its own uninstallation wizard. Ensure that you have completed the Tivoli Storage Productivity Center uninstalling procedures before completing the IBM Tivoli Storage Productivity Center for Replication uninstalling wizards.

Note:

- If you are using the LDAP server for authentication, you must make sure that the LDAP server is up and running before uninstalling IBM Tivoli Storage Productivity Center for Replication.
- Windows users can also uninstall IBM Tivoli Storage Productivity Center for Replication by clicking **Start** > **Settings** > **Control Panel** > **Add or Remove Programs**.

.

To uninstall IBM Tivoli Storage Productivity Center for Replication:

1. To start the uninstallation wizard, go to the *base_installation_directory*\TPCRM_uninst*base_installation_dir* folder, where *base_installation_directory* is the directory in which you installed the product. Then run the `uninstaller.bin` file (AIX and Linux) or the `uninstaller.exe` (Windows).On Windows, the default location is `C:\Program Files\IBM\replication\TPCRM_uninst`. On AIX and Linux, the default location is `/opt/IBM/replication/TPCRM_uninst`.
2. InstallShield displays the welcome page. Click **Next** to continue.
3. InstallShield displays a list of the features that can be uninstalled. There is only one feature, IBM Tivoli Storage Productivity Center for Replication. A summary of the previous steps including the features that you selected to uninstall. If the summary is correct, click **Uninstall**. A progress bar shows percentage complete.

Uninstalling IBM Tivoli Storage Productivity Center for Replication manually

You can manually uninstall IBM Tivoli Storage Productivity Center for Replication.

Uninstalling IBM Tivoli Storage Productivity Center for Replication manually on Windows

You can manually uninstall IBM Tivoli Storage Productivity Center for Replication on Windows.

IBM Tivoli Storage Productivity Center has its own uninstallation program. Ensure that you have completed the IBM Tivoli Storage Productivity Center uninstalling procedures before completing the IBM Tivoli Storage Productivity Center for Replication uninstalling procedure.

Note: If you are using the LDAP server for authentication, ensure that the LDAP server is up and running before uninstalling IBM Tivoli Storage Productivity Center for Replication.

1. Click **Start** > **Settings** > **Control Panel** > **Administrative Tools** > **Services**.
2. If IBM WebSphere Application Server V6.1 -CSM is listed, stop the service by right-clicking and selecting **Stop**.
3. From a command prompt, issue the following command:

   ```
   install_directory/eWAS/bin/wasservice.exe -remove CSM
   ```

   where *install_directory* is the directory where IBM Tivoli Storage Productivity Center for Replication is installed.

   Note: The default installation directory is `c:\Program Files\IBM\replication\eWAS\bin`.
4. Delete the directory where you installed IBM Tivoli Storage Productivity Center for Replication.
5. Delete the RM_Install.jacl file from the `C:\Documents and Settings\Administrator\Local settings\Temp` directory.
6. If IBM Tivoli Storage Productivity Center for Replication previously used DB2 for the database, open the DB2 command window and issue the following command:

   ```
   drop database dbname
   ```

   where *dbname* is the name of the IBM Tivoli Storage Productivity Center for Replication database.

Uninstalling IBM Tivoli Storage Productivity Center for Replication manually on AIX and Linux

You can manually uninstall IBM Tivoli Storage Productivity Center for Replication on AIX and Linux.

IBM Tivoli Storage Productivity Center has its own uninstallation wizard. Ensure that you have completed the IBM Tivoli Storage Productivity Center uninstalling procedures before completing the IBM Tivoli Storage Productivity Center for Replication uninstalling procedure.

Note: If you are using the LDAP server for authentication, you must make sure that the LDAP server is up and running before uninstalling IBM Tivoli Storage Productivity Center for Replication.

1. From a command prompt, issue the following command to stop the WebSphere Application Server process:

   ```
   installRoot/eWAS/profiles/CSM/bin/stopServer.sh server1
   ```

   where *installRoot* is the location where WebSphere Application Server is installed. If this command fails, perform the following steps:

   a. From a command prompt, issue the following command to obtain the process ID (PID) for the IBM Tivoli Storage Productivity Center for Replication service:

```
ps -ef | grep CSM
```

b. To stop the service, issue the following command:

```
kill -9 pid
```

where *pid* is the PID obtained in the previous step.

c. Verify that the process is no longer running by issuing this command:

```
installRoot/eWAS/profiles/CSM/bin/stopServer.sh server1
```

2. Delete the directory where the IBM Tivoli Storage Productivity Center for Replication was installed.
3. If IBM Tivoli Storage Productivity Center for Replication previously used DB2 for the database, issue the following command:

```
drop db dbname
```

where *dbname* is the name of the IBM Tivoli Storage Productivity Center for Replication database.

a. Change to the following directory:

```
. /home/db2_instance_name/sqllib/db2profile
```

b. Issue the following command:

```
db2 drop db dbname
```

4. Go to the `/etc/inittab` directory and edit the `inittab` file by removing the following entry:

```
/opt/IBM/replication/eWAS/bin/startServer.sh
```

Uninstalling the agents

When you installed the agents, you installed them locally or remotely. These topics describe how to uninstall these agents, both locally or remotely.

The Storage Resource agents are uninstalled using the Tivoli Storage Productivity Center GUI.

Deleting or uninstalling an agent using the GUI

You can delete or uninstall a Storage Resource agent and all the data that is collected by that agent from the database repository. The information that is collected by the agent is no longer available within Tivoli Storage Productivity Center reports. When you delete an agent, the agent is uninstalled and you can no longer activate the agent.

To delete or uninstall a Storage Resource agent, complete the following steps:

1. In the navigation tree, click **Administrative Services** > **Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select an agent and under the Select Action menu, click **Delete**.
3. In the confirmation dialog, select **Yes Delete the agent**.

Uninstalling the Storage Resource agent manually

This section describes how to uninstall the Storage Resource agent manually.

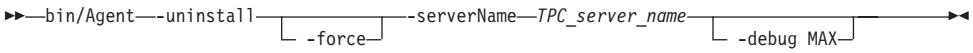
To uninstall the Storage Resource agent, follow these steps:

1. If the Storage Resource agent is on the Virtual I/O system, follow this log in procedure:

a. Log in with the **padmin** user ID.
b. Run the following command to set up the AIX environment:
   `oem_setup_env`

2. Go to the directory where the agent has been installed:
   agent_install_location

   Where *agent_install_location* is where the Storage Resource agent has been installed.
3. Run the uninstallation command:



   Where:

   -uninstall
   Uninstalls the agent.

   -force
   This optional parameter forces an uninstallation. If you use this parameter, do not provide the serverName parameter.

   -serverName *TPC_server_name*
   TPC_server_name is the Data server name as defined in IBM Tivoli Storage Productivity Center. You can check the configuration file for the server name:
   agent_install_location`/config/Agent.config`

   where *agent_install_location* is where is the agent is installed.

   -debug MAX
   This optional parameter is for debugging purposes. If you set the **-debug** parameter, then some files are not deleted.

   Note: If you run the uninstallation program from the bin directory, then the bin directory is not deleted.

   If you run the uninstallation program outside of the agent installation directory, then you must specify the full path.

   If the uninstallation fails, see the return codes in the Information Center. Search for **Return codes used by Storage Resource agent**.

Uninstalling the Language Pack

This procedure describes how to uninstall the Language Pack.

You must have administrator authority on Windows or root authority on UNIX or Linux to uninstall the Language Pack.

To uninstall the IBM Tivoli Storage Productivity Center components, follow these steps:

1. Log on to the system as a local administrator with the "Act as part of the operating system " user right on Windows. You must also have database administrative authority.
2. Exit the Tivoli Storage Productivity Center GUI.
3. Go to the following installation directory:

TPC_install_directory/_uninst_langpack

For UNIX and Linux, run the following command: **./uninstall.sh**. For Windows, run the following command: **uninstall.exe**.

4. The Select a language panel is displayed. Select a language and click **OK**.
5. The Welcome panel is displayed. Click **Next**.
6. The Select one or more Tivoli Storage Productivity Center products to uninstall panel is displayed. Select the components you want to uninstall. Click **Next**.

   Note: Tivoli Storage Productivity Center detects the language pack you installed and preselects the language pack to uninstall. You do not have to select a language pack to uninstall.
7. The summary information panel is displayed. Review the information and click **Uninstall** to uninstall the components.
8. The successfully uninstalled panel is displayed. Click **Finish**.

Uninstalling the Tivoli Storage Productivity Center Monitoring Agent

This procedure provides information about uninstalling the Tivoli Storage Productivity Center Monitoring Agent.

Uninstalling the IBM Tivoli Storage Productivity Center Monitoring Agent on Windows

This procedure provides information about uninstalling the IBM Tivoli Storage Productivity Center Monitoring Agent on Windows.

To uninstall the Tivoli Storage Productivity Center Monitoring Agent on Windows, complete the following steps:

1. Stop the IBM Tivoli Monitoring agent before uninstalling the agent. For information about stopping the agent for Windows, go to http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/index.jsp?topic=/com.ibm.omegamon.mes.doc_7.0.1/t_stop_win.html.
2. Log on to the system as a local administrator with administrative authority on Windows.
3. Exit the Tivoli Storage Productivity Center GUI.
4. For Windows, click **Start > Settings > Control Panel > Add/Remove Programs**. Highlight **IBM Tivoli Storage Productivity Center** and click **Remove**.
5. The select a language panel is displayed. Select a language and click **OK**.
6. The Welcome panel is displayed. Click **Next**.
7. Select one or more components to uninstall panel is displayed. Select **IBM Tivoli Storage Productivity Center Monitoring Agent**. Click **Next**.
8. The Summary Information panel is displayed. Review the information and click **Uninstall**.
9. The Uninstalling panel is displayed. This panel indicates what component is being uninstalled.
10. The status panel for uninstallation is displayed.
11. The successfully uninstalled panel is displayed. Click **Finish**.
12. The Restart Your Computer panel is displayed (for Windows only). To restart your computer, click **Yes, restart my computer**. Click **Finish** to complete the uninstallation operation.

Uninstalling the IBM Tivoli Storage Productivity Center Monitoring Agent on UNIX or Linux

This procedure provides information about uninstalling the IBM Tivoli Storage Productivity Center Monitoring Agent on UNIX or Linux.

To uninstall the Tivoli Storage Productivity Center Monitoring Agent on UNIX or Linux, complete the following steps:

1. Stop the IBM Tivoli Monitoring agent before uninstalling the agent. For information about stopping the agent on UNIX or Linux, go to http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/index.jsp?topic=/com.ibm.omegamon.mes.doc_7.0.1 and search for *Starting or stopping an agent on UNIX and Linux systems*.
2. For UNIX and Linux, you must have root authority.

   Note: For UNIX and Linux systems, you must source the DB2 profile before uninstalling the Tivoli Storage Productivity Center Monitoring Agent.
3. Open a terminal session and go to the Tivoli Monitoring Services installation directory. Run the uninstallation program from this default directory:

   ```
   /opt/IBM/ITM/bin/uninstall.sh
   ```
4. On the Manage Tivoli Enterprise Monitoring Services window, select **Monitoring Agent for TPC** and press **Enter**.
5. Confirm your selection by entering **1** and press **Enter** again.
6. Wait for the uninstallation program to complete.

For more information about uninstalling an agent from Tivoli Monitoring Services, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.1/itm_install427.htm.

Uninstalling DB2

Uninstall DB2 after you have uninstalled IBM Tivoli Storage Productivity Center and Agent Manager. Use the DB2 uninstallation program to uninstall DB2.

For information about uninstalling DB2 V9.1 on Windows systems, go to http://publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp?topic=/com.ibm.db2.udb.uprun.doc/doc/t0007436.htm.

For information about uninstalling DB2 V9.5 on Windows systems, go to http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/index.jsp?topic=/com.ibm.db2.luw.qb.server.doc/doc/t0007436.html.

Uninstalling DB2 on UNIX or Linux

Use these steps to uninstall DB2 on UNIX or Linux systems.

To uninstall DB2 on UNIX and Linux, follow these steps. This example applies to DB2 V9.7.

1. As the DB2 instance user (typically db2inst1), drop the databases and stop DB2.
   a. If you are the user with root authority, run **su - db2inst1** to switch the user to the DB2 instance user. Otherwise, log in as the DB2 instance user.
   b. Run **db2 list db directory** to list any databases that might not have been removed by the Tivoli Storage Productivity Center uninstallation program.

c. For each database name listed by the command in the previous step, run **db2 drop db** *database_name*. For example:

```
Command:
   $ db2 drop db tcpdb
Response:
   DB20000I  The DROP DATABASE command completed successfully.
   $ db2 drop db ibmcdb
Response:
   DB20000I  The DROP DATABASE command completed successfully.
```

d. Run **db2stop force**. The following messages are displayed:

```
08/10/2010 21:46:13     0   0   SQL1064N  DB2STOP processing was successful.
SQL1064N  DB2STOP processing was successful.
```

e. Run **db2 terminate**. The following message is displayed:

```
DB20000I The TERMINATE command completed successfully.
```

f. Run **exit** to close the su shell or log out.

2. Drop the DB2 instance.

   Note: This step must be performed as a user with root authority and run in a shell that has not had the db2profile sourced.

   a. Log in as a user with **root** authority.

   b. Change to the DB2 instance directory. The default path depends on your operating system.

      On AIX, run **cd /opt/IBM/db2/V9.7/instance**.

      On Linux, run **cd /opt/ibm/db2/V9.7/instance**.

   c. Run **./db2ilist** to list the db2instance name (typically db2inst1).

   d. Run **./db2idrop -f db2inst1** where "db2inst1" is the instance name reported by the **db2ilist** command in the previous step. The following messages are displayed:

```
DBI1324W  Support of the -f command is deprecated. For more information,
          see the DB2 Information Center.
DBI1070I  Program db2idrop completed successfully.
```

3. As the DB2 administration server (DAS) user (typically dasusr1), stop DAS.

   a. If you are the user with root suthority, run **su - dasusr1** to switch user to the DAS user. Otherwise, log in as the DAS user.

   b. Run **db2admin stop**. The following message is displayed:

```
SQL4407W The DB2 Administration Server was stopped successfully.
```

   c. Run **exit** to close the su shell or log out.

4. Remove the DB2 administration server. You must remove the DB2 administration server before you remove the DB2 product.

   a. Log in as a user with **root** authority.

   b. Change to the DB2 instance directory. The default path depends on your operating system.

      On AIX, run **cd /opt/IBM/db2/V9.7/instance**.

      On Linux, run **cd /opt/ibm/db2/V9.7/instance**.

   c. Run **./dasdrop**. The following messages are displayed:

```
SQL4410W The DB2 Administration Server is not active.
DBI1070I Program dasdrop completed successfully.
```

5. Uninstall DB2. Use the **db2_deinstall** command to remove DB2 products or DB2 components. The **db2_deinstall** command removes all DB2 products from your system.

a. Change to the DB2 installation directory. The default path depends on your operating system.

   On AIX, run **cd /opt/IBM/db2/V9.7/instance**.

   On Linux, run **cd /opt/ibm/db2/V9.7/instance**.

b. Run **./db2_deinstall -a**. The following messages are displayed:

```
DBI1016I Program db2_deinstall is performing uninstallation.
Please wait.
The execution completed successfully.
For more information see the DB2 uninstallation log at
"/tmp/db2_deinstall.log.pid".
```

6. Delete DB2 users and groups. The default user and group names are shown in the commands below.

   a. Run the following commands to remove the users:

```
userdel -r dasusr1
userdel -r db2inst1
userdel -r db2fenc1
```

   b. Run the appropriate commands for your operating system to remove the groups.

      On AIX run the following commands:

```
rmgroup dasadm1
rmgroup db2iadm1
rmgroup db2fadm1
```

      On Linux run the following commands:

```
groupdel dasadm1
groupdel db2iadm1
groupdel db2fadm1
```

7. Remove the DB2 directory.

   a. Run **cd**.

   b. Run the command to remove the DB2 directory on your operating system:

      On AIX, run **rm -r /opt/IBM/db2**.

      On Linux, run **rm -r /opt/ibm/db2**.

8. Remove any DB2 definitions from `/etc/services`.

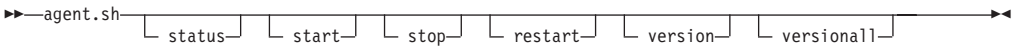
Chapter 6. Reference

agent.sh command

The **agent.sh** command lets you start, stop, and restart the Storage Resource agent. You can also display the status and version of the Storage Resource agent.

Note:

- You must have root authority to run this command.
- For Windows, the status, stop, and start functions are handled through the Windows Services panel.



Parameters:

status
Displays the current status of the Storage Resource agent. The status is returned is running or not running.

start
Starts the Storage Resource agent.

stop
Stops the Storage Resource agent.

restart
Stop and then start the Storage Resource agent.

version
Displays the current version of the Storage Resource agent.

versionall
Displays the version of the Storage Resource agent and its related components.

Estimating tablespace allocation

This topic describes how you can use the IBM Tivoli Storage Productivity Center to estimate the tablespace size in the repository database.

When you install Tivoli Storage Productivity Center, you can specify the tablespace size in the repository database. The space needed for the Tivoli Storage Productivity Center database schema varies significantly with storage network configuration, data collection, data retention period, and other factors. Unless you are an experienced administrator, you should use the default values for tablespace allocation that Tivoli Storage Productivity Center provides.

Table 45 on page 428 provides space estimations for a storage configuration containing 5000 volumes with some general assumptions.

© Copyright IBM Corp. 2001, 2012

Table 45. Tablespace allocation for the Tivoli Storage Productivity Center database schema

Tablespace	Description of tablespace usage	Recommended size for a 5000 volume configuration	Assumptions
KEY	This tablespace is used for configuration data which is constantly used. For example, the key entity and relationships data (T_RES_STORAGE_ SUBSYSTEM, T_RES_STORAGE_VOLUME, and the normalization tables, and so forth)	500 MB	A table that uses significant space is T_RES_DATA_PATH. This table uses about 300 bytes for each record for the relationship between the storage volumes and host ports. There could be dozens to hundreds of data paths for a volume.
NORMAL	This tablespace is used for snapshots and miscellaneous data	500 MB	A table that uses significant space is T_RES_STORAGE_VOLUME_ SNAPSHOT. This table uses about 2500 bytes for each record. The number of snapshots depends on the data collection activities.
BIG	This tablespace is used for performance statistics	2 to 3 GB or 400 MB per day of performance data	The data collected for performance data for storage volumes can use a significant amount of space (about 200 bytes for each record). For 5000 volumes, if performance data is collected every 5 minutes, the data for one day would be 300 MB. If the data is kept for 7 days, the data collected would take about 2 to 3 GBs. If the data is kept longer, the storage must be scaled up accordingly.
TEMP	This tablespace is used for temporary data for query processing and other temporary tables	1 GB	

Valid characters for user IDs and passwords

This section provides information about valid user IDs and passwords for Tivoli Storage Productivity Center.

If you are performing a typical or custom installation, the valid characters for the common user ID and password are as follows. This includes the user IDs and passwords for the DB2 administrator, DB2 user, WebSphere administrator, host authentication password, Windows service for the Storage Resource agent, and NAS filer.

For the user ID, you can have these characters:

- A through Z (uppercase characters)
- a through z (lowercase characters)
- 0 through 9 (numeric characters)
- Special characters: - _ .

Do not use these special characters:

```
! @ $ * = + [ ] \ | ; : " , < > / ? .
```

Tivoli Storage Productivity Center does not check the minimum and maximum length for the user ID and password for installation. The user ID and password to install Tivoli Storage Productivity Center must conform to the following rules:

- DB2 naming rules. For information about DB2 naming rules, see http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp. Search for **General naming rules**.
- WebSphere naming rules. For information about WebSphere naming rules, see http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp. Search for **WebSphere MQ naming restrictions**.
- Operating system.

For Windows, the password names are based on the Local Password Policy.

Important: For Tivoli Storage Productivity Center, do not use a password that begins with a dash character.

For custom installation, the valid characters are shown in the following table.

Table 46. Valid characters for user IDs and passwords

User ID and password for	Valid characters
WebSphere administrator or host authentication password	• A through Z (uppercase characters) • a through z (lowercase characters) • 0 through 9 (numeric characters) • Special characters: - _ .
Windows Service user ID (for Storage Resource agent) or NAS Filer user ID	• A through Z (uppercase characters) • a through z (lowercase characters) • 0 through 9 (numeric characters) • Special characters: ` ~ # % ^ & ( ) - _ { } ' .
Windows Service password (for Storage Resource agent) or NAS Filer password	For Windows, the password is based on the Local Password Policy.

For information about user IDs and passwords required for installing and using Tivoli Storage Productivity Center, see "User IDs and passwords for installing and using Tivoli Storage Productivity Center" on page 37.

TCP/IP ports used by the IBM Tivoli Storage Productivity Center family

This topic lists the default ports that must be opened through the firewall when you install the IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication servers. You must disable the firewall program or open the ports to allow incoming requests to the IBM Tivoli Storage Productivity Center ports. Review these ports before installing the IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication.

TCP/IP ports used by IBM Tivoli Storage Productivity Center

See Table 12 on page 34 for the IBM Tivoli Storage Productivity Center ports.

Table 47. TCP/IP ports used by IBM Tivoli Storage Productivity Center

Component	Session initiator (server perspective)	Inbound/ outbound (server perspective)	Port	Inbound/ outbound (agent perspective)	Session initiator (agent perspective)
Data server		Both	Default: 9549 (see notes 1 on page 35 and 3 on page 35)		
Device server		Both	9550 (see notes 2 on page 35 and 3 on page 35)		
Remote installation of UNIX agent	Yes	Outbound	SSH (22)	Both	No
Remote installation of Windows agent	Yes	Outbound	NetBIOS Sessions Service (139)		
Remote installation of UNIX agent	Yes	Outbound	RSH (514)	Both	No
Remote installation of UNIX agent	Yes	Outbound	REXEC (512)	Both	No
Remote installation of UNIX agent	Yes	Inbound	601		
Remote installation of all agents	Yes	Inbound	Data Server 9549		
Device server to CIM agent	Yes	Outbound	Default: HTTP: 5988 HTTPS: 5989		
VMware VI Data Source to Virtual Center or ESX server	Yes	Outbound	Default: HTTP: 80 HTTPS: 443		
Tivoli Integrated Portal server. See note 6 on page 35.	Yes	Outbound	389		

Table 47. TCP/IP ports used by IBM Tivoli Storage Productivity Center (continued)

Component	Session initiator (server perspective)	Inbound/ outbound (server perspective)	Port	Inbound/ outbound (agent perspective)	Session initiator (agent perspective)
Tivoli Integrated Portal server	Yes	Inbound	16310, 16311, 16312, 16313, 16315, 16316, 16318, 16320, 16322, 16323		
Device server. See note 6 on page 35.	Yes	Outbound	389		
IBM Tivoli Storage Productivity Center for Replication server. See note 6 on page 35.	Yes	Outbound	389		

1. Data server inbound server port plus another port that is 10 greater than the Data server port (9549). For example, if the Data server port is 9549, then another port would be 9549+10 or 9559.
2. If you select the default port (9550) for the Device server during installation, the following ports are also used: 9551, 9552, 9556, 9557. If you use a port other than the default port for the Device sever, make sure that port and the ports above it are available for use. For example, if N represents the open port you want to use, ensure that ports N+1, N+2, N+6, and N+7 are open as well (for example, if you select 9570, then the following ports must be open as well: 9571, 9572, 9576, 9577).
3. You must select different ports for the Data server and Device server. For example, if N is the device server port, and Z the data server port, make sure that N+1, N+2, N+6 and N+7 are not the same ports as Z and Z+10.
4. When you install the Storage Resource agent, you can specify a listener port for a remote probe. This port is used by the program to communicate back to the installation program.
5. You can find the port numbers used on your system by running the following command:

   ```
   netstat -an
   ```

6. If you intend to install IBM Tivoli Storage Productivity Center so that it authenticates users against an LDAP-compliant repository, then the embedded WebSphere Application Servers for Tivoli Integrated Portal, IBM Tivoli Storage Productivity Center Device server, and IBM Tivoli Storage Productivity Center for Replication need to be able to initiate an outbound connection to the LDAP-compliant repository, which typically listens on port 389.

Other TCP/IP ports used by IBM Tivoli Storage Productivity Center

Table 48. Other TCP/IP ports used by IBM Tivoli Storage Productivity Center

Port	Default
DB2	50000

Table 48. Other TCP/IP ports used by IBM Tivoli Storage Productivity Center (continued)

Port	Default

• SAN Volume Controller
• IBM Storwize V7000
• IBM Storwize V7000 Unified
	For native interface:
	`Inbound: none`
	`Outbound (on SVC or Storwize V7000 cluster):`
	`22 for SSH`

CIM agent for IBM TotalStorage Enterprise Storage Server (ESS)	5989

DS8000	For native interface:
	`Inbound for events: 1750`
	`Outbound (on HMC): 1750`
	`Inbound and outbound on HMC for offloading logs: 1755`

IBM XIV Storage System	For native interface:
	`Inbound: none`
	`Outbound (on XIV storage system): 7777, 7778`

IBM Tivoli Storage Productivity Center with DS8000 GUI	`8451` `8452`
IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication	162 (default SNMP listening port)
WebSphere® Application Server	9043, 9060, 9353, 5060, 5061

TCP/IP ports used by IBM Tivoli Storage Productivity Center for Replication

Table 14 on page 36 shows the ports used for incoming communication. Table 15 on page 37 shows the ports used for outgoing communication.

Table 49. TCP/IP ports used by IBM Tivoli Storage Productivity Center for Replication for incoming communication

Port	Description
3080 and 3443	A Web browser typically communicates with the IBM Tivoli Storage Productivity Center for Replication GUI using HTTP ports 3080 and 3443. These ports are configurable from the IBM Tivoli Storage Productivity Center for Replication installer during installation.
5110	The IBM Tivoli Storage Productivity Center for Replication command line and GUI typically communicate with the IBM Tivoli Storage Productivity Center for Replication server using port 5110.
5120	IBM Tivoli Storage Productivity Center for Replication uses port 5120 for communication with other IBM Tivoli Storage Productivity Center for Replication servers for high-availability purposes.
Note: If you changed ports 3080 or 3443 during the installation or changed other port settings, note the values to which these settings were changed.	

Table 50. TCP/IP ports used by IBM Tivoli Storage Productivity Center for Replication for outgoing communication

Port	Description
2433	For communication with ESS/DS (direct connection).
1750	For communication with DS8000 (HMC connection).
443 and 22	For communication with SAN Volume Controller, Storwize V7000 Unified or Storwize V7000 clusters.
9930	For communication with the AIX host for the Open HyperSwap.

Note:

- If you changed the port configuration of your storage controller, the ports would be different.
- Your network configuration must allow for IBM Tivoli Storage Productivity Center for Replication to send outgoing TCP/IP packets to the storage controllers. It is possible when adding the storage controllers to IBM Tivoli Storage Productivity Center for Replication to set a specific port number for your storage controller.
- Since there are typically multiple applications running on a server, it is possible that port conflicts might arise if other applications attempt to use the same ports that IBM Tivoli Storage Productivity Center for Replication is using. You can find the port numbers used on your system by running the following command:

  ```
  netstat -an
  ```

- If firewalls are being used in your configuration, make sure that none of these ports are being blocked. Ensure that not only is the IBM 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 IBM Tivoli Storage Productivity Center for Replication server.
- If you are running Windows Server 2008, you need to configure the inbound and outbound rules for IBM Tivoli Storage Productivity Center for Replication. To create an outbound rule, launch the New Outbound Rule wizard from the Windows Firewall with Advanced Security menu.
- If you are installing or deploying Storage Resource agents on a system, you must first disable the firewall on that system before installing or deploying an agent.

Appendix A. Work sheet for user IDs and passwords

This section provides a work sheet to help plan your user IDs and passwords during the installation of IBM Tivoli Storage Productivity Center.

Table 51. User IDs and passwords for Tivoli Storage Productivity Center installation

Item	Description	Your input
DB2 administrator user ID and password	This user ID and password are required to install Tivoli Storage Productivity Center. On Windows, this user ID must be a member of the DB2ADMNS group and Administrators group. On UNIX, the user ID must be the instance owner of the instance you want to use. This user ID and password are created when you install DB2.	
Storage Resource agent service logon user ID and password	This user ID and password are for Windows only and is optional. This user ID and password creates a new service account for the Storage Resource agent to run under. If your enterprise has a security policy that requires long passwords, the installation can fail. This user ID and password allow you to run under a local account user name and password that meets the security policy for your enterprise. This user ID and password are created if you do not have one.	No default. If this user ID and password do not exist, they are created at installation time.
Host authentication password	This password is required for: • the Storage Resource agent to communicate with the Device server • the subordinate server when added to the master server (for enterprise-wide reporting) Specified when you install the Device server. You must specify this password when you install the Storage Resource agent.	This password is created when you install the Device serverr.

© IBM Corporation 2009 © IBM 2001, 2012

Table 51. User IDs and passwords for Tivoli Storage Productivity Center installation (continued)

Item	Description	Your input
NAS filer login user ID and password	This user ID and password are for Windows only. This user ID and password are the default user name and password used by Storage Resource agents during NAS discovery.	
WebSphere administrator user ID and password	This user ID and password are required by the Device server to communicate with embedded WebSphere. Embedded WebSphere runs as a Windows Service. The Windows Service runs under the authority of this user ID and password. This user ID and password are only used when you install the Device server.	If this user ID and password do not exist, they are created at installation time.
LDAP Tivoli Storage Productivity Center Administrator username and password and LDAP Tivoli Storage Productivity Center Administrator group	These values are relevant if you install Tivoli Storage Productivity Center so that it authenticates users against an LDAP-compliant repository. The username, password, and group must exist in your LDAP-compliant repository in the "branches" specified by the "Relative Distinguished Name for usernames" and the "Relative Distinguished Name for groups" values. The user must also be a member of the group.	During the Tivoli Storage Productivity Center installation, the LDAP Tivoli Storage Productivity Center Administrator username value is used to set the WebSphere Administrative ID value for the Tivoli Integrated Portal server, the Tivoli Storage Productivity Center Device server, and the Tivoli Storage Productivity Center for Replication server. Also during the Tivoli Storage Productivity Center installation, the LDAP Tivoli Storage Productivity Center Administrator group value is mapped to the "superuser" role in Tivoli Integrated Portal, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication to allow initial access to the applications once the installation is complete.

For information about valid user IDs and passwords, see “Valid characters for user IDs and passwords” on page 428

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:

- Tivoli Storage Productivity Center functions are available using the keyboard for navigation instead of the mouse. You can use keys or key combinations to perform operations that can also be done using a mouse. 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.
- You can use screen readers to read the user interface.
- The user interface communicates all information independently of color.
- The *IBM Tivoli Storage Productivity Center Information Center*, and its related publications are accessibility-enabled and include the following accessibility features:
  - The information center is provided in XHTML 1.0 format, which is viewable in most Web browsers. XHTML allows you to view documentation according to the display preferences set in your browser. With XHTML, you can use screen readers and other assistive technologies.
  - All documentation for Tivoli Storage Productivity Center is available in Adobe Portable Document Format (PDF) 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/v4r1/index.jsp.
  - All images are provided with alternative text, so that users with vision impairments can understand the contents of the images.

Keyboard navigation

This product uses standard Microsoft Windows navigation keys.

Displaying the interface

Click **Preferences** > **Look and Feel** to select how to display the Tivoli Storage Productivity Center user interface. To do so, complete the following steps:

1. Start the Tivoli Storage Productivity Center user interface.
2. Click **Preferences** > **Look and Feel** to change the visual appearance of the user interface to best suit your visual needs:
   - Windows Classic
   - Windows
   - CDE/Motif
   - Metal

© Copyright IBM Corp. 2001, 2012

Limitations for sight-impaired users on IBM Tivoli Storage Productivity Center for Replication

There some limitations in the product graphical user interface (GUI) that requires special attention before you start using Tivoli Storage Productivity Center for Replication with a screen reader:

Dojo components are not read by all screen readers.

The JAWS screen reader does not read some Dojo components on Internet Explorer 7. Use the command-line interface instead of the GUI with JAWS on Internet Explorer 7.

IBM and accessibility

See the IBM Human Ability and Accessibility Center Web site for more information about the commitment that IBM has to accessibility.

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

See the IBM Human Ability and Accessibility Center website at www.ibm.com/able for more information about the commitment that IBM has for accessibility.

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 (the right side), 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 in the left side, press Alt+R, and then press Enter or Up Arrow to enter the view.
- To go directly to the Navigation (Table of Contents) view in the left side, press Alt+C, and then press Enter or Up Arrow to enter the view.
- To expand and collapse a node in the navigation tree, press the Right and Left Arrow.
- To move to the next topic node, press the Down Arrow or Tab.
- To move to the previous topic node, press the Up Arrow or Shift+Tab.

© Copyright IBM Corp. 2001, 2012

- To go to the next link, button, or topic node from inside on of the views, press Tab.
- 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 above 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 above 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 Firefox and Internet Explorer GUIs. 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 Internet Explorer 7. Use the command-line interface instead of the GUI with JAWS on Internet Explorer 7.

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.

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.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 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, 2012

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.

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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

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.s

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corporation, 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 is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Intel Xeon, Itanium, and Pentium 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 are trademarks of Microsoft Corporation in the United States, other countries, or both.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Glossary

This glossary includes terms and definitions for IBM Tivoli Storage Productivity Center.

The following cross-references are used in this glossary:

- *See* refers the reader from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- *See also* refers the reader to a related or contrasting term.

To view glossaries for other IBM products, go to: http://www.ibm.com/software/globalization/terminology/.

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.

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.

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 SAN Volume Controller or Storwize V7000, a pair of nodes that provides a single configuration and service interface. 2. 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 information. The Distributed Management Task Force (DMTF) develops and maintains CIM specifications.

community name
The part of a Simple Network Management Protocol (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 source and target volumes in a session that contain copies of the same data. A session can contain one or more copy sets.

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.

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*.

discovery
The process of finding resources within an enterprise, including finding the new location of monitored resources that were moved. See also *schedule*.

enterprise repository
A component of the Data server that

© Copyright IBM Corp. 2001, 2012

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 Data server using Java Database Connectivity (JDBC).

extended count key data (ECKD™)
An extension of the count key data (CKD) architecture.

FlashCopy
In storage systems, a feature that creates a point-in-time copy of a volume or set of volumes on a specified target volume.

flush-through mode
A process in which data is written to a storage device at the same time as the data is cached.

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*.

GiB See *gibibyte*.

gibibyte (GiB)
A base-2 unit of measurement that is equal to 1,073,741,824 bytes (2 to the 30th power). See also *gigabyte*.

gigabyte (GB)
In decimal notation, 1,073,741,824 bytes, when referring to memory capacity; in all other cases, it is defined as 1,000,000,000.

Global Copy
A non-synchronous long-distance copy option for data migration and backup. See also *remote mirror and copy* and *Metro Mirror*.

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 *Metro Mirror* and *remote mirror and copy*.

Globally Unique Identifier (GUID)
An algorithmically determined number that uniquely identifies an entity within a system.

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*.

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 where colors are used to differentiate values in a data set.

high availability
Pertaining to a clustered system that is reconfigured when node or daemon failures occur, so that workloads can be redistributed to the remaining nodes in the cluster.

high-availability solution
A combination of hardware, software, and services that fully automates the recovery process and does not disrupt user activity. High-availability solutions must provide an immediate recovery point with a fast recovery time.

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.

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. Contrast with *out-of-band discovery*.

job A unit of work that monitors, manages, or collects data about a storage resource.

job scheduler
A component of the Data server that deploys all monitoring activities.

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.

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
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.

logical unit number (LUN)
In the Small Computer System Interface (SCSI) standard, a unique identifier that is used to differentiate devices, each of which is a logical unit (LU).

LUN See *logical unit number*.

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.

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.

Metro Mirror
A method of synchronous remote replication that operates over distances up to 300 kilometers. This method constantly updates a secondary copy of a volume to match changes that are made to a source volume. See also *remote mirror and copy* and *Global Copy*.

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.

MiB See *mebibyte*.

MIB See *Management Information Base*.

native interface
An interface that is specific to a system or subsystem.

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 Simple Network Management Protocol (SNMP) Management Information Base (MIB) queries, which are invoked over an IP network. Contrast with *in-band discovery*.

ping 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. See also *schedule*.

pool A storage pool or pool is 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 deallocated on a storage device. Storage pools are created by allocating storage capacity from primordial pools.

probe A job that itemizes and creates an inventory of assets, such as computers,

controllers, disk drives, file systems, and logical units. See also *schedule*.

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 *Metro Mirror* and *Global Mirror*.

resource
In a storage environment, an entity that is monitored. Resources can include fabrics, switches, computers, storage systems, and tape libraries.

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).

scan job
A job that monitors the usage and consumption of storage. See also *schedule*.

schedule
A monitoring task that collects data about storage resources. Monitoring tasks include scans, probes, pings, discovery, and performance monitors. For example, define a probe schedule to collect data about storage systems; define a scan schedule to collect file-level information about hosts.

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 that consists 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, tape libraries, and other storage resource groups.

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 (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.

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*.

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 that you want to manage together.

zone set
A group of zones that function together on a fabric.

Index

Numerics

64-bit environment supported 20

A

about this document xiv
accessibility features
  for Tivoli Storage Productivity Center 437
accessibility features for users with disabilities 439
adding
  installation license 227
Agent Manager
  registering 307
agent.sh command
  syntax 427
agents
  assigning agents for scan and probe jobs 302
  configuration file 306
  configuring for Storage Resource agent upgrade 326
  deploying Storage Resource agents 308
  deployment considerations for Storage Resource agents 312
  hardware requirements 4
  installing 232
  overview 255
  planning for CIM agents 77
  planning for Storage Resource agents 49
  planning for TagmaStore CIM agents 78
  preparing Storage Resource agents for installation 229, 312
alerts, overview 258
authorization
  for users 33

B

backing up
  system 397
Baroc file location 299
Brocade
  planning for DCFM 101, 103
browsers supported 17

C

CAP file 141, 143
certificates
  creating for Storage Resource agents 317
  replacing for Storage Resource agents 322
CIM agents
  planning for 77
  planning for TagmaStore 78
  software requirements 19
CLI
  installing 227
cluster resource group
  for PowerHA 124
cluster resource groups 122
commands
  agent.sh 427
comments, sending xiv
configuration
  analysis tool 63
  authorizing users 291
  editing nas.config file 306
  for Device server 111
  for NetWare 302
  for Network Attached Storage 302
  for switches 335
  guidelines 31
  History Aggregator 308
  history settings 335
  history tool 65
  NetApp 129
  overview 255
  planning for Fabric Manager 111
  planning for HBAs 115
  planning for roll up reports 62
  planning for VMware 135
  Removed Resource Retention 333
  Resource History Retention 332
  role-to-group mapping 291
  the Device server 335
configuration files
  agent..config file 305
  server.config file 305
configuration utility
  how to access Tivoli Storage Productivity Center for Replication 265
  how to create alert definitions 263
  how to display availability reports 262
  how to manage tape library reports 264
  how to monitor storage assets, storage systems and fabrics 262
  learn how to manage element managers and external tools 265
  learn how to manage servers and agents 261
  show managed devices 262
  tasks 261
configuring
  Data server for Windows domain 342
  DB2 JDBC driver 352
  inactive zone sets for fabric data source 338
  Java Web Start 347
configuring *(continued)*
  Monitoring agent 370
  Storage Resource agent for upgrade 326
credentials, migrating 381

D

Data agents
  migrating using installation wizard 409
  uninstalling using GUI 421
data collection 185, 380
Data Manager
  agent..config file 305
  databases that can be monitored 44
Data Path Explorer view 72
Data server
  installing using local database 205
  installing using remote database 220
  preparing for using remote database 214
database repository
  installing 203
  planning to retrieve information from DB2 views 139, 140
  software requirements 15
databases
  histories, specifying how long to keep 334
  specifying how long to keep information 334
  that Data Manager can monitor 44
DB2
  installing on AIX 162, 164
  installing on Linux 159
  installing on UNIX 159
  installing on Windows 157
  migrating repository database on AIX 391
  migrating repository database on Linux 394
  migrating repository database on Windows 387
  monitoring the DB2 environment 141
  preparing to install on UNIX 156
  preparing to install on Windows 156
  uninstalling 424
  upgrading 383, 386
  verifying installation 166
  verifying installation using CLP 166
  verifying installation using First Steps 166
DCFM
  planning for 101, 103
Device server
  configuration 111, 335
  installing using local database 205
  installing using remote database 220
  preparing for using remote database 214

© Copyright IBM Corp. 2001, 2012

Device server *(continued)*
setting timeout values 339
disabilities, accessibility features for users with 439
Discovery Library Adpater 141
domain account
how to grant local administrative privileges 343
drop-down lists, limitations of to the sight-impaired 439
DS8000
how to manage DS8000 element managers 265
planning for 86
planning for performance 87

E

editing configuration files 305
EMC PowerPath Multipathing
planning for 49

F

Fabric agents
migrating using installation wizard 409
fabric management
planning for
Brocade and McData 99
QLogic and Cisco 103
Fabric Manager
directors 114
planning for 98
planning for configuration 111
planning for HBAs 115
switches 114
fabrics 65
file systems
supported 45
fully qualified host names
required 154
fully-qualified host names
host names
fully-qualified host names
required 145
required 145

G

glossary 447
GUI
installing 226

H

hardware requirements 2
Tivoli Storage Productivity Center for Replication 2
Tivoli Storage Productivity Center GUI 5
HBA 65
HBAs
planning for Fabric Manager 115
History Aggregator, configuring 308
host port 65
hosts 65

I

IBM Tivoli Storage Productivity Center
installing 167
migrating 147
upgrading 147
IBM Tivoli Storage Productivity Center for Replication
installing 167
installing Two Site or Three Site BC license 185
installing using custom installation 202
installing using typical installation 182
migrating 147
upgrading 147
installing
agents 228, 232
CLI 227
database repository 203
DB2 on AIX 162
DB2 on AIX POWER7 164
DB2 on Linux 159
DB2 on UNIX 159
DB2 on Windows 157
DB2, preparing for installation on UNIX 156
DB2, preparing for installation on Windows 156
Device server using local database 205
GUI 226
TIP 144
IBM Tivoli Storage Productivity Center for Replication using custom installation 202
IBM Tivoli Storage Productivity Center for Replication using typical installation 182
Language Pack 244
Monitoring agent 236
new installation 144
overview 153
preparing for 167
starting installation 169
Storage Resource agent considerations 312
Storage Resource agents 308
Storage Resource agents locally 229
Storage Resource agents remotely 229, 312
Storage Resource agents using installation program 232
TIP
installing 144
Tivoli Integrated Portal using typical installation 176
Tivoli Storage Productivity Center 172, 185
Tivoli Storage Productivity Center for Replication 172, 185
Two Site or Three Site Business Continuity license 185
installing *(continued)*
using custom mode 185
using typical mode 172
InstallShield
limitations 155
interoperability namespaces 84
IPv6
configuring
AIX for IPv6 377
DB2 on AIX for IPv6 377
DB2 on Linux for IPv6 378
configuring DB2 on Linux 378
configuring for AIX 377
configuring for DB2 on AIX 377
limitations 54
planning for 52

J

Java Web Start
configuring 347
JDBC driver
configuring 352
jobs, overview 257

L

Language Pack
installing 244
uninstalling 244, 422
LDAP
planning for 58
LDAP servers
software requirements 19
license
adding 227
licenses 30
log files
setting retention period 300

M

MIB file location 299
migrating
Data agents 407
Fabric agents 407
overview 147, 379
monitoring
databases 44
file systems 45
MSCS clustered nodes and cluster resource groups 121
networked file systems 46
overview 185, 380
PowerHA clustered nodes and cluster resource groups 126
volume managers 47
Monitoring agent
configuring 370
configuring for Windows 370
configuring on UNIX 371
installing 236
installing and configuring 236, 240
starting 371
stopping 371

Monitoring Agent
uninstalling 423
uninstalling on UNIX 424
uninstalling on Windows 423
MSCS support 121
multipath 65
multipath subsystem device drivers
planning for 47
upgrading 49

N

NAS
support of 128
NAS configuration file, editing 306
nas.config file, editing 306
native interfaces
planning for 57
NetApp
configuration 129
identifying in user interface 131
limitations 130
planning for 129
removing a device 130
NetWare
configuring for 302
logging into NDS tree 332
Network Attached Storage
configuring for 302
networked file systems
supported 46
Novell Directory Services
logging into 332

O

operating systems supported 5

P

partitiontables.bat file 381
partitiontables.sh file 381
passwords
valid characters 428
work sheet 40, 435
path planning 65
performance
configuration 119
planning for DS8000 87
performance database migration tool 381
performance metrics 117
performance monitor 65
planning 33
for an MSCS environment 121
for CIM agents 77
for Data Path Explorer view 72
for DS8000 86
for MSCS support 121
for NAS system support 128
for retrieving information from DB2 views 139, 140
for roll up reports configuration 62
for rollup reports 61
for SAN configuration tools 63
planning *(continued)*
for Solaris Storage Resource agents 75
for System Storage N Series Gateway servers 132
for TagmaStore CIM agents 78
for tape library support 44
for Tivoli Storage Productivity Center for Disk Select 82
for Tivoli Storage Productivity Center for Replication 1
for VMware configuration 135
for VMware support 132
for volume manager support 47
for XIV system 96
IPv6 support 52, 54
NetApp support 129
overview 1
storage system support 42
switches 43
Virtual I/O Server support 136
PowerHA
environment 126
planning for 124
upgrading agents 126
probe 65

Q

quota and constraint e-mail address rules 301

R

RAID 65
reader feedback, sending xiv
replication planning 65
reports
configuring for data aggregation 308
specifying how long to keep information 334
VMware capacity 135
resource histories
setting retention periods 332
resources
specifying how long to keep 333
roll up reports
planning for configuration 62
rollup reports
planning for 61

S

SAN configuration tools
planning 63
SAN Planner 65
SAN Volume Controller
publications xi
SDD 65
sending comments xiv
sight-impaired 439
single sign-on
planning for 60
user authentication method used with 60
SMI Providers 84
SNMP
sending traps 299
software requirements
database repository 15
for CIM agents 19
for LDAP servers 19
operating systems 5
SSH protocol
creating a certificate 317
SSL certificates
replacing a certificate 322
storage capacity of volumes 82
Storage Optimizer
analysis report 71
optimization report 71
supported storage systems 71
storage pools 65
storage port 65
Storage Resource agents
configuring for upgrade 326
creating certificates for 317
deploying 308
deployment considerations 312
importing authentication information 358
installing 308
installing locally 229
installing using installation program 232
monitoring by Tivoli Agent Management Services 141, 143
planning for 49
planning for Solaris 75
preparing for 229
protocol support 50
replacing SSL certificates 322
uninstalling manually 421
uninstalling using GUI 421
upgrading manually 406, 408
storage resource groups 65
Storage Subsystem Credential Migration tool 381
storage system supported 42
subsystem device drivers
planning for 47
upgrading 49
switch 65
switch zone 65
switches
configuring 335
planning for 114
supported 43
using private networks 116
System Storage N Series Gateway servers
planning for 132
upgrading 410

T

tape library support 44
TCP/IP ports
planning for 34, 35, 430, 431
used by IBM Tivoli Storage Productivity Center for Replication 36, 432
TEC events
sending 299

time zone 65
timeout values
  setting for Device server 339
Tivoli Agent Management Services
  monitoring Storage Resource agents 141, 143
Tivoli Common Reporting
  installing and configuring 248
Tivoli Integrated Portal
  installing using customized installation 197
  installing using typical installation 176
Tivoli Storage Productivity Center
  disabling 368
  hardware requirements 2
  installing overview 153
  installing using custom mode 185
  installing using typical mode 172
  Lightweight Directory Access Protocol SAN Volume Controller remote authentication support 366
  migrating 379
  software requirements 5
  uninstalling 417, 418
  uninstalling Monitoring Agent 423
  uninstalling Monitoring Agent on UNIX 424
  uninstalling Monitoring Agent on Winsows 423
  uninstalling the agents 421
  upgrading 379
Tivoli Storage Productivity Center for Disk Select
  planning for 82
Tivoli Storage Productivity Center for Replication
  disabling 368
  how to access through Tivoli Storage Productivity Center 265
  installing overview 153
  installing using custom mode 185
  installing using typical mode 172
  migrating 379
  planning for 1
  software requirements 5
  uninstalling 418
  uninstalling with InstallShield 419
  upgrading 379, 383
topology 65
TotalStorage Productivity Center
  upgrading 382
trademarks 444
translations
  browser locale requirement xiii

U

uninstalling
  AIX 420
  Data agent 421
  DB2 424
  Language Pack 244, 422
  Linux 420
  Storage Resource agent 421
  Storage Resource agent manually 421
uninstalling *(continued)*
  Tivoli Storage Productivity Center 417, 418
  Tivoli Storage Productivity Center agents 421
  Tivoli Storage Productivity Center for Replication 418, 420
    on Windows 420
  Tivoli Storage Productivity Center Monitoring Agent 423
  Tivoli Storage Productivity Center Monitoring Agent on UNIX 424
  Tivoli Storage Productivity Center Monitoring Agent on Windows 423
  with InstallShield 419
upgrading
  overview 147, 379
  Storage Resource agents manually 406, 408
  System Storage N Series Gateway servers 410
  Tivoli Storage Productivity Center for Replication 380, 383
  TotalStorage Productivity Center 3.x 382
user IDs
  valid characters 428
  work sheet 40, 435
user interface
  installing Tivoli Integrated Portal using custom installation 197

V

views
  planning to retrieve information 139, 140
Virtual I/O Server
  planning for 136
VMware
  capacity reports 135
VMware support 132
  configuration 135
volume managers
  supported 47
volume planning 65
VPA 65

W

Windows domain
  adding Tivoli Storage Productivity Center 344
  configuring Data server for 342
work sheets 40, 435
workload profile 65
WWPN 65

X

XIV system
  planning for 96

Z

zone 65
zone planning 65
zone set 65
zone sets
  configuring for inactive zones ets 338



Product Number: 5608-WB1, 5608-WC0, 5608-WC3, 5608-WC4, 5725-E18,
5725-A01
5608-WB2, 5608-WB3

Printed in USA

SC27-2337-06

