IBM Tivoli Storage Productivity Center

User's Guide

Version 4.1

SC27-2338-00

IBM Tivoli Storage Productivity Center

User's Guide

Version 4.1

SC27-2338-00

Note:
Before using this information and the product it supports, read the information in "Notices" on page 835.

This edition applies to version 4, release 1, modification 0 of IBM Tivoli Storage Productivity Center (product numbers 5608-WB1, 5608-WC0, 5608-WC3, and 5608-WC4) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

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 have already installed the product and are now using IBM Tivoli Storage Productivity Center. This book provides an overview of how to use the product and detailed information about a variety of topics.

Administrators and users should be familiar with the following topics:

- SAN concepts
- Tivoli Storage Productivity Center concepts (or plan to read the Concepts chapter)
- IBM Database 2 (DB2) Universal Database (UDB)
- Simple Network Management Protocol (SNMP) concepts
- IBM Tivoli Enterprise Console®

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/shop/publications/order.

IBM Tivoli Storage Productivity Center publications

Use these publications for information about how to install, configure, and use IBM Tivoli Storage Productivity Center.

The Tivoli Storage Productivity Center publications are available from the IBM Tivoli Storage Productivity Center Information Center at http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp. Click **Tivoli Storage Productivity Center**.

For PDF documents, click **IBM Tivoli Storage Productivity Center > Printable documentation**.

Publication Title	Order Number
IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide	SC27-2337
IBM Tivoli Storage Productivity Center User's Guide	SC27-2338
IBM Tivoli Storage Productivity Center Messages	SC27-2340
IBM Tivoli Storage Productivity Center Command-Line Interface Reference	SC27-2339
IBM Tivoli Storage Productivity Center Problem Determination Guide	GC27-2342

Publication Title	Order Number
IBM Tivoli Storage Productivity Center Workflow User's Guide	SC27-2341

IBM Tivoli Storage Productivity Center for Replication publications

Use these publications for information about how to install, configure, and use IBM Tivoli Storage Productivity Center for Replication.

The following table lists the IBM Tivoli Storage Productivity Center for Replication publications. These publications are available in the Information Center at http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp.

Click **Tivoli Storage Productivity Center for Replication > Reference > Publications**.

Information for installing, upgrading, and uninstalling IBM Tivoli Storage Productivity Center for Replication is documented in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide*.

Publication Title	Order Number
IBM Tivoli Storage Productivity Center for Replication for System z v4.1 Installation and Configuration Guide	SC27-2321-00
IBM Tivoli Storage Productivity Center for Replication V4.1 Command-Line Interface User's Guide	SC27-2323-00
IBM Tivoli Storage Productivity Center for Replication V4.1 Problem Determination Guide	GC27-2320-00
IBM Tivoli Storage Productivity Center for Replication V4.1 User's Guide	SC27-2322-00

IBM System Storage Productivity Center publications

Use these publications for information about how to install, configure, and use IBM System Storage™ Productivity Center.

These publications are available in the Information Center at http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp.

Click **System Storage Productivity Center**.

For PDF documents, click **System Storage Productivity Center > Printable documentation**.

Publication Title	Order Number
IBM System Storage Productivity Center Introduction and Planning Guide	SC23-8824
IBM System Storage Productivity Center Documentation CD	SCD7-1477
Read This First: Installing the IBM System Storage Productivity Center	GI11-8938
IBM System Storage Productivity Center User's Guide	SC27-2336

IBM System Storage DS3000 publications

Use these publications for information about how to install, configure, and use the IBM DS3000.

To see the IBM DS3000 publications, follow these steps:

1. Go to http://www.ibm.com/servers/storage/support/.
2. Under Select your product, in the Product Family field, select **Disk systems**.
3. In the Product field, select **DS3200**, **DS3300**, or **DS3400**, as applicable.
4. Click **Go**.
5. In the Support and Download pane, click **Documentation**.
6. Under Documentation, click on a document title.

Publication Title	Part or Order Number
IBM System Storage DS3000 Storage Manager Version 10 Installation and Support Guide for Windows, Linux, NetWare, and VMware	46M1364
IBM System Storage DS3200 Storage Subsystem Installation, Maintenance, and User's Guide	46M1361
IBM System Storage DS3300 Storage Subsystem Installation, Maintenance, and User's Guide	46M1362
IBM System Storage DS3400 Storage Subsystem Installation, Maintenance, and User's Guide	46M1363
IBM System Storage DS3000 Storage Manager Version 10 Installation and Support Guide for AIX, Linux on POWER, and Sun Solaris	46M1365

IBM System Storage DS4000 and DS5000 publications

Use these publications for information about how to install, configure, and use the IBM DS4000® and IBM DS5000.

To see the IBM DS4000 or IBM DS5000 publications, follow these steps:

1. Go to http://www.ibm.com/systems/support/storage/disk.
2. Under Select your product, in the Product Family field, click **Disk systems**.
3. In the Product field, click the appropriate storage system.
4. Click **Go**.
5. Click **Documentation**.
6. Click a document.

Publication Title	Order Number
IBM System StorageDS4000 Concepts Guide	GC26-7734
IBM System Storage DS4000/DS5000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview	GC53-1137
IBM System Storage DS® Storage Manager Version 10 IBM System Storage DS Storage Manager Installation and Host Support Guide	GC53-1135
IBM System Storage DS Storage Manager Version 10.50 Copy Services User's Guide	GC53-1136
IBM System Storage DS4800 Storage Subsystem Quick Start Guide	GC27-2148
DS5100, DS5300 and EXP5000 Quick Start Guide	GC53-1134

Publication Title	Order Number
IBM System Storage DS4800 Storage Subsystem Installation, User's, and Maintenance Guide	GC26-7845
IBM System Storage DS5100 and DS5300 Storage Subsystems Installation, User's, and Maintenance Guide	GC53-1140
IBM System Storage DS4000/DS5000 Hard Drive and Storage Expansion Enclosure Installation and Migration Guide	GC53-1139
IBM System Storage DS5000 EXP5000 Storage Expansion Enclosure Installation, User's, and Maintenance Guide	GC53-1141
IBM System Storage DS3000, DS4000, and DS5000 Command Line Interface and Script Commands Programming Guide	GC52-1275
IBM System Storage DS4000/DS5000EXP810 Storage Expansion Enclosure Installation, User's and Maintenance Guide	GC26-7798

IBM System Storage DS6000 publications

Use these publications for information about how to install, configure, and use the DS6000™.

These publications are available from the DS6000 Information Center on the following Web site:

http://publib.boulder.ibm.com/infocenter/dsichelp/ds6000ic/index.jsp

Publication Title	Order Number
IBM System Storage DS6000: Host Systems Attachment Guide **Note:** No hardcopy is produced for this publication.	GC26-7680
IBM System Storage DS6000: Introduction and Planning Guide	GC26-7679
IBM System Storage Multipath Subsystem Device Driver User's Guide	SC30-4096
IBM System Storage DS6000 : Messages Reference	GC26-7682
IBM System Storage DS6000 Installation, Troubleshooting, and Recovery Guide	GC26-7678
IBM System Storage DS6000 Quick Start Card	GC26-7659

IBM System Storage DS8000 publications

Use these publications for information about how to install, configure, and use the DS8000® system.

These publications are available from the DS8000 Information Center at http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp.

Publication Title	Order Number
IBM System Storage DS8000: Host Systems Attachment Guide **Note:** No hardcopy is produced for this publication.	SC26-7917
IBM System Storage DS8000: Introduction and Planning Guide	GC35-0515
IBM System Storage DS8000: Command-Line Interface User's Guide	GC53-1127
IBM System Storage DS8000: Messages Reference	GC26-7914

IBM System Storage DS Open Application Programming Interface publications

Use these publications for information about how to install, configure, and use the DS CIM agent.

These publications are available at http://www.ibm.com/servers/storage/support/software/cimdsoapi/.

Click the **Install** tab **> Documentation**. Make sure you reference the correct document for the CIM agent version.

Publication Title	Order Number
IBM System Storage DS Open Application Programming Interface 5.4.1 and 5.4.2 Installation and Reference	GC35-0516-04
IBM System Storage DS Open Application Programming Interface 5.3 Installation and Reference	GC35-0516-03
IBM System Storage DS Open Application Programming Interface Reference for CIM agent 5.2	GC35-0516-01
IBM Tivoli Storage Productivity Center DS Open Application Programming Interface Reference for CIM agent 5.1	GC35-0493

IBM System Storage SAN Volume Controller publications

Use these publications for information about how to install, configure, and use IBM System Storage SAN Volume Controller.

The following table lists the SAN Volume Controller publications. These publications are available in the SAN Volume Controller Information Center at http://publib.boulder.ibm.com/infocenter/svcic/v3r1m0/index.jsp.

Publication Title	Order Number
IBM System Storage SAN Volume Controller CIM Agent Developer's Guide	SC26-7904
IBM System Storage SAN Volume Controller Command-Line Interface User's Guide	SC26-7903
IBM System Storage SAN Volume Controller Software Installation and Configuration Guide	SC23-6628
IBM System Storage SAN Volume Controller Host Attachment Guide	SC26-7905
IBM System Storage SAN Volume Controller Planning Guide	GA32-0551
IBM System Storage SAN Volume Controller Troubleshooting Guide	GC27-2227
IBM System Storage SAN Volume Controller Hardware Maintenance Guide	GC27-2226
IBM System Storage SAN Volume Controller Model 2145-8G4 Hardware Installation Guide	GC27-2220
IBM System Storage SAN Volume Controller Model 2145-8A4 Hardware Installation Guide	GC27-2219
IBM System Storage SAN Volume Controller Model 2145-4F2 Hardware Installation Guide	GC27-2222
IBM System Storage SAN Volume Controller Models 2145-8F2 and 8F4 Hardware Installation Guide	GC27-2221

IBM DB2 Database for Linux, UNIX, and Windows publications

Use these publications for information about how to install, configure, and use DB2®.

The following table lists some of the IBM DB2 Database for Linux®, UNIX®, and Windows® product publications for Version 9.5.

For a complete list of DB2 publications, go to http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/index.jsp.

Publication Title	Order Number
IBM DB2 Version 9.5 for Linux, UNIX, and Windows, Getting started with DB2 installation and administration on Linux and Windows	GC23–5857
IBM DB2 Version 9.5 for Linux, UNIX, and Windows, Command Reference	SC23–5846
IBM DB2 Version 9.5 for Linux, UNIX, and Windows, Message Reference Volume 1	GI11–7855
IBM DB2 Version 9.5 for Linux, UNIX, and Windows, Message Reference Volume 2	GI11–7856
IBM DB2 Version 9.5 for Linux, UNIX, and Windows, Migration Guide	GC23–5859
IBM DB2 Version 9.5 for Linux, UNIX, and Windows, Troubleshooting Guide	GI11–7857

IBM XIV Storage System publications

For information about how to install, configure, and use the IBM XIV® Storage System, use the following link:

http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp

Related publications

Use these publications for information about related components that IBM Tivoli Storage Productivity Center interfaces with, for example, the storage area network.

The following table lists related IBM product publications.

Publication Title	Order Number
IBM Tivoli Storage Area Network Manager: A Practical Introduction	SG24-6848
Introduction to Storage Area Networks	SG24-5470
Designing an IBM Storage Area Network	SG24-5758
IBM Tivoli Storage Common Information Model Agent for the Enterprise Storage Server®: Installation and Configuration Guide	GC35-0485
IBM Tivoli Storage SAN File System Planning Guide	GA27-4344
IBM Tivoli Storage SAN File System Installation and Configuration Guide	GA27-4316
IBM Tivoli Storage SAN File System Administrator's Guide and Reference	GA27-4317
IBM Tivoli Storage SAN File System, System Management API Guide and Reference	GA27-4315
IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide	SC26-7650
IBM SAN Fibre Channel Switch: 2108 Model S08 User's Guide	SC26-7349

Publication Title	Order Number
IBM SAN Fibre Channel Switch: 2108 Model S16 Installation and Service Guide	SC26-7352
IBM SAN Fibre Channel Switch: 2108 Model S16 User's Guide	SC26-7351
Implementing Fibre Channel Attachment on the Tivoli Storage Enterprise Storage Server	SG24-6113

IBM International Technical Support Organization publications

The IBM International Technical Support Organization (ITSO) publishes IBM Redbooks®, which 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.

For information about IBM System Storage Productivity Center, see *IBM System Storage Productivity Center Deployment Guide*. Search for **SG24-7560**.

For information about IBM Tivoli Storage Productivity Center, see *TotalStorage® Productivity Center V3.3 Update Guide*. Search for **SG24-7490**.

Translations

Translated publications are available within the IBM Tivoli Storage Productivity Center Information Center. The IBM Tivoli Storage Productivity Center Information Center is available in certain translated languages, and is displayed in the language that is appropriate for the Web 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.

See the ″Printable documentation″ section of the information center for links to PDFs.

Contact your IBM Support Center for more information about the translated publications and whether these translations are available in your country.

Accessing publications online

This topic provides information on how to access the IBM Tivoli Storage Productivity Center Information Center.

You can access publications in the IBM Tivoli Storage Productivity Center Information Center at http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp.

The IBM Tivoli Storage Productivity Center Information Center contains the most recent version of the books in the product library in PDF or HTML formats, or both. Translated documents are also available for some products.

Note: If you print PDF documents on other than letter-sized paper, select the **Fit to page** check box in the **Adobe Acrobat Print** dialog. This option is available when

you click **File** → **Print**. **Fit to page** ensures that the full dimensions of a letter-sized page print on the paper that you are using.

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/shop/publications/order.

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

Providing feedback about publications

This topic provides information on where to send feedback about the publications.

If you have comments or suggestions about the product and documentation, complete the customer feedback survey at http://www.ibm.com/systems/support/storage/software/tpc.

On the left side of the Web page, click **Feedback**.

Contacting IBM Support Center

This topic provides information on how to contact IBM Support Center for information.

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

- Go to the IBM Tivoli Storage Productivity Center technical support Web site at http://www.ibm.com/systems/support/storage/software/tpc/.

  To receive future support notifications, go to the right and under **Stay informed**, click **Subscribe**. You will be required to enter your IBM ID and password. Once authenticated, you will be able to configure your subscription for Tivoli Storage Productivity Center technical support Web site updates.
- Customers in the United States can call 1-800-IBM-SERV (1-800-426-7378).
- International customers should go to the Tivoli Storage Productivity Center technical support Web site for customer support telephone numbers.

You can also review the *IBM Software Support Handbook*, which is available on our Web site at http://techsupport.services.ibm.com/guides/handbook.html.

The support Web site 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 Web site. See “Reporting a problem” on page xv.

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 on 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

- Lowercase 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:

- < > (less than, greater than symbols) are used to indicate a variable value. Do not type the < > symbols.
- # is the prompt for the root user on UNIX platforms.

- Uppercase and lowercase characters do matter. Type in commands exactly as shown.

New for IBM Tivoli Storage Productivity Center Version 4.1

Use this information to learn about new features and enhancements in IBM Tivoli Storage Productivity Center version 4.1. This section highlights the changes since IBM TotalStorage Productivity Center 3.3.2.

For more information about each of the features, go to the Tivoli Storage Productivity Center Information Center and search for **Planning for the IBM Tivoli Storage Productivity Center family**. For information about how to use the features, see the *IBM Tivoli Storage Productivity Center User's Guide*.

Tivoli Storage Productivity Center 4.1 adds the following new features, functions, and enhancements:

Name change

IBM Tivoli Storage Productivity Center V4.1 has been renamed from IBM TotalStorage Productivity Center. All user interfaces, documentation, online help, and messages have also been changed to reflect the name change.

Licensing changes

These are the licenses available for IBM Tivoli Storage Productivity Center:

- IBM Tivoli Storage Productivity Center Basic Edition
- IBM Tivoli Storage Productivity Center Standard Edition
- IBM Tivoli Storage Productivity Center for Disk
- IBM Tivoli Storage Productivity Center for Data

If you have an IBM TotalStorage Productivity Center for Fabric license only, you can upgrade to IBM Tivoli Storage Productivity Center Standard Edition.

If you have an IBM TotalStorage Productivity Center for Basic Edition license only, you can upgrade to IBM Tivoli Storage Productivity Center Basic Edition, IBM Tivoli Storage Productivity Center for Disk, IBM Tivoli Storage Productivity Center for Data, or IBM Tivoli Storage Productivity Center Standard Edition.

If you have an IBM TotalStorage Productivity Center for Data license only, you can upgrade to IBM Tivoli Storage Productivity Center for Data or IBM Tivoli Storage Productivity Center Standard Edition.

If you have an IBM TotalStorage Productivity Center for Disk license only, you can upgrade to IBM Tivoli Storage Productivity Center for Disk, IBM Tivoli Storage Productivity Center for Data (Disk plus Data), or IBM Tivoli Storage Productivity Center Standard Edition.

If you have an IBM TotalStorage Productivity Center Standard Edition license, you can upgrade to IBM Tivoli Storage Productivity Center Standard Edition.

Integration features

Tivoli Storage Productivity Center provides these integration features.

Integration of Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication

Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication, previously separated products,

are now integrated. You can start the IBM Tivoli Storage Productivity Center for Replication user interface from within the Tivoli Storage Productivity Center user interface.

The *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* also includes the installation, upgrade, and uninstallation information for IBM Tivoli Storage Productivity Center for Replication.

This integration enables you to:

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

Integration of Tivoli Storage Productivity Center and IBM Tivoli Integrated Portal

Tivoli Integrated Portal is a standards-based architecture for Web administration. Tivoli Integrated Portal enables developers to build administrative interfaces for IBM and independent software products as individual plug-ins to a common console network. The installation of Tivoli Integrated Portal is required to enable single sign-on for Tivoli Storage Productivity Center.

Single sign-on is an authentication process that enables you to enter one user ID and password to access multiple applications. Single sign-on integrates with the launch in context feature to enable you to move smoothly from one application to a specific location in a second application.

Launch in context feature

The launch in context feature enables you to access external applications from the Tivoli Storage Productivity Center GUI. Element managers are the most prevalent external applications that use the launch in context feature. An element manager is usually the vendor-specific software that is used to administer a particular storage device. The launch in context feature provides starting points in the Tivoli Storage Productivity Center GUI so you can click a button or select a menu item to start an element manager.

When you install Tivoli Storage Productivity Center, Tivoli Integrated Portal, and Tivoli Storage Productivity Center for Replication, the components are automatically configured to use launch in context. You can access Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication from

the Tivoli Integrated Portal GUI and you can access Tivoli Storage Productivity Center for Replication from the Tivoli Storage Productivity Center GUI.

There are three levels of launch in context ability:

Simple launch
This level exists in TotalStorage Productivity Center 3.3.2. Tivoli Storage Productivity Center discovers basic information about the device and the management of the device.

Launch with parameters
You can specify additional parameters in the URL or command-line interface when starting an application. The parameters that are passed enable you to navigate to a particular panel or state of the application that was started. You can also identify objects to operate on and possibly provide values to use in the operation.

Launch with single sign-on

You can enhance the launch in context feature to include single sign-on. Single sign-on can be used when an external application can perform authentication against the same user repository as Tivoli Storage Productivity Center. A directory that is Lightweight Directory Access Protocol (LDAP) compliant is a common example of such a user repository.

External applications that do not include the WebSphere Application Server (WAS), require the authentication service that is provided by Tivoli Integrated Portal. For example, the element manager for IBM System Storage DS8000, DS8000 Storage Manager, uses the authentication service to handle launch in context with single sign-on from the Tivoli Storage Productivity Center GUI.

Single sign-on

Single sign-on is an authentication process that enables you to enter one user ID and password to access multiple applications. Single sign-on enables you to access:

- Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication from the Tivoli Integrated Portal GUI.
- Tivoli Storage Productivity Center for Replication from the Tivoli Storage Productivity Center GUI.
- External applications such as element managers from the Tivoli Storage Productivity Center GUI.

The single sign-on feature requires a centralized user and group repository, such as an LDAP-compliant directory, that all participating applications can access.

Tivoli Storage Productivity Center uses Lightweight Third Party Authentication (LTPA) tokens to pass 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 a determined amount of time. When the tokens expire, you must re-enter your user ID and password information.

If you select operating system authentication, then the use of the single sign-on feature is limited. Operating system authentication does not support single sign-on for element managers, even when the element manager is installed on the same machine as Tivoli Storage Productivity Center.

Storage Resource agents

Tivoli Storage Productivity Center now supports Storage Resource agents on Microsoft® Windows, AIX®, and Linux. The Storage Resource agent probe is equivalent to the information that is collected by probes using the Data agent.

The Storage Resource agents do not require the Agent Manager and can be deployed to other systems using the Tivoli Storage Productivity Center GUI on the server system.

You can use the following functions:

- Asset reports (including HBA)
- Capacity reports
- Subsystem to host storage correlation including multipathing information
- Topology and Data Path explorer functions

This support does not include file system scans, NAS discovery or topology, zoning and zone control functions or subsystem device driver configuration. You can still use the Data agent and Fabric agent for this information.

SQL access to Tivoli Storage Productivity Center database

Tivoli Storage Productivity Center will provide a set of DB2 views that represent key information that has been collected by monitoring jobs and stored in the database repository. A *view* is a way of describing data that exists in one or more tables within the database repository. It does not contain data but, instead, is a stored set of SQL commands that define a subset of rows and columns in the base tables.

You can use the Structured Query Language (SQL) to retrieve the information from the views and create reports using your own tools, such as Business Intelligence and Reporting Tools (BIRT) or Microsoft Excel. Other applications can also use these views to gather and import information that is collected by Tivoli Storage Productivity Center.

The following categories of views will contain information collected by Tivoli Storage Productivity Center:

Storage entity views

These views include information about the properties of the entity. For example, the name, capacity, freespace, and so forth for a storage subsystem.

Entities defined by Tivoli Storage Productivity Center

These entities include Data agents, Fabric agents, alert log, Tivoli Storage Productivity Center server, computer groups, storage subsystem groups, file system groups, storage resource groups, and so forth.

Aggregated views
These views provide summary information for the database history, data in a database instance, and the Data agent file system.

Reporting views
These views combine several different entities in one view for a report.

Rollup views
These views include rollup report information from the master and subordinate Tivoli Storage Productivity Center servers, Data agents and Fabric agents, host cluster data, computer group, host, database computer groups, fabric SAN assets, switch assets, storage subsystem group, storage subsystems, and Tivoli Storage Productivity Center for Databases.

Storage Optimizer
The Storage Optimizer is a tool to help you analyze your storage networks to identify hot spots or bottlenecks, plan for storage growth, improve performance, and help develop storage migration or storage consolidation plans. Using the data in the Tivoli Storage Productivity Center database, the Storage Optimizer enables you to create an analysis report and an optimization report. The analysis report analyzes your data storage environment and recommends changes to improve your environment. Based on the analysis report, the optimization report includes storage migration or storage consolidation recommendations.

This feature requires an IBM Tivoli Storage Productivity Center Standard Edition license.

Storage resource groups
Storage resource groups are new objects provided to help storage administrators plan, monitor, and report on the managed environment.

A storage resource group 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. Storage resource groups can be a group of heterogeneous objects and can also contain other storage resource groups without any connectivity.

Policies for provisioning (volume creation and selection, workload profiles, zoning and multipathing configuration) can be specified and associated with storage resource groups. These policies are used by the SAN Planner to populate default settings.

Storage resource groups are used primarily for planning functions but is also available with the Tivoli Storage Productivity Center Basic Edition license. With the basic license, you can create and view storage resource groups in the topology. With the Standard Edition license, the planner function is enabled and you can use storage resource groups as input.

Storage resource groups also work with these profiles:

Workload profiles
Describes the requirements that define the performance characteristics of newly provisioned capacity.

Provisioning profiles
Describes the requirements such as total capacity, number of

volumes, Redundant Array of Independent Disks (RAID) level, volume name prefix, multipathing options, zoning options, and so forth.

IBM General Parallel File System™

Tivoli Storage Productivity Center supports the monitoring of the IBM General Parallel File System (GPFS™) 3.2 on AIX. GPFS provides access to critical file data. GPFS also provides concurrent high-speed file access to applications that are running on multiple nodes of an AIX cluster, a Linux cluster, or a heterogeneous cluster of AIX and Linux nodes. In addition to providing file storage capabilities, GPFS provides storage management, information life-cycle tools, centralized administration and allows for shared access to file systems from remote GPFS clusters.

Installation changes

IBM Tivoli Storage Productivity Center for Replication

The *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* also includes the installation, upgrade, and uninstallation information for IBM Tivoli Storage Productivity Center for Replication. IBM Tivoli Storage Productivity Center for Replication is now installed with IBM Tivoli Storage Productivity Center.

IBM DB2 Database for Linux, UNIX, and Windows

Tivoli Storage Productivity Center now supports DB2 9.5. You will be able to migrate your Tivoli Storage Productivity Center databases from DB2 9.1 or DB2 8.2 to DB2 9.5. DB2 9.5 is optional. Tivoli Storage Productivity Center still supports DB2 9.1.

Installation of IBM Tivoli Integrated Portal

Tivoli Storage Productivity Center now installs IBM Tivoli Integrated Portal along with Tivoli Storage Productivity Center.

Embedded WebSphere® 6.1 and JRE 1.5

The Device server is upgraded to run under Embedded WebSphere 6.1 (from Embedded WebSphere 6.0.2). The Data server, GUI, and CLI is upgraded to use JRE version 1.5. InstallShield uses JRE 1.5 during the installation and uninstallation process when Tivoli Storage Productivity Center is installed using the disk1 image. The image to perform local agent installations uses JRE version 1.4.2.

Silent installation

Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication do not support silent installation except for the Data agents and Fabric agents.

New device and application support

IBM System Storage DS8000 4.2

This release supports DS8000 4.2 with these additional items:

- Storage pool striping
- Dynamic volume expansion
- Internet Protocol Version 6 (IPv6)
- Redundant Array of Independent Disks (RAID 6)
- Variable logically partitioned mode (LPARs)
- Space-efficient FlashCopy®

IBM System Storage SAN Volume Controller 4.3.1

This release supports SAN Volume Controller 4.3.1 with these additional items:

- Embedded CIM agent
- 64-bit logical block address (LBA) for the back end array
- 2 TB virtual disks (VDisks) and managed disks (MDisks)

Microsoft SQL Server 2005 and Microsoft SQL Server 2008 databases

Tivoli Storage Productivity Center can now monitor the Microsoft SQL Server 2005 and Microsoft SQL Server 2008 databases. You must configure Microsoft SQL Server before you can monitor the database. For information about configuration, see the Information Center. Search for **Configuring Microsoft SQL Server 2005 or 2008**.

EMC PowerPath

With Tivoli Storage Productivity Center, you can now use EMC PowerPath storage systems like CLARiiON and Symmetrix. Using these storage systems, you can discover host volume information and display detailed information for the volume for capacity planning purposes. Connection reports can show the connectivity from the host to the storage subsystems.

EMC PowerPath version 4.0 or later is supported.

Network Appliance (NetApp)

With Tivoli Storage Productivity Center, you can use the Network Appliance SMI-S agent to support block storage devices. The SMI-S agent supports the SMI-S 1.2 array profile.

IBM XIV Storage System

Important: The XIV Storage System information provided in the Tivoli Storage Productivity Center 4.1 documentation is only for planning purposes until the supported XIV Storage System software is available. Tivoli Storage Productivity Center support is targeted for a futureXIV Storage System software release. A flash will be issued when Tivoli Storage Productivity Center support for XIV Storage System is available.

XIV Storage System will have an embedded CIM agent that Tivoli Storage Productivity Center will use to run discovery and probe jobs.
You will be able to start the XIV Storage System GUI from within Tivoli Storage Productivity Center if the GUI is installed on the same system as the Tivoli Storage Productivity Center GUI. The XIV Storage System GUI will be supported on Windows and Linux.
Both the Data agent and Storage Resource agent will support XIV Storage System.

Multipath subsystem device drivers

Tivoli Storage Productivity Center supports these subsystem device drivers (SDD):

- AIX SDD
- Windows SDD
- Windows SDD DSM
- Linux SDD

- HP SDD
- Solaris SDD
- Novell SDD (reporting only)
- AIX SDD PCM
- Linux DM_Multipath

IBM System Storage N Series Gateway servers
IBM Tivoli Storage Productivity Center supports IBM System Storage N Series Gateway servers as **Other NAS**. This support allows you to monitor and report on file systems through the Windows CIFS or UNIX NFS shares that are accessible to the scan or probe jobs for the Data agent. No backend storage information such as controllers, disks, and logical volumes is collected or reported.

High-Availability Cluster Multi-Processing
This release provides additional support for High-Availability Cluster Multi-Processing version 5.5.

Tivoli Enterprise Portal
A Universal Agent for Tivoli Storage Productivity Center that utilizes a set of Tivoli Storage Productivity Center Web services calls to gather information and provide results files that will display enhanced information such as job status and Tivoli Storage Productivity Center status in the IBM Tivoli Integrated Portal.

Terminology
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.

Chapter 1. Getting started (roadmap)

Use this roadmap as a getting started guide for installing and using IBM Tivoli Storage Productivity Center. This topic is organized into a number of sequential tasks that help you learn how the product works and how to get started using it to administer and manage your storage environment.

Complete the following tasks as shown below and click on the included links to learn more about the capabilities of Tivoli Storage Productivity Center. Follow the tasks as outlined: Planning, Installing, Configuring, Administering, managing storage resources, and reporting. Troubleshooting contains procedures for debugging any problems you might have.

Task	Description	Additional information
Determine prerequisites and plan for installation.	The first installation task is planning for Tivoli Storage Productivity Center. Planning is very important. The time you invest in planning directly affects your ease of implementation. Planning includes: • Ensuring your systems meet hardware and software requirements • Determining which agents to install and where to install them • Ensuring ports are open for use by Tivoli Storage Productivity Center components • Granting privileges to the user ID used to install the Tivoli Storage Productivity Center components • Understanding Agent Manager requirements • Setting up for LDAP Support	**When and how often should I do this?**: Perform this task before you install Tivoli Storage Productivity Center.
Install Tivoli Storage Productivity Center.	The second task is installing Tivoli Storage Productivity Center. You can install all the Tivoli Storage Productivity Center components using typical installation or custom installation. It is recommended that for a production environment you use custom installation. Tivoli Storage Productivity Center has the following installable components: • Database schema • Data server • Device server • graphical user interface (GUI) • command line interface (CLI) • Data agent • Fabric agent • IBM Tivoli Integrated Portal	**When and how often should I do this?**: Perform this task after you have completed planning and ensuring that you meet the installation prerequisites. You can also install components separately depending on your needs, so you might perform a portion of the installation process a number of times after the initial installation. For example, you can install agents at any time on the machines whose storage you want to monitor.

Task	Description	Additional information
Start and Configure Tivoli Storage Productivity Center.	Perform this task after installing Tivoli Storage Productivity Center to configure it according to the standards within your organization, including: • Starting the Tivoli Storage Productivity Center GUI. • Running the Configuration Utility as a guide to help you through the steps required to configure your system. • Use the items in the Help menu to launch the product information center and learn about the appropriate levels of software, firmware, and hardware that will be needed to run the product. • Assigning roles to individuals who will use the product. • Entering licensing information • Setting SNMP and e-mail alerts • Entering retention settings for log files, resource history, removed resources • Defining rules for generating e-mail addresses of quota violators • Scheduling history aggregation and data agent upgrades • Setting up Novell NetWare and NAS filer access • Setting up IBM Tivoli Storage Enterprise Storage Server or SAN Volume Controller access • Using the configuration history feature to capture and analyze historical data that identifies possible problems with a storage area network (SAN) configuration (for system administrators only) • Using the configuration analysis feature to determine whether an existing SAN configuration complies with predefined best practices (for system administrators only)	**When and how often should I do this?**: Typically, you perform this step immediately after installing the product. Once Tivoli Storage Productivity Center is configured, you do not have to perform this step again unless changes within your organization require it.

Task	Description	Additional information
Administer DB2 and Tivoli Storage Productivity Center.	Once Tivoli Storage Productivity Center is installed and configured, start the necessary services for using the database and the GUI. Information in this section helps you perform the following tasks: • Authorizing users • Starting and stopping the console • Viewing and managing services • Managing Agent Manager • Administering data sources (agents): CIMOM, Data, Native, Inband Fabric, Out of Band Fabric, Tivoli Storage Productivity Center Servers, VMWare VI • Starting and stopping Tivoli Storage Productivity Center services • Increasing the memory allocation and checking the status of Tivoli Storage Productivity Center components • Setting timeout values for the Device server • Changing passwords for Tivoli Storage Productivity Center components • Configuring the launch of the IBM Tivoli Storage Productivity Center for Replication GUI and other web GUIs or local executables, • Using the Repository Copy tool • Collecting information with the service tool	**When and how often should I do this?**: Perform administrative tasks at any time during the use of Tivoli Storage Productivity Center. These tasks are meant to provide you with the tools to maintain a successful implementation of the product during its lifecycle.
Discover storage resources.	Perform discovery to have Tivoli Storage Productivity Center detect the storage resources within your environment, including detection of network topology. Discovery jobs collect basic information about the storage resources in your environment, including computers, CIMOMs, storage subsystems managed by CIMOMs, fabrics, NetWare trees, filers, clusters, and tape libraries.	**When and how often should I do this?**: You must run discovery jobs before you can collect more detailed information about storage resources. Once discovery has run, you can use monitoring jobs such as probes, scans, and pings to collect detailed information about the discovered resources for use in reports. See Discover resources in your system for more information.

Task	Description	Additional information
Manage storage resources: define monitoring groups to represent the storage entities within your organization.	Organizing the users, computers, storage subsystems, filesystems, and directories within your organization into groups enables you to quickly target specific entities against which to collect information, set quotas and constraints, and view reports. Use the Monitoring > Groups facility within each of Tivoli Storage Productivity Center managers to define a group that corresponds to that manager. For example, use Disk Managerr to create storage subsystem monitoring groups; use Fabric Manager to create fabric groups; use Data Manager to create computer, filesystem, directory, user, and OS user monitoring groups. After creating groups, you can include them in data collection jobs.	**When and how often should I do this?**: Create monitoring groups before you: • define monitoring jobs (probes, scans, and pings) to collect information about your storage resources • define Data Manager quotas and constraints to enforce your storage policies Creating groups is not required, but they are meant to help you better organize your storage resources for monitoring in a large environment. See the following topics for more information: 1. "Data collection groups" on page 233 2. "Working with groups of resources" on page 236

Task	Description	Additional information
Manage storage resources: determine what information you want to gather about the storage resources in your environment and define the appropriate monitoring jobs.	Determining what information you want to gather about the storage resources within your enterprise is critical to helping you implement a storage management strategy. Once you have determined the types of information you want to collect about your storage resources, you can use different monitoring jobs to collect information about those resources. The monitoring jobs that you can run include: • Probes: collect detailed statistics about the storage assets in your enterprise, such as computers, storage subsystems, fabrics, volumes, disk controllers, hard disks, and file systems. • Scans: collect statistics about the usage and trending of your storage consumption. • Pings: collect information about the availability of the storage assets in your enterprise. • Storage Subsystem Performance Monitors: collect statistics about the performance of storage subsystems. • Fabric Performance Monitors: collect statistics about the performance of fabrics. • Tivoli Storage Productivity Center Server Probes: In an environment where multiple Tivoli Storage Productivity Center servers are deployed, use Tivoli Storage Productivity Center Server probes to collect the information that has been gathered by those servers and view that information through a single interface.	**When and how often should I do this?**: Define monitoring jobs after determining the types of information you want to collect about your storage resources. You must define and run monitoring jobs before you can perform the other tasks within Tivoli Storage Productivity Center, such as viewing reports, enforcing storage policies, setting alerts, using the topology viewer, and managing storage subsystems. Once you define a job to collect specific information about your storage, you do not need to redefine that job in the future (unless you want to make changes to it such as adding new computers or groups)-defining a monitoring job is a one-time only task. Once defined, the job will run according the schedule that you define. See the following topics for more information: • Collect data from system resources • "Collecting storage statistics (probes)" on page 203 • "Collecting storage usage and trending information (scans)" on page 210 • "Determining the availability of storage resources (pings)" on page 230 • "Collecting performance data" on page 241 • "Managing fabrics" on page 430 • "Collecting storage statistics from IBM Tivoli Storage Productivity Center servers (TPC server probes)" on page 207

Task	Description	Additional information
Manage storage resources: determine what conditions must occur before you are alerted to a problem within your storage environment.	Determining when and how you will be alerted to conditions within your storage environment is important to helping you maintain and administer your resources. You can define alerts that will notify you when certain conditions are detected on a computer, filesystem, directory, storage subsystem, fabric, switch, and endpoint device. When a condition is met and alert is triggered, you can choose to be notified through an SNMP trap, IBM Tivoli Enterprise Console event, login notification, event log (Windows®), syslog (Unix), or e-mail. If you need a way to find and fix problems—or potential problems—as alerts are triggered, you can use the triggered action facility associated with alerts to automatically address these problems.	**When and how often should I do this?** Define alerts: • After defining data collection (monitoring) jobs • Based on the information gathered by data collection jobs Alerts are checked whenever a probe or scan job is run against your storage resources. Computer, fabric, switch, and endpoint alerts are triggered by probes. Directory and storage subsystem alerts are triggered by scans. Filesystem freespace alerts are triggered by both probes and scans. All other filesystem alerts are triggered only by probes. Once you define an alert, you do not need to redefine it in the future (unless you want to make changes to it)—defining an alert is a one-time only task. See “Working with alerts” on page 247 for more information.

Task	Description	Additional information
Manage storage resources: determine the storage policies you want to enforce.	Enforcing storage policies within your organization can be critical to ensuring the proper and most cost-effective use of your storage devices. Use: • Quotas to define limits on the amount of storage that a user or a group of users can consume • Network Appliance Quotas to import, view, and report on the NetApp Quotas defined on NAS filers • Constraints to define the acceptable and unacceptable file types, file sizes, and file owners for a computer or a set of computers in your environment • Filesystem extensions to automatically extend filesystems when utilization reaches a specified threshold • Scheduled actions to run scripts against storage devices on a schedule that you define When a condition within a quota or constraint is met, you can choose to be notified through a SNMP trap, login notification, event log (Windows), syslog (Unix), or e-mail.	**When and how often should I do this?**: Schedule the jobs to enforce your organization's storage policies after running the scan jobs to collect the information that you want checked. • Quotas: Once defined, a quota will run according to the schedule that you specified. Because quotas look at the information gathered by scans to determine quota violations, you should schedule quotas to run after the appropriate scan jobs. For example, if you want to restrict the amount of space consumed by users on specific computers, run a scan on those computers to collect usage information before running the quota to enforce that usage restriction. • Constraints: constraints are automatically checked whenever a scan is run against the storage resources to which the constraint applies. • Filesystem extensions: run filesystem extensions after probes of the related storage assets and filesystems. • Scheduled actions: run scheduled actions: immediately, once at a specified time, or repeatedly according to a schedule you specify. See the following topics for more information: 1. Creating and enforcing storage policies 2. "Setting storage usage quotas" on page 302 3. "Setting file constraints" on page 327 4. "Using filesystem extension" on page 336 5. "Scheduling script-based actions" on page 350

Task	Description	Additional information
Manage storage resources: managing storage subsystems.	Manage the storage systems discovered and monitored by Tivoli Storage Productivity Center in the following ways: • Define storage resource groups to logically group related storage entities such as hosts, storage subsystems, fabrics, and switches. • Use the SAN Planner, workload profiles, provisioning profiles, and storage resource groups to most effectively provision IBM Enterprise Storage Server subsystems. • Define element managers for storage subsystems • Manage virtual disks: assign host ports, create virtual disks, delete virtual disks, add a managed disk to a managed-disk group, view virtual-disk information • Manage volumes: assign host ports, create volumes, delete volumes, view information about volumes	**When and how often should I do this?**: You can manage storage subsystems at any time after collecting information about them through discovery and monitoring jobs. See the following topics for more information: 1. "Working with workload profiles" on page 406 2. 3. "Managing storage subsystems" on page 411 4. "Managing virtual disks" on page 417 5. "Managing volumes" on page 424
Manage storage resources: charging for storage usage.	Use the storage usage information gathered by monitoring jobs to generate invoices that charge back for storage usage.	**When and how often should I do this?**: Generate invoices whenever you want to charge users or groups of users in your organization for their storage usage. See "Working with invoices" on page 454 for more information.
View information about storage resources: using the topology viewer.	Use the Topology viewer to see an overall, high-level view of your storage environment. The topology viewer provides a central location for viewing your storage assets, and for monitoring, troubleshooting, and performing storage management tasks.	**When and how often should I do this?** : You can use the topology viewer at any time after running monitoring jobs to collect the information about your storage environment. See "Using the topology viewer" on page 368 for more information.
View information about storage resources using any of the following report types: predefined, user-defined, batch, and rollup reports.	Use predefined (system reports), user-defined, batch, and rollup reports to view the information collected by Tivoli Storage Productivity Center data collection and quota jobs. Consider the following steps when determining how to best use Tivoli Storage Productivity Center reports: 1. Determine how you want to generate and view reports by selecting a reporting type that best meets your needs. 2. Identify which reports contain the most important information about your storage environment. 3. Run the appropriate data collection jobs and generate the reports you want to view.	**When and how often should I do this?** : Perform this task after the appropriate data collection jobs have run and you want to view information about your storage resources. Generate any of the many available reports that show resource activity. See the following topics for more information: 1. "Choosing a reporting type" on page 464 2. "Choosing a report category" on page 466 3. "What can I find out using reports?" on page 472 4. "Working with reports" on page 477

Task	Description	Additional information
Troubleshooting	Use troubleshooting tips to resolve any problems you might have with Tivoli Storage Productivity Center.	**When and how often should I do this?**: Perform troubleshooting tasks when you encounter problems while working with the product. Troubleshooting information is meant to help you with resolving issues that might occur during the use of the product.

Chapter 2. Product Overview

This chapter contains general and conceptual information about IBM Tivoli Storage Productivity Center.

Introducing IBM Tivoli Storage Productivity Center

IBM Tivoli Storage Productivity Center provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance and assets in your enterprise environment, including NetWare, NAS, SSA, Tivoli Storage Enterprise Storage Server, and Microsoft Cluster Server technologies, as well as RDBMSs such as Oracle, Sybase SQL Server, Microsoft SQL Server, and DB2 UDB. Tivoli Storage Productivity Center provides storage management from the host and application to the target storage device. It provides disk and tape subsystem configuration and management, performance management, SAN fabric management and configuration, and usage reporting and monitoring.

Tivoli Storage Productivity Center can help you identify, evaluate, control and predict your enterprise storage management assets. Because it is policy-based, it can detect potential problems and automatically make adjustments based on the policies and actions that you define. For example, it can notify you when your system is running out of disk space or warn you of impending storage hardware failure. By alerting you to these and other issues related to your stored data, it enables you to prevent unnecessary system and application downtime.

Tivoli Storage Productivity Center:

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

Architecture

The IBM Tivoli Storage Productivity Center consists of several key components. This topic identifies these components and shows how they are related.

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

Device server
This component discovers, gathers information from, analyzes performance

of, and controls storage subsystems and SAN fabrics. It coordinates communication with and data collection from agents that scan SAN fabrics and storage devices.

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

Agents
Data agents, Storage Resource agents, and fabric agents gather host, application, and SAN fabric information and send this information to the Data Server or Device server.

GUI The graphical user interface (GUI) lets you enter or receive information for all Tivoli Storage Productivity Center components.

CLI The command-line interface (CLI) lets you issue commands for major Tivoli Storage Productivity Center functions.

Tivoli Integrated Portal
The IBM Tivoli Storage Productivity Center installation program includes IBM Tivoli Integrated Portal. Tivoli Integrated Portal is a standards-based architecture for Web administration. Tivoli Integrated Portal enables developers to build administrative interfaces for IBM and independent software products as individual plug-ins to a common console network. The installation of Tivoli Integrated Portal is required to enable single sign-on for Tivoli Storage Productivity Center. *Single sign-on* is an authentication process that enables you to enter one user ID and password to access multiple applications. Single sign-on integrates with the launch in context feature to enable you to move smoothly from one application to a functionally-related location in a second application.

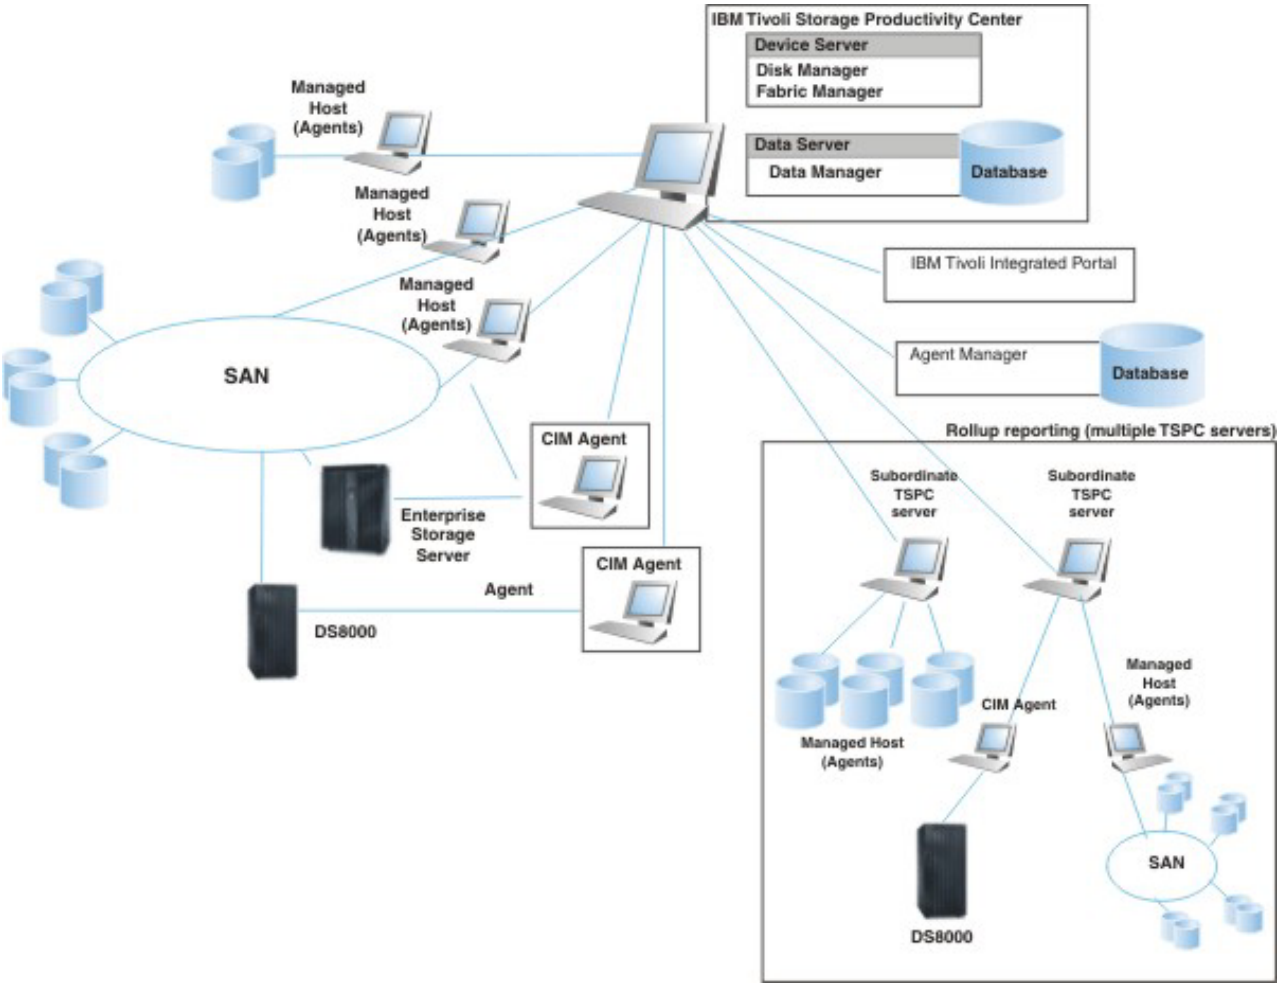

Figure 1. Tivoli Storage Productivity Center components

IBM Tivoli Storage Productivity Center summary

This topic provides a summary of the functionality provided by IBM Tivoli Storage Productivity Center.

The Tivoli Storage Productivity Center functions are shown in Table 1.

Table 1. Tivoli Storage Productivity Center summary of functions

Data management	Disk management	Fabric management	Tape management
Host-centric: • Discovery • Monitoring • Filesystem extension • Enterprise- wide reporting Application-centric: • Monitor DB2, Oracle, SQL Server, Sybase • Discovery • Monitoring • Chargeback	For storage subsystems: • Discovery • Monitoring • Configuration (for example, creating volumes) • Performance management	For fabrics: • Discovery • Monitoring • Configuration (for example, zoning) • Performance management	For tape libraries: • Discovery • Monitoring

IBM Tivoli Storage Productivity Center licenses

The following product licenses are available for IBM Tivoli Storage Productivity Center: Basic, Data, Disk, and Standard. 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

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 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 2. Available functions in each Tivoli Storage Productivity Center license

Function	Basic Edition	Disk Edition	Data Edition	Standard Edition
Administrative Services				
Device discovery	Yes, but not all[1]	Yes, but not all[1]	Yes	Yes

Table 2. Available functions in each Tivoli Storage Productivity Center license (continued)

Function	Basic Edition	Disk Edition	Data Edition	Standard Edition
Data Sources	Yes, but not all[2]	Yes, but not all[2]	Yes	Yes
Configuration	Yes, but not all[3]	Yes, but not all[3]	Yes	Yes
IBM Tivoli Storage Productivity Center				
Configuration Utility	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 2 below for a list of available reports.	Yes Not all reports are available in the Disk edition. See table 2 below for a list of available reports.	Yes Not all reports are available in the Data edition. See table 2 below for a list of available reports.	Yes
Topology Viewer	Yes	Yes	Yes	Yes
Monitoring • Probes • TPC Server Probes	 Yes[4] No	 Yes[4] No	 Yes Yes	 Yes Yes
Storage Resource Group Management	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	 Yes Yes Yes Yes
Alerting • Authentication Configuration Alerts • Alert Log	 Yes Yes, but not all[5]	 Yes Yes, but not all[5]	 Yes Yes	 Yes Yes
Data Manager				
Monitoring • Groups • Pings • Scans • Profiles	 Yes Yes No No	 Yes Yes No No	 Yes Yes Yes Yes	 Yes Yes Yes Yes
Alerting	Yes	Yes	Yes	Yes
Policy Management	No	No	Yes	Yes
Reporting	Yes Not all reports are available in the Basic edition. See table 2 below for a list of available reports.	Yes Not all reports are available in the Disk edition. See table 2 below for a list of available reports.	Yes	Yes
Data Manager for Databases	No	No	Yes	Yes

Table 2. Available functions in each Tivoli Storage Productivity Center license (continued)

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

Table 2. Available functions in each Tivoli Storage Productivity Center license (continued)

Function	Basic Edition	Disk Edition	Data Edition	Standard Edition
Notes: 1. You cannot run the following discovery jobs in the Basic and Disk editions: Netware Filer, Windows Domain, NAS, SAN FS, and VMware VI Data Source. 2. You cannot associate VMware VI data sources with Tivoli Storage Productivity Center in the Basic and Disk editions. 3. The following nodes under **Configuration** are not available in the Basic and Disk editions: • License Keys • Quota and Constraint e-mail Address Rules • Scan/Probe Agent Administration • Manual NAS/Netware Server Entry • Resource History Retention for Databases • Removed Resource Retention for Databases 4. Use probes and Storage Resource agents to collect asset and capacity information about storage entities in the Disk and Basic 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 and Disk editions. 5. The following nodes under **Alert Log** are not available in the Basic and Disk editions: • Directory • User • OS User Group • Hypervisor				

Table 3. Available reports in each IBM Tivoli Storage Productivity Center license

Reports	Basic Edition	Disk Edition	Data Edition	Standard Edition
Batch reports	No	Yes[1, 2]	Yes[2]	Yes
System Reports: • Data • Fabric • Disk	 Yes Yes, but not all [3] No	 Yes Yes, but not all [3] Yes	 Yes Yes, but not all [3] No	 Yes Yes Yes
Asset reports: • By Cluster • By Computer • By Hypervisor • By OS Type • By Storage Subsystem • System-wide	 Yes Yes Yes Yes No Yes, but not all [4]	 Yes Yes Yes Yes No Yes, but not all [4]	 Yes Yes Yes Yes Yes Yes	 Yes Yes Yes Yes Yes Yes
Availability reports	No	No	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
Usage reports	No	No	Yes	Yes

Table 3. Available reports in each IBM Tivoli Storage Productivity Center license (continued)

Reports	Basic Edition	Disk Edition	Data Edition	Standard Edition
Usage Violation reports	No	No	Yes	Yes
Backup reports	No	No	Yes	Yes
Monitored Computer Storage Space reports	Yes	Yes	Yes	Yes
Storage Subsystem reports	Yes	Yes	Yes	Yes
Storage Subsystem Performance reports	No	Yes	No	Yes
Switch Performance reports	No	No	No	Yes
Tape Library report	Yes	Yes	Yes	Yes

Notes:

1. You must install Agent Manager and a Data agent to enable batch reports for these licenses. The Data agent must be installed on the computer on which 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
   - Data edition: Storage Subsystem Performance, Switch Performance
3. The following Fabric system reports are not available in the Basic, Disk, and Data editions: Top switch Ports Data Rate Performance, Top Switch Ports Packet Rate Performance.
4. The following Asset > System-wide reports are not available in the Basic and Disk editions: Storage Subsystems, Disks, Volumes.

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 to identify trends and prevent problems.
- Provides a wide variety of standardized reports about filesystems and storage infrastructure to track usage and availability.
- Provide file analysis across platforms to help you to identify and reclaim space used by non-essential files.
- Provide policy-based management and automated capacity provisioning for filesystems 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. View when files are created, accessed, and modified, and by what group or user. This type of information helps system administrators map storage resources to the consumers of the resource. The ability to map storage consumption to storage hardware has become increasingly important as the size of open systems environments has increased.

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 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 make adjustments 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.

For example, Data Manager can notify you when your system is running out of disk space or warn you of impending storage hardware failure. By alerting you to these and other issues related to your stored data, Data Manager enables you to prevent unnecessary system and application downtime.

Data Manager for Databases

Data Manager for Databases provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance and assets in your relational databases. It can help you identify, evaluate, control and predict the storage needs of Relational Database Management Systems (RDBMSs), which include Oracle, Sybase SQL Server, Microsoft SQL Server, and UDB/DB2.

Data Manager for Databases is policy-based and through its autonomic self-healing capabilities it can detect potential problems and automatically make adjustments based on the policies and actions you have established. For example, it can notify you when your database tables are running out of storage space or warn you of a dropped tablespace. By alerting you to these and other issues related to your stored data, it enables you to prevent unnecessary system and application downtime.

The following table lists the RDBMS objects whose storage usage is monitored by Data Manager for Databases.

Oracle	SQL Server	Sybase	DB2/UDB
Instances	**Instances**	**Instances**	**Instances**
Databases	**Databases**	**Devices**	**Databases**
Tablespaces	**Datafiles**	**Databases**	**Tablespaces**
Datafiles	**Tables**	**Fragments**	**Containers**
Tables	**Indexes**	**Tables**	**Tables**
Indexes	**Log Files**	**Indexes**	**Indexes**
Extents			

Oracle	SQL Server	Sybase	DB2/UDB
Segments			
Redo Log files			
Control Files			
Archive Log			
Directories			

Note: Data Manager for Databases does not currently support the monitoring of clustered database applications.

Data Manager for Databases collects information about RDBMS storage assets and performs tasks against the RDBMS storage assets. Methods for managing and automating capacity utilization of your databases help you:

- Monitor storage assets associated with enterprise-wide databases and notification of potential problems before they occur
- Make intelligent capacity management decisions based on current and trended historical data
- See how much storage is being consumed by users, groups of users and OS's within the database
- Create policy-based management for databases when user-defined threshold is reached
- View storage utilization management from a database and application perspective

Data Manager for Chargeback

Data Manager for Chargeback uses the storage usage information gathered by Data Manager and Data Manager for Databases to generate invoices that chargeback for storage usage.

With Data Manager for Chargeback you can collect usage information on a departmental, group, or user level. You can allocate costs by storage usage by user, disk capacity by computer, tablespace, or file system/physical device. You can create cost centers by creating user, computer, or tablespace groups, allowing organization to chargeback individuals or business units for their storage usage. By understanding the costs associated with existing and future storage usage, you can improve the use of that storage and reduce the costs associated with its maintenances and upgrades.

Using Data Manager, you can run monitoring jobs that gather detailed statistics about the usage and trending of the storage consumed by the users and user groups within your organization. Using Data Manager for Databases, you can run monitoring jobs that gather detailed statistics about the usage and trending of the storage consumed within RDBMS Instances by the users and user groups within your organization.

In addition to providing invoicing for storage usage, Data Manager for Chargeback also integrates with the chargeback systems already implemented in your environment. It provides you with a higher level, application-specific CIMS output format which can be imported into CIMS applications. With Data Manager for Chargeback, you can export chargeback data for direct import into CIMS, giving you the ability to integrate your storage usage data with other enterprise chargeback information and processes.

Disk Manager

Disk Manager helps you manage SANs and heterogeneous storage from a single console.

With Disk Manager you can manage network storage components based on SMI-S, such as:

- IBM System Storage SAN Volume Controller
- IBM Tivoli Storage Enterprise Storage Server (Tivoli Storage Enterprise Storage Server)
- Tivoli Storage Productivity Center Disk Subsystems (DS4000, DS6000, and DS8000 Series)
- IBM XIV Storage System (Note that some disk management features are not supported for XIV Storage System as described in the planning for XIV Storage System information in the IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide.)
- Other storage subsystems that support the SMI-S standards

Device discovery is done using the Service Location Protocol (SLP), as specified by SMI-S. Configuration of the discovered devices is possible in conjunction with CIM agents associated with those devices, using the standard mechanisms defined in SMI-S. Disk Manager also gathers events, and can launch an element manager specific to each device.

Disk Manager can:

- Collect and store performance data and provides alerts
- Provide graphical performance reports
- Help optimize storage allocation
- Provide volume contention analysis

Through the use of data collection, setting of thresholds and use of performance reports, performance can be monitored for the Tivoli Storage Enterprise Storage Server, DS4000, DS6000, DS8000, SAN Volume Controller, and any other storage subsystem that supports the SMI-S block server performance subprofile. The performance function starts with the data collection task, which captures performance statistics for the devices and stores the data in the database.

After data collection is done, you can set thresholds to identify exceptions for certain performance metrics, depending on the type of device. Threshold checking is performed during data collection, and when performance is outside the specified bounds, alerts can be generated.

After performance data has been collected, you can configure Disk Manager to present graphical or text reports on the historical performance behavior of specified devices, or of specific internal components of those devices. The performance reports provide information about past period performance metrics or current performance in graphical form.

Note: Tivoli Storage Productivity Center supports discovery of and reporting on both FB and CKD volumes (and pools and arrays), but can create only FB volumes. Tivoli Storage Productivity Center does not support the creation of CKD volumes.

Managing TagmaStore CIM agents

The TagmaStore CIM agents are provided by Hitachi Data Systems for the TagmaStore storage subsystem. The TagmaStore CIM agent collects information from the TagmaStore storage subsystem.

IBM Tivoli Storage Productivity Center now 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 will support the TagmaStore as a Storage Virtualizer. Tivoli Storage Productivity Center will be able to display information for virtual disks and local disks.

Tivoli Storage Productivity Center cannot provide correlation information between the virtual storage used by TagmaStore and the volumes created on the storage due to an existing limitation of the CIM agent 5.8 from Hitachi Data Systems. However, Tivoli Storage Productivity Center reports correctly display the correlation between volumes created on the local storage and the local disks. This limitation has no impact on the topology but it does affect several reports that show the correlation:

- **Data Manager > Reporting > Asset > By Storage Subsystem > <HDS_device> > Disks > Volumes** (will not show the relation between the disk and Volume or volume in the tree)
- **Disk Manager > Reporting > Storage Subsystem > Volume to Backend Volume Assignment**
- **Disk Manager > Reporting > Storage Subsystem > Computer Views** (because Tivoli Storage Productivity Center cannot populate the information for volumes created on the virtual disks)

All volumes are created from a storage pool that is allocated from a primordial storage pool and an imported primordial pool. A volume cannot be created over both local and virtual extents.

For volume correlation, the host machine must have the Data agent installed and the TagmaStore device should be in the same SAN fabric. There should also be a zone configured in the active zoneset between the ports of the host machine and the ports of the TagmaStore device. The Fabric 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.

Tape Manager

IBM Tivoli Storage Productivity Center Tape Manager helps you manage your SMI-S compliant tape libraries.

Tape Manager lets you manage those tape libraries which are based on SMI-S 1.1 profile for tape libraries.

You must have at least Tape Operator authority to access the functions in the Tape Manager node. To view your authority level, open the ″Role to Group Mappings″ window in the **Administrative Services** → **Configuration** node in the navigation tree.

With Tape Manager you can:

- Discover SMI-S compliant tape libraries attached to your SAN fabrics
- View tape library information
- Monitor tape libraries by probing for data and viewing alerts
- Create and view asset and capacity reports using data collected from your tape libraries
- Use the topology viewer to view information about tape libraries

Fabric Manager

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

Fabric Manager is an enterprise scalable solution architected to ANSI SAN standards, allowing you to choose the products right for your environment. It helps you:

- Simplify the task of SAN management and configuration
- Ensure SAN availability
- Improve SAN return on investment

Fabric Manager is able to do SAN fabric performance and capacity management reporting and monitoring. Zoning is one of many SAN security options. Zoning is the only configuration supported for this release. Along with subsystem based volume masking, it is almost universally employed to ensure that only systems that are authorized can access selected data. Zoning is provided by the switch at the port level, so that, for example, a host on port X can access the subsystem connected by port Y. When business needs change, often zones must change also.

Zone control is enabled from a centralized location. Existing zones can be discovered and the members that make up the zone can be viewed and modified by adding or deleting them. New zones can be created and existing zones can be deleted. Fabric Manager makes working with zones easy by providing a GUI that discovers SAN devices and makes them available to add or remove from zones as individual devices.

The ability to do switch performance and capacity management reporting and monitoring can help you to determine if more bandwidth is needed. Wide-area and local-area IP networks, and storage area networks (SANs) move data from one place to another. The management of the bandwidth is needed to continually monitor for link utilization and errors. Data needs to be gathered to tune resources, balance workloads, and do capacity planning.

With Fabric Manager, the comprehensive management of multi-vendor SANs can help simplify the IT infrastructure by providing automatic resource and topology discovery, monitoring, and alerts, and zone control. It brings all the sources of information about SAN topology and configuration into a single place, and creates topology mapping of the SANs. This topology offers both host-centric and device-centric views that can be displayed graphically. The SAN topology display tracks all topology and configuration changes through in-band, out-of-band, and SMI-S monitoring. Without this type of centralized topology, information from a number of sources, such as Element Management tools, device logs, and SNMP traps would have to be continually monitored and manually correlated to determine the current SAN configuration and topology.

Fabric Manager provides you with the ability to view events happening in your SAN environment and records state changes. The events are displayed in a color-coded fashion and can be further customized to reflect organizational priorities. It will forward events signaling topology changes or updates to the IBM Tivoli Enterprise Console, another SNMP manager, or both. Source side filtering by the Fabric Manager enables you to control what events get sent to your centralized console. Filtering helps ensure that the administrators get the information they need and are not overwhelmed by extraneous information.

Graphical user interface

The tree-based navigation system of the user interface lets you see a hierarchical organization of the product features in the left pane while viewing detailed reports and charts in the right pane.

Figure 1 shows the main window, which is displayed when you log into IBM Tivoli Storage Productivity Center.

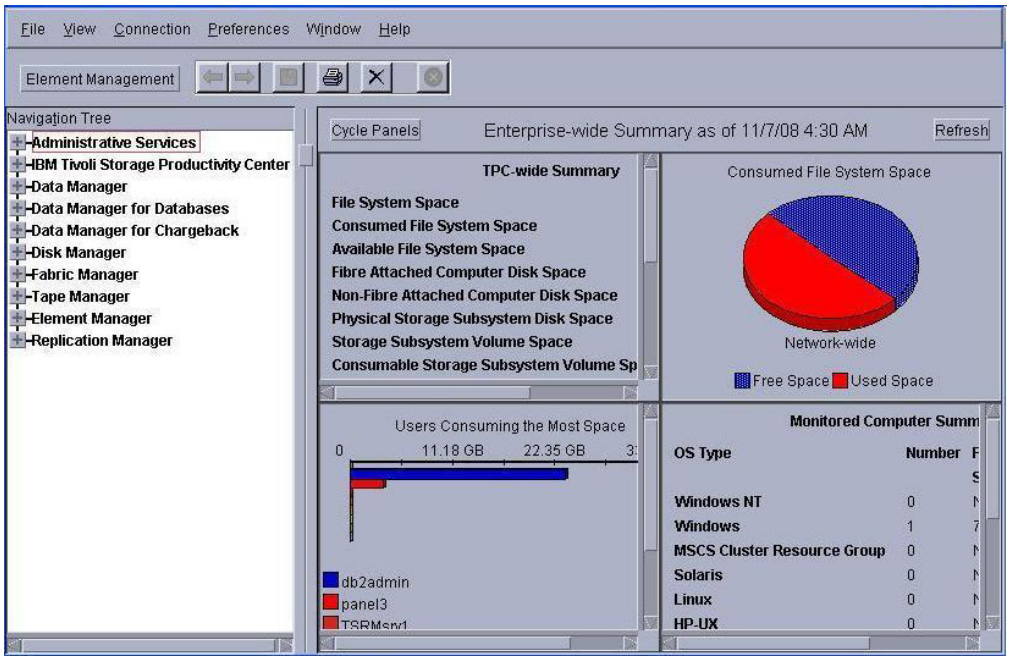

Figure 2. Tivoli Storage Productivity Center Main Window

The Tivoli Storage Productivity Center user interface consists of a navigation tree and a content pane.

- **Navigation tree**

The navigation tree is on the left side of the main window. You can expand, select, and collapse the tree nodes to navigate through the interface.

When you start Tivoli Storage Productivity Center, the navigation tree is expanded to show all the high level functions. You can drill down on an element in the tree by clicking on it or by clicking on the expand icon.

When you right-click a node a pop-up menu displays, which lets you perform additional actions for the node. **Expanding and collapsing the navigation tree:** If the navigation tree gets too large, or if you want to return it to its original state, right-click on the major nodes of the tree and select **Collapse Branch** or **Expand Branch**. If you right-click on the **Tivoli Storage Productivity Center** node and select **Collapse Branch**, that entire subtree collapses. Then, right-click on the main **Tivoli Storage Productivity Center** node and select **Expand Branch** to return the navigation tree to its original state—expanded only to show the main functions.

Printing the navigation tree: Right-click on a node and choose **Print Branch** to print all the sub-items within that node. For example, to print a listing of all your defined scans, expand the **Scans** node, right-click on **Scans**, and select **Print Branch** from the pop-up menu.

- **Content pane.** The content pane is on the right side of the main window. When you select a node in the navigation tree, the corresponding function window opens in the content pane. You can use the windows that open in the content pane to define and run the different functions (for example, monitoring jobs, alerts, reports, etc.) available within Tivoli Storage Productivity Center. The information shown on the pages in the content pane will vary depending on the function with which you are working.

The GUI works like a Web browser, whether you are accessing it through the Web or from the installed interface application. The  and  buttons on the tool bar move control back and forth through Tivoli Storage Productivity Center windows that you have been working on during a product session.

Navigation tree

The navigation tree shows all of the available resource management tasks, organized by major topic nodes for Administrative Services, IBM Tivoli Storage Productivity Center, Data Manager, Disk Manager, Fabric Manager, and Tape Manager.

Administrative Services

Provides product configuration functions for Tivoli Storage Productivity Center.

This node is always present.

Tivoli Storage Productivity Center

Provides operational functions for Tivoli Storage Productivity Center.

This node is always present.

Data Manager
Provides operational functions for the Data Manager.

This node is present when Data Manager is installed and licensed.

For Data Manager functions, you must use Data agents or Storage Resource agents to collect data about your storage resources. For example, use scans, pings, and probes to collect data about the computers on which the agents are installed.

Data Manager for Databases
Provides operational functions for Data Manager for Databases.

This node is present when Data Manager is installed and licensed.

Data Manager for Chargeback
Provides operational functions for Data Manager for Chargeback.

This node is present when Data Manager is installed and licensed.

Disk Manager
Provides operational functions for the Disk Manager.

This node is present when Disk Manager is installed and licensed. A subset of the functions is provided when you install the Data Server.

For Disk Manager functions, the subsystem CIMOMs and data collection are required (for example, through probes).

Fabric Manager
Provides operational functions for the Fabric Manager.

This node is present when Fabric Manager is installed and licensed.

For Fabric Manager functions, the available agent types are CIM agents, in-band fabric agents, and out-of-band fabric agents. You must run a data collection job (for example, through probe jobs). If you want to perform zone control actions, the in-band fabric agents for QLogic and Cisco fabrics must be deployed. For Fabric Manager features such as zone control and performance monitoring, you must perform data collection from performance monitors. For a list of current agent types supported for the Switch Performance Management functionality and the Fabric Zone Configuration functionality, see "Fabric Manager" on page 775.

Tape Manager
Provides operational functions for the tape library.

This node is present when the Device server is installed.

For Tape Manager functions, the library CIMOMs and data collection are required (for example, probes).

Element Manager
Allows you to access the **Element Manager** tab of the Configuration Utility where you can configure and administer IBM System Storage DS8000 and other storage subsystems.

Replication Manager
Allows you to access the **Replication Manager** tab of the Configuration Utility where you can manage IBM Tivoli Storage Productivity Center for Replication and to set the triggering condition and actions for alerts associated with Tivoli Storage Productivity Center for Replication.

Dashboard

The dashboard provides a concise, yet detailed overview of the health of your storage environment. It quickly points out potential problem areas that need further investigation.

When you first log in to IBM Tivoli Storage Productivity Center you are asked if you want to collect statistics. You must answer yes in order for the dashboard information to be displayed. The dashboard will appear whenever you start the product or close all the active windows in a running session.

The dashboard contains four display areas and seven panels that you can cycle through to view your environment's storage information.

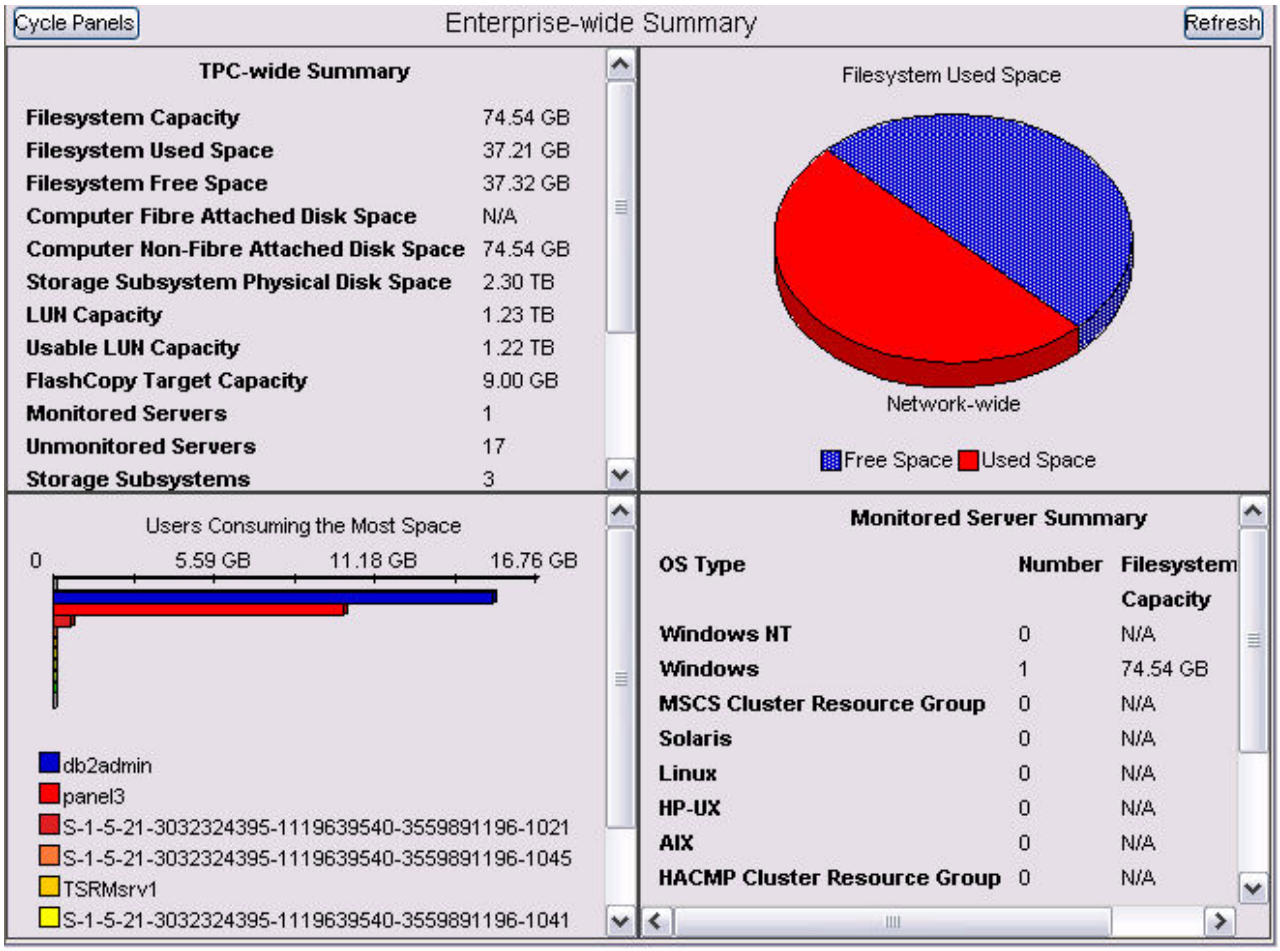

Using the dashboard:

Use the dashboard to monitor the overall health of your storage resources.

When you first login to IBM Tivoli Storage Productivity Center or close all the active windows, the content pane displays the "dashboard". The dashboard provides a concise, yet detailed overview of the health of your storage environment. It enables you to quickly point out potential problem areas that need further investigation. The dashboard contains four displayable areas and eight panels that you can cycle through to view your environment's storage information.

Keep in mind the following considerations when viewing the information on the dashboard:

- The information displayed in the dashboard is collected by monitoring jobs and aggregated by the History Aggregator. We recommend that you run a history aggregator job regularly to ensure that the latest information is displayed in the dashboard. See the IBM Tivoli Storage Productivity Center for information about how to create and schedule history aggregator jobs.
- The value in the **All** field on the Alerts Pending panel might be higher than the sum of the alerts shown in the other fields on that panel. This occurs because some alerts are not associated with a specific category, and are reflected in the value that appears in the **All** field only.
- The names of some fields have been updated. For your reference, the original names of those fields appear within parentheses ( ) next to their new names in this topic.
- Click **Refresh** to update the displayed panels. You must click **Refresh** to ensure that the latest information from the history aggregator is displayed in the dashboard.
- Click **Cycle Panels** to cycle through the panels in the dashboard and view information about your storage assets.

  TPC-wide Summary panel
  This panel provides an enterprise-wide summary of your storage through the following values:

  File System Space (Capacity)
  Total amount of unique file system storage space on computers where a Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:
  - File systems that become missing after a probe.
  - File systems discovered on virtual machines where the Data agent is installed.
  - File systems that are mounted remotely and are not identified to a network attached storage server.
  - File systems that are not mounted.

  Consumed File System Space (Filesystem Used Space)
  Total amount of unique file system storage space that is used or consumed by the computers where a Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:
  - File systems that become missing after a probe.
  - File systems on discovered hypervisors that are assigned to virtual machines where the Data agent is installed.
  - File systems on virtual machines where the Data agent is installed, but the hypervisors for the virtual machines have not been discovered and the total used space exceeds the capacity of the total file system capacity.
  - File systems that are mounted remotely and are not identified to a network attached storage server.
  - File systems that are not mounted.

  Available File System Space (Filesystem Free Space)
  Total amount of unique file system storage space that is not used

or available to the computers where a Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Fibre Attached Computer Disk Space (Computer Fibre Attached Disk Space) Total amount of unique disk space on computers where a Data agent is installed and the disks are attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this field:

- Computer disks that reside behind a fibre channel port that has not been identified by a Data agent.
- Computer disks that do not reside behind a fibre channel port.
- Computer disks discovered on virtual machines where a Data agent is installed.
- Computer disks that become missing after a probe.

Non-Fibre Attached Computer Disk Space (Computer Non-fibre Attached Disk Space)

Total amount of unique disk space on computers where a Data agent is installed and the disks are not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this field:

- Computer disks that reside behind a fibre channel port.
- Computer disks discovered on virtual machinewhere a Data agent is installed.
- Computer disks that become missing after a probe.

Physical Storage Subsystem Disk Space (Storage Subsystem Physical Disk Space)

Total amount of unique internal physical disk space on the storage subsystem arrays that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this field:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller) that are attached from a storage subsystem.
- Storage subsystem disks that are on storage subsystems that have not been probed by the subordinate servers or the master server.

Storage Subsystem Volume Space (LUN Capacity)

Total amount of unique storage subsystem volume space and virtual storage volume space on the storage subsystem arrays that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this field:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller).
- Storage subsystem disks that are on storage subsystems that have not been probed by the subordinate servers or the master server.

Consumable Storage Subsystem Volume Space (Usable LUN Capacity)
Amount of file system storage space for a file system. Total amount of unique storage subsystem volume space and virtual storage volume space on monitored storage subsystem arrays that can be assigned or are assigned to systems within the network . Tivoli Storage Productivity Center does not include the following columns in its calculation of the value for this field:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are used as flash copy targets.
- Storage subsystem volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller).

FlashCopy Target Storage Subsystem Volume Space (FlashCopy Target Capacity)
Total amount of unique storage subsystem volume space and virtual storage volume space on monitored storage subsystem arrays that is identified as flash copy target space. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this field:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are not used as flash copy targets.

Monitored Computers (Servers)
Total number of computers where a Data agent is installed and are detected by the Data server. Tivoli Storage Productivity Center does not include the following systems in its calculation of the value for this field:

- Systems where a Data agent is installed but was not started or was unable to communicate with the Data server.
- Systems where a Data agent is not installed.

- Systems where a Data agent is installed but that agent communicates with a different Tivoli Storage Productivity Center server.
- Systems that have been removed from the Administrative Services > Data Sources > Data Agents panel.
- Hypervisors

Unmonitored Computers (Servers)
Total number of computers found in a Windows domain that do not have a Data agent communicating with the Data server. Tivoli Storage Productivity Center does not include the following systems in its calculation of the value for this field:

- Systems that are not running a Windows operating system.
- Systems that are not in a Windows domain.
- Systems where a Data agent is installed and is communicating with a Data server.

Storage subsystems
Total number of storage subsystems that have been discovered and probed. Tivoli Storage Productivity Center does not include the following storage subsystems in its calculation of the value for this field:

- Storage Subsystems that have been discovered but not probed.
- Storage Subsystem that have not been discovered.
- Storage Subsystems that have been removed from the Disk Manager > Storage Subsystems panel.

Users Total number of user accounts discovered by Data agents.

Disks Total number of unique disks discovered by Data agent probes, hypervisor probes, network storage filer probes, and storage subsystem probes. This value does not include disks that have become missing after a successful probe.

Storage Subsystem Volumes (LUNS)
Total number of unique storage volumes and virtual storage volumes discovered by storage subsystem probes. This value does not include storage volumes or virtual storage volumes that have become missing after a successful storage subsystem probe.

Filesystems
Total number of unique file systems discovered by probes. This value excludes file systems that have become missing after a successful probe.

Directories
Total number of directories discovered by scans. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this field:

- Directories within file systems that were discovered by a probe but were not scanned.
- Directories that have become missing after a successful probe.

Files Total number of files discovered by scans. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this field:

- Files within file systems that were discovered by a probe but were not scanned.
- Files that have become missing after a successful probe.

Rollup Summary panel

This panel provides a summary of information collected by the master server and by Tivoli Storage Productivity Center servers that are subordinate servers.

File System Space (Capacity)

Total amount of unique file system storage space on the computers that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems discovered on virtual machines where the Data agent is installed.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.
- File systems that have not been probed by subordinate servers or the master server.

Consumed File System Space (Used Space)

Total amount of unique file system storage space that is used or consumed by the systems that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems on discovered hypervisors that are assigned to virtual machines where the Data agent is installed.
- File systems on virtual machines where the Data agent is installed, but the hypervisors for the virtual machines have not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.
- File systems that have not been probed by subordinate servers or the master server.

Available File System Space (Filesystem Free Space)

Total amount of unique file system storage space that is not used or available to the systems that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

– File systems that have not been probed by subordinate servers or the master server.

Physical Storage Subsystem Disk Space (Storage Subsystem Physical Disk Space)

Total amount of unique internal physical disk space on monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this field:

– Storage subsystem disks that become missing after a storage subsystem probe.
– Back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller) that are attached from a storage subsystem.

Storage Subsystem Volume Space (LUN Capacity)

Total amount of unique storage subsystem volume space and virtual storage volume space on monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this field:

– Storage subsystem volumes or virtual storage volumes that become missing storage subsystem probe.
– Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller).

Consumable Storage Subsystem Volume Space (Usable LUN Capacity)

Total amount of unique storage subsystem volume space and virtual storage volume space on storage subsystem arrays that are monitored by subordinate servers and the master server and can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this field:

– Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
– Storage subsystem volumes or virtual storage volumes that are used as flash copy targets.
– Storage subsystem volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
– Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
– Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller).
– Storage subsystem disks that are on storage subsystems that have not been probed by the subordinate servers or the master server.

Monitored TPC Servers

Total number of subordinate servers that are monitored by the master server.

Monitored Computers (Servers)
Total number of computers that are monitored by Tivoli Storage Productivity Center and are detected by the subordinate and the master server. Tivoli Storage Productivity Center does not include the following systems in its calculation of the value for this field:

- Systems where a Data agent is installed but was not started or was unable to communicate with the Data server.
- Systems where a Data agent is not installed.
- Systems where a Data agent is installed but that agent communicates with a Tivoli Storage Productivity Center server that is not a subordinate server or the master server.
- Systems that have been removed from the Administrative Services > Data Sources > Data Agents panel.

Storage subsystems
Total number of storage subsystems that have been discovered and probed by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following storage subsystems in its calculation of the value for this field:

- Storage Subsystems that have been discovered but not probed.
- Storage Subsystem that have not been discovered.
- Storage Subsystems that have been removed because of the storage subsystem settings on the Removed Resource Retention panel.

Consumed File System Space (Filesystem Used Space panel)
This chart displays the total amount of used file system storage space across the entire network.

Users Consuming the Most Space (chart) panel
This panel displays the following information:

- A list of the users consuming the most space within your environment
- A graphical view of the amount of space they are consuming

Note: The space usage of administrative users within your environment is not included as part of this dashboard report.

Monitored Computer Summary panel
This panel displays the following information:

- The number of monitored servers running each Operating System type
- Total amount of file system and disk capacity per Operating System type

Filesystems with Least Free Space Percentage panel
This panel shows the monitored machines with the least percentage of free space, including the following:

- Percent of free space on each computer
- Total file system capacity on each computer
- Name of the computers, file systems, and mount points

Users Consuming the Most Space Report panel

This panel displays the following information:

- A list of the users consuming the most space within your environment
- The total amount of storage each user is consuming

Alerts Pending panel

Displays the active alerts that are still pending only.

Note: Because of this, the number of alerts on this panel might be less than the number of alerts shown in the IBM Tivoli Storage Productivity Center > Alerting > Alert Log. Alerts for failed jobs (pings, probes, and scans) are not listed in separate rows on this panel, but are included in the total displayed in the **All** row.

These alerts are organized into the following categories:

- All
- Computer
- Disk
- Filesystem
- Directory
- User Quota
- OS User Group Quota
- Directed to *user*
- Storage Subsystem
- Switch
- Fabric
- Endpoint Device
- External
- Tape Library
- Hypervisor
- Authentication Configuration
- Replication
- Configuration Analysis

Working with the navigation tree and content pane

This section describes some methods for working with the nodes and windows in the navigation tree and content pane.

Right-clicking and pop-up menus:

Right-click nodes in the navigation tree to access pop-up menus that contain actions you can perform.

As you navigate through the IBM Tivoli Storage Productivity Center user interface, notice that you can right-click nodes in the navigation tree to access pop-up menus that contain actions related to that node. This enables you to quickly access many product functions while at the same time keeping the user interface uncluttered.

For example, depending on the object type you right-click, the pop-up menu that appears will contain the valid operations that you can apply to that object type. If you right-click on a **Data Manager** > **Monitoring** > **Scan** > **scan_job** (or any other

data collection job node), a pop-up menu with the following options appears:

Option	Description
Delete	Delete the job definition.
Rename	Rename the job.
Refresh Job List	Refresh the list of job runs for the selected job.
Update Job Status	Update the status of a job. When you select this option, Tivoli Storage Productivity Center will update the visual indicator next to a job that indicates its status. For example, select this option to have Tivoli Storage Productivity Center turn the status of any non-running jobs to red. **Note:** While updating job statuses, Tivoli Storage Productivity Center contacts all the agents that currently have running jobs. In large environments with many agents, this process might take significant time to complete. During this time, an hourglass appears on the user interface to indicate that the update job status process is running. If processing is taking too long to complete, you can click the red ″X″ icon in the toolbar to continue working with the user interface. If you cancel an update job status process, we recommend waiting before attempting to restart the process. This will provide the server time to complete the previously canceled request and prevent any performance issues.
History	Show a job run history on the right pane.
Run Now	Run the job now.
Print Branch	Print this node and any expanded node in the branch under it.

Selecting windows in the content pane:

Navigate the open windows in the content pane.

When you click a node in the navigation tree, a window for that node opens in the content pane. Each time you select a new node, the open window for the previously selected node is not deleted. Instead, it moves into the background and is accessible by clicking  and  on the toolbar. You can also select currently active windows from the Window list menu > Show Window List.

When viewing the active windows listed in the Window menu, a > appears next to the current window (the window that currently occupies the content pane). Objects that you are currently editing (such as a Scan definition) will be marked with a black arrow if they have no pending changes that have not been saved; a solid red arrow indicates objects with pending changes that have not been saved.

Menu bar

Use the menu bar at the top of the main window to access some of the main functions within IBM Tivoli Storage Productivity Center.

The following sections describe the options available from the menu bar.

File menu:

Use the options in the File menu to save objects and jobs, export report data, print, refresh alerts, close the current window, and exit the product session.

Menu Option	Description
Save	Save the changes for the current object (for example, groups, profiles, etc.) or job (for example, probes, scans, pings, etc.) that you are editing.
Save As	Save the changes for the current object or job that you are editing under a different name.
Export Data	Export the report data to other formats: • Formatted Report file (.txt) • Comma delimited file (.csv) • Comma delimited file with headers • HTML file This allows you to import the data to another application, such as Excel, if needed.
Print	Print the current data in the content pane. You can print tabular reports or report charts. You can also save the output to a PDF file or HTML file, and print those files at a later time.
Print Tree	Print the expanded contents of the navigation tree. You can also save output to a PDF or HTML file, and print those files at a later time.
Refresh Alerts	Refresh all the alerts from the database repository and update the GUI by highlighting the appropriate objects and showing all the new alerts in the alert logs.
Close	Close the current window that occupies the content pane.
Exit	End the current session.

View menu:

Use the options in the View menu to hide or show the navigation tree and highlight the node in the navigation tree for the currently open window.

Menu Option	Description
Tree	Hide and show the navigation tree. If you want to increase the size of the content pane to occupy the entire window, clear the check box for the **Tree** option. To redisplay the navigation tree next to the content pane in the main window, select the **Tree** option.
Current Page in Tree	Display the highlighted node on the navigation tree that was responsible for displaying the data on the content pane. Use this option to reorient the navigation tree with the window that is currently displayed in the content pane.

Connection menu:

Use the option in the Connection menu to end the current session and establish a new connection to the same IBM Tivoli Storage Productivity Center server or to a different server.

Menu Option	Description
New Connection	Initiate a new connection to a server from the current GUI session. For example, use this option to connect to the same server using a different user ID, or to a completely different server. All the current changes that are unsaved in the open windows will be lost when you establish a new connection.

Preferences menu:

Use the options in the Preferences menu to change the appearance of the user interface, edit the way the product manages open panels, determine what alerts are shown upon login, select what tab appears first when generating reports, and specify what unit of measurement to use when displaying storage sizes in reports.

Menu Option		Description
Look and Feel	**Windows**	Set the look and feel of the GUI to appear like Microsoft Windows GUI.
	CDE/Motif	Set the look and feel of the GUI to appear like the UNIX Motif interface.
	Metal	Set the look and feel of the GUI to appear like Metal, the cross-platform Java look and feel.

Menu Option		Description
Edit General	**Panel Retention**	The **Panel recall depth** value specifies the number of windows that are linked in the content pane for access via the  and  buttons on the toolbar. The default is 20. The **Panels kept in memory field** specifies the number of windows to store in memory, which helps to optimize the system performance. The default is 2. **Note:** Any open views in the topology viewer (for example, an L2:Storage Subsystem view, an L0:Computers view) are considered part of the topology viewer panel and are listed as tabs at the top of the topology viewer panel. If the value in this field is a low number, such as the default of 2, then opening additional non-topology panels can exceed the number of panels that are kept in memory. This can result in the topology viewer panel refreshing and any tabbed views being closed. To avoid this situation, change the value in this field to a higher number, such as 10.
	On login, show	Control which alerts you will view upon logging in to a user interface session. You can select: • No Alerts • Alerts Directed to <Userid> • All Active Alerts
	Initial Reporting Tab to Display	Control what tab will be displayed when you first generate and view a report. You can select: • Selection tab • Report tab • Chart tab (if Available)

Menu Option	Description
Advanced Options	Select the unit of measurement that should be used to reflect the storage sizes included in IBM Tivoli Storage Productivity Center reports. This setting affects the results in online, exported, and Batch reports. You can choose from the following options: • **As Appropriate:** Choose this option to automatically determine the unit measurement (kb, mb, gb, tb) to display for storage sizes in a report. • **Kilobytes (KB):** Choose this option to display storage sizes in reports using the kb unit of measurement. For example, if the storage capacity for a computer is 100 megabytes, it would be shown as 100000kb in a capacity report. • **Megabytes (MB):** Choose this option to display storage sizes in reports using the mb unit of measurement. For example, if the storage capacity for a computer is 1 gigabyte, it would be shown as 1000mb in a capacity report. • **Gigabytes (GB):** Choose this option to display storage sizes in reports using the gb unit of measurement. For example, if the free space on a computer is 500 megabytes, it would be shown as .5gb in a free space report. • **Terabytes (TB):** Choose this option to display storage sizes in reports using the tb unit of measurement. For example, if the free space on a computer is 350 gigabytes, it would be shown as .35tb in a free space report.

Window menu:

Use the options in the Windows menu to close all the open windows that do not have any changes, display the dashboard, and list all the open windows.

Menu Option	Description
Close all Windows Without Changes	Close all of the open windows that do not contain any modifications. Any windows that contain active changes to editable objects will not be closed.
Show Dashboard	Display the IBM Tivoli Storage Productivity Center dashboard.
Show Window List	Lists names of all open windows active in the content pane. Select one of these options to view the corresponding window. Beside each window name is an icon that indicates the status of the window: • **no icon** - No changes were made to the window • **solid red circle** - Changes were made, but have not been saved • **black arrow** (>) - Current window displayed. No changes were made • **solid red arrow** (>) - Current window displayed. Changes were made, but have not been saved

Help menu:

Use the options in the Help menu to access online help for the product, launch the information center Web page, launch the Supported Products Web pages, launch the demonstrations Web page, and display product information (such as the version number).

Menu Option	Description
Help Topics	Select this menu item to launch IBM Tivoli Storage Productivity Center's online Help system. The Help system will appear in its own window and display a table of contents from which you can access information about how to use the product.
Help for Displayed Panel	Select this menu item to launch the online Help system and display help information for the product window that is currently displayed. You can also press F1 to launch the help topic for the product window that you are viewing.
Launch Information Center	Select this menu item to launch the information center for IBM Tivoli Storage Productivity Center Information Center in a Web browser. The IBM Tivoli Storage Productivity Center Information Center contains detailed instructions that explain how to install, upgrade, configure, maintain, and use the Tivoli Storage Productivity Center product. See "Web browser support" on page 42 for information about Web browser support for launching the IBM Tivoli Storage Productivity Center Information Center.
Supported Products Data	Select the links under this menu item to launch Web pages that contain additional documentation for Disk Manager, Fabric Manager, and Data Manager. This documentation includes matrixes that show the devices supported by each of the managers. • **IBM Tivoli Storage Productivity Center for Disk Supported Products**. Click this item to view the Support for Tivoli Storage Productivity Center for Disk Web page. To view the devices supported by Disk Manager, scroll to the **Integration/Operation** section and click the link that matches your version of Tivoli Storage Productivity Center. • **IBM Tivoli Storage Productivity Center for Fabric Supported Products**. Click this item to view the Support for Tivoli Storage Productivity Center for Fabric Web page. To view the devices supported by Fabric Manager, scroll to the **Integration/Operation** section and click the link that matches your version of Tivoli Storage Productivity Center. • **IBM Tivoli Storage Productivity Center for Data Supported Products.** Click this item to view the Support for Tivoli Storage Productivity Center for Data Web page. To view the devices supported by Disk Manager, scroll to the **Integration/Operation** section and click the link that matches your version of Tivoli Storage Productivity Center. See "Web browser support" on page 42 for information about Web browser support for launching the supported products data Web pages.

Menu Option	Description
Tivoli Storage Productivity Center Demos	Select this menu item to launch a Web page that contains Flash demos of Tivoli Storage Productivity Center. These demos provide an introduction to the product, the business and storage administrator challenges that are met by the product, and detailed information about the features and functions that the product provides. See "Web browser support" for information about Web browser support for launching the product demos Web page.
Launch IBM Tivoli Storage Productivity Center Information Center	Select this menu item to display product information including: • version information • copyright information • trademark information

Web browser support:

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

The following Web browsers are supported by Tivoli Storage Productivity Center:

AIX

- Mozilla Firefox 2.0

Linux

- Mozilla Firefox 2.0

Windows

- Internet Explorer 7
- Mozilla Firefox 2.0 and 3.0

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.

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 do this, 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 occurs when you use Tivoli Storage Productivity Center to start a session of Tivoli Storage Productivity Center for Replication in an Internet Explorer 7 Web browser and then open and close a wizard or secondary window in that Tivoli Storage Productivity Center for Replication session.

To workaround this problem, start Tivoli Storage Productivity Center for Replication from Tivoli Storage Productivity Center using Firefox, start Tivoli Storage Productivity Center for Replication from IBM Tivoli Integrated Portal, or start Tivoli Storage Productivity Center for Replication by entering its Web address directly in a Firefox or Internet Explorer 7 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 machine where you are running the Tivoli Storage Productivity Center user interface and have a default browser configured to successfully access these Web pages.

Tool bar

The tool bar at the top of the IBM Tivoli Storage Productivity Center main window provides one-click access to key commands.

The buttons on the tool bar are described below:

Button	Description
Element Management	Click this button to open the Element Manager tab in the Configuration Utility. Use the Element Manager tab to configure and start the element managers that are visible to Tivoli Storage Productivity Center. Element managers are programs that enable you to configure and maintain your storage devices.
Back 	Click this button to navigate backward through the active windows in the content pane.

Button	Description
Forward 	Click this button to navigate forward through the active windows in the content pane.
Save 	Click this button to save the current object that you are editing. The object changes are saved to the to the database repository.
Print 	Click this button to print the current data or report chart displayed in the content pane. You can send the output directly to a printer, HTML file, CSV file, formatted file, or a PDF.
Cancel Edit Session 	Click this button to discard content pane. If content pane is an edit session and changes have been made, a prompt notifying you that a change has been made is displayed.
Stop Current Server Request 	Click this button to stop the current server request. This button becomes active when a request is made to the server and the user interface is waiting for a response. During the server request wait time, the cursor arrow in the user interface will appear as an hourglass. If you decide the current request is taking too long or you do not wish to wait any longer, click this button to cancel that request and regain control of the GUI.

Interface elements and themes

The IBM Tivoli Storage Productivity Center user interface has repeating elements and consistent themes that help you perform storage management tasks. By understanding these elements and themes, you can learn to quickly navigate and use the available functions.

User-created objects and jobs:

Create named *objects* to help simplify the data collection, alerting, and reporting process. Objects represent reusable definitions you can run to perform specific tasks. Each object belongs to a specific *object type*.

Each object that you can create belongs to a specific *object type*. For example, some objects can be run as jobs. When you expand the navigation tree under an object, the job runs for that object are displayed. If you create a probe named *sample_probe*, each time that job runs a subnode and it appears for that run under the **IBM Tivoli Storage Productivity Center** > **Monitoring** > **Probes** > ***sample_probe*** node.

The main object types appear in the navigation tree and are defined in the following table:

Object Type	Function and Use	Where does an entry for a job or object appear after I create one?
Monitoring Groups *(Data Manager, Data Manager for Databases, Disk Manager, Fabric Manager, Tape Manager)*	Create named groups of computers, storage subsystems, filesystems, fabrics, tape libraries, directories, tables, databases, tablespaces, users, and OS user groups. This allows you to more easily refer to those entities when you are defining monitoring jobs that collect information about your storage resources.	• **Data Manager > Monitoring > Groups >** *group_type > group_name* • **Data Manager for Databases > Monitoring > Groups >** *group_type > group_name* • **Disk Manager > Monitoring > Groups > Storage Subsystem >** *storage_subsystem_group_name* • **Fabric Manager > Monitoring > Groups > Fabric >** *group_name*
Discovery *(Administrative Services > Discovery)*	Run discovery jobs to: • Find new Windows computers that have been introduced into your environment. Discovery jobs will locate new Windows computers within your environment only--they do not discover computers running non-Windows operating systems. For IBM Tivoli Storage Productivity Center to recognize a new non-Windows computer, you must install and license an agent on that computer. • Identify the servers and volumes within your environment's NetWare trees (NDS trees). A Discovery job will: (for NetWare) log in to NDS trees and enumerate the NetWare servers in those trees (for NetWare) log in to the NetWare servers in the NDS trees and gather volume/disk information. • Discover the CIMOMs in your environment and the storage subsystems, fabrics, switches, and tape libraries managed by those CIMOMs. • Discover out-of-band fabric agents in your environment.	• **Administrative Services > Discovery > CIMOM >** *CIMOM_job_name* • **Administrative Services > Discovery > Out of Band Fabric >** *OutofBandFabric_job_name* • **Administrative Services > Discovery > NetWare Filer >** *NetWareFiler_job_name* • **Administrative Services > Discovery > Windows Domain, NAS, and SAN FS >** *WindowsDomain_NAS_SANFS_job_name*
Probes *(IBM Tivoli Storage Productivity Center > Monitoring > Probes)*	Create probes to itemize and create an inventory of your storage assets, such as clusters, computers, storage subsystems, controllers, switches, disk drives, file systems, logical units, tape libraries, RDBMSs, and so forth. You must run probes before you can view information about your storage assets in the reports available within each of the managers.	**IBM Tivoli Storage Productivity Center > Monitoring > Probes >** *probe_name*
Tivoli Storage Productivity Center Server Probes *(IBM Tivoli Storage Productivity Center > Monitoring > IBM Tivoli Storage Productivity Center Server Probes)*	Create Tivoli Storage Productivity Center Server probes from a master server to gather storage information collected by its associated subordinate servers. You must run Tivoli Storage Productivity Center server probes before you can view information inIBM Tivoli Storage Productivity Center > Rollup Reports.	**IBM Tivoli Storage Productivity Center > Monitoring > IBM Tivoli Storage Productivity Center Server Probes >** *probe_name*

Object Type	Function and Use	Where does an entry for a job or object appear after I create one?
Pings *(Data Manager)*	Create pings to track the availability of your storage assets (not available for running against NetWare Servers and NAS filers). You can view the information gathered by pings in Data Manager reports.	**Data Manager > Monitoring > Pings >** *ping_job_name*
Scans *(Data Manager, Data Manager for Databases)*	Create scans to monitor the usage and consumption of your storage. You can view the information gathered by scans in Data Manager reports.	• **Data Manager > Monitoring > Scans >** *scans_job_name* • **Data Manager for Databases > Monitoring > Scans >** *scans_job_name*
Profiles *(Data Manager, Data Manager for Databases)*	Create profiles to define what files you want to scan and what statistical information you want to gather.	• **Data Manager > Monitoring > Profiles >** *profile_name* • **Data Manager for Databases > Monitoring > Profiles >** *profile_name*
Quotas *(Data Manager, Data Manager for Databases)*	Create quotas to control how much space a user, or group of users, can consume. This consumption can be controlled at three levels: the file system level, the computer level, or the entire network level.	• **Data Manager > Policy Management > Quotas > User/User Group >** *quota_job_name* • **Data Manager for Databases > Policy Management > Quotas > Network/Instance/Database-Tablespace >** *quota_job_name*
Constraints *(Data Manager)*	Use constraints to: • Define the acceptable and unacceptable file types, file sizes, and file owners for a computer or a set of computers in your environment. Constraints enable you to do such things as restrict users from putting certain files (for example, MP3 files) on monitored servers. • Request an IBM Tivoli Storage Productivity Center archive and backup of the largest violating files identified by a constraint. IBM Tivoli Storage Productivity Center protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage.	**Data Manager > Policy Management > Constraints >** *constraint_name*
Alerts *(Data Manager, Data Manager for Databases, Disk Manager, Fabric Manager)*	Create alerts to define events for which you want IBM Tivoli Storage Productivity Center to watch. By defining the events that you want to be alerted about, you will be notified by the product in the manner you specify when one of those events occurs.	• **Data Manager > Alerting >** *alert_type > alert_job_name* • **Data Manager for Databases > Alerting >** *alert_type > alert_job_name* • **Disk Manager > Alerting > Storage Subsystem Alerts >** *storage_subsystem_alert_job_name* • **Fabric Manager > Alerting >** *alert_type > alert_job_name*
Aggregator job	Aggregator jobs run on a periodic basis and sum up storage usage across computers to allow for enforcement of quotas that span multiple computers.	

To view the objects defined for an object type:

- Click on the desired object type name (for example, **Data Manager** > **Monitoring** > **Scans**) and the tree will expand to display a list of all the objects defined for that type.

To create an object for an object type:

1. Right-click on that object type node in the product navigation tree (for example, right-click the **IBM Tivoli Storage Productivity Center** > **Monitoring** > **Probes** node). A pop-up menu appears.
2. Select the **Create <*object type*>** option. The create window for that object type appears in the content pane on the right side of the main window.
3. Enter information about that object.
4. Click

   

   in the toolbar, select **File > Save**, or select **File > Save As** to save the object definition.
5. Enter the name of the object and click **OK**. The name for the object appears as a node under the corresponding object type. For example, **IBM Tivoli Storage Productivity Center** > **Monitoring** > **Probes** > **userID.nameofprobe**.

To edit an object definition:

1. Click on the icon next to the desired object type name (for example, **Scans**) to expand the tree and view a list of all the objects defined for that type.
2. Click the object you want to edit. The edit window for that object type appears in the content pane on the right side of the main window.
3. Edit the information defined for that object.
4. Click

   

   in the toolbar, select **File > Save**, or select **File > Save As** to save the object definition. The object definition is saved in the enterprise repository.

To view information about the runs for a job:

1. Expand the object type whose jobs you want to view (for example, **Probes**) to view a list of all the jobs defined for that type.
2. Expand the job whose runs you want to view.
3. Click on a job run to view information about that run in the content pane.

Job runs:

The term *job run* refers to whenever a job executes. Each job run is recorded as an entry in the product navigation tree and is displayed under the corresponding scheduled job definition.

For example, if you define a scan named **Daily Scan** to run once every day, an entry for each day's job run will appear under the **Data Manager** > **Monitoring** > **Scans** > **Daily Scan** node.

Beside each job run is a visual indicator that reflects the status of the job:

Job Status Indicator	Description
	The job ran successfully on all the computers on which it was deployed. Additional information for performance monitoring jobs: If you stop a running job (blue) manually or the job's duration elapses, the status turns to green (completed successfully).
	The job failed. If an alert was defined for the job, you will be notified by the alert. You can also view the job log to view the error. To do this, click the job to view the run detail for the job that shows each computer on which the job ran. The status column in the detail window indicates if the job ran successfully. Click the magnifying glass icon for the computer that experienced an error to then view the log. Additional information for performance monitoring jobs: If you stop a running job with errors (yellow) manually or the job's duration elapses, the status turns to red (completed with error).
	The job is currently running on one or more computers. Additional information for performance monitoring jobs: • If the job is successful and it is collecting data from the device (CIMOM is available), the status is blue.
	The job completed but contained warnings. Additional information for performance monitoring jobs: • If in the middle of the data collection, IBM Tivoli Storage Productivity Center cannot collect the expected data from subsystem, the running job (blue) turns to yellow (running with errors) and turns back to blue when data is returned successfully. • If (in the middle of the collection) the CIMOM becomes unavailable and Tivoli Storage Productivity Center cannot collect data, the running job (blue) turns to yellow (running with errors) and turns back to blue when the CIMOM is again available.

Note:

- While updating job status, Tivoli Storage Productivity Center contacts all the agents that currently have ″inflight″ jobs. In large environments with many agents, this process might take significant time to complete. During this time, an hourglass will appear on the user interface to indicate that the update job status process is running. If processing is taking too long to complete, you can click the red ″X″ icon in the toolbar to continue working with the user interface. If you cancel an update job status process for any reason, we recommend waiting before attempting to restart the process. This will provide the server component time to complete the previously canceled request and prevent any performance issues.
- Additional information for performance monitoring jobs: If the Tivoli Storage Productivity Center server machine is rebooted or the Device server is restarted, the previous running job tries to start again if its duration has not elapsed. If it

starts successfully and can collect data, the status turns blue. If it cannot connect to the CIMOM or cannot return data, the status turns yellow. If there are other problems that prohibit the job from starting successfully, the status turns red.

IBM Tivoli Storage Productivity Center for Replication overview

This topic provides a brief overview of IBM Tivoli Storage Productivity Center for Replication.

IBM Tivoli Storage Productivity Center now includes IBM Tivoli Storage Productivity Center for Replication, specifically:

- IBM Tivoli Storage Productivity Center for Replication Two Site Business Continuity (BC)
- IBM Tivoli Storage Productivity Center for Replication Three Site BC
- IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z®
- IBM Tivoli Storage Productivity Center for Replication for System z

IBM Tivoli Storage Productivity Center for Replication provides you with continuous availability and disaster recovery solutions with both point-in-time replication, and continuous replication. Other options include the ability to practice what you would do in the event of a disaster, and management servers that enable you to manage your disaster recovery solution if you experience a site outage.

For detailed documentation, refer to the IBM Tivoli Storage Productivity Center for Replication Information Center.

Tivoli Integrated Portal

The IBM Tivoli Storage Productivity Center installation program includes IBM Tivoli Integrated Portal. Tivoli Integrated Portal is a standards-based architecture for Web administration. Tivoli Integrated Portal enables developers to build administrative interfaces for IBM and independent software products as individual plug-ins to a common console network. The installation of Tivoli Integrated Portal is required to enable single sign-on for Tivoli Storage Productivity Center.

Tivoli Integrated Portal and single sign-on

Single sign-on is an authentication process that you can use to enter one user ID and password to access multiple applications. For example, you can access 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 in context feature so you can move smoothly from one application to a functionally-related location in a second application.

Related concepts

"Single sign-on" on page 54
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 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 in context feature to enable you to move smoothly from one application to a functionally-related location in a second application.

"Launch in context" on page 440
The Launch in context feature is available in two modes: inbound and outbound. With the inbound launch in context feature, you can start IBM Tivoli Storage Productivity Center at a specified point in the graphical user interface (GUI). With the outbound launch in context feature, you can use various points in the Tivoli Storage Productivity Center GUI to start an application such as an element manager at a specified point in the manager.

Related tasks

"Administering Tivoli Integrated Portal" on page 157
This topic provides information on how to perform administrative tasks in IBM Tivoli Integrated Portal that are reflected in IBM Tivoli Storage Productivity Center.

Key concepts

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

Starting and stopping the Tivoli Storage Productivity Center GUI

This topic describes how to start and stop the IBM Tivoli Storage Productivity Center graphical user interface (GUI).

Starting the Tivoli Storage Productivity Center GUI

Learn how to start the IBM Tivoli Storage Productivity Center graphical user interface (GUI) from one of the following locations: the Microsoft Windows Start menu, IBM Tivoli Integrated Portal, the Productivity Center icon on your desktop (Windows), or the command line (AIX, UNIX, or Linux).

Starting the Tivoli Storage Productivity Center GUI from Tivoli Integrated Portal:

This section describes how to start the Tivoli Storage Productivity Center graphical user interface (GUI) from Tivoli Integrated Portal. With the new single sign-on feature, you can log on to the Tivoli Storage Productivity Center GUI without having to explicitly enter your username and password.

Before you start Tivoli Integrated Portal and Tivoli Storage Productivity Center, ensure that you are using one of the following Web browsers:

- AIX: Firefox 2.0
- Linux and UNIX: Firefox 2.0
- Windows: Internet Explorer 7, Firefox 2.0, Firefox 3.0

Ensure that a fully qualified domain name is defined for the Tivoli Storage Productivity Center server. To verify the Tivoli Storage Productivity Center server name and modify the name if required, see the Troubleshooting section of the Tivoli Storage Productivity Center Information Center. Contact your Tivoli Storage Productivity Center administrator before you modify the file.

To start Tivoli Storage Productivity Center from Tivoli Integrated Portal, complete the following steps.

1. 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.
2. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator permissions.
3. In the Tivoli Integrated Portal navigation tree, click **Tivoli Storage Productivity Center**.
4. On the Tivoli Storage Productivity Center portlets page, click **Start Storage Productivity Center**. One of the following actions occur:
   - If single sign-on is successful, Tivoli Storage Productivity Center starts without displaying the logon window.
   - If single sign-on is not successful, an error message and the Tivoli Storage Productivity Center logon window are displayed.
   - If the Tivoli Storage Productivity Center Data server or Device server is not accessible, an error message is displayed. The status of both the Data server and the Device server also display.
   - If the status of the Device server and Data server is inaccessible and you started Tivoli Integrated Portal from a remote computer, it is possible that a fully qualified domain name was not defined for the Tivoli Storage Productivity Center server during installation. To check the Tivoli Storage Productivity Center server name and modify the name if required, see the Tivoli Storage Productivity Center troubleshooting documentation.
   - If you are using a Lightweight Directory Access Protocol (LDAP) compliant directory for Tivoli Storage Productivity Center user authentication and the directory is not available, an error message is displayed.

Once you have logged on to Tivoli Integrated Portal, a Lightweight Third-Party Authentication (LTPA) token is created. This token is passed to other applications that you start from Tivoli Integrated Portal for single sign-on authentication purposes.

During the period between when you log on to Tivoli Integrated Portal and when you start another application such as Tivoli Storage Productivity Center from Tivoli Integrated Portal, the following conditions might occur:

- The user password that was used to log on to Tivoli Integrated Portal is changed in the user repository.
- The user ID that was used to access Tivoli Integrated Portal is changed in the repository or removed from the user repository.

- The user repository is not accessible.

Under the first condition, the original user credentials that were used to access Tivoli Integrated Portal are used to access other applications until the timeout period for the LTPA token that is used for single sign-on expires. When the LTPA token expires, you are prompted to re-enter your user ID and password when you attempt to start another application using single sign-on.

Under the second and third conditions, the single sign-on feature does not work. You are always prompted to re-enter your user ID and password when you attempt to start another application.

Starting the Tivoli Storage Productivity Center GUI as a stand-alone application:

This section describes how to start the IBM Tivoli Storage Productivity Center from the Microsoft Windows Start menu, or the Productivity Center icon on your desktop (Windows), or from the command line (AIX, UNIX, or Linux).

To start the Tivoli Storage Productivity Center graphical user interface (GUI) on Windows, click **Start → Programs → IBM Tivoli Storage Productivity Center → Productivity Center**. You can also double-click the IBM Tivoli Storage Productivity Center icon if it is installed on your desktop.

To start the Tivoli Storage Productivity Center GUI on UNIX or Linux, type the following path and command at the command line:

```
/opt/IBM/TPC/gui/TPCD.sh
```

To start the Tivoli Storage Productivity Center GUI on AIX, type the following path and command at the command line:

```
/opt/IBM/TPC/gui/TPCD.sh
```

Logging on to Tivoli Storage Productivity Center:

Use the logon window to specify the user ID and password that you want to use to log on to a IBM Tivoli Storage Productivity Center server.

Define logon information for Tivoli Storage Productivity Center in the following fields:

User ID
: Enter the user ID that you want to use to log on to Tivoli Storage Productivity Center. The roles that are assigned to that user ID determine what nodes in the navigation tree you can see and act upon.

Password
: Enter the password that is associated with the **User ID**.

Server
: Enter the IP address or Domain Name System (DNS) name of the computer on which the Tivoli Storage Productivity Center server is installed.

  When you enter an IP address, you can use either an IPv4 or IPv6 address format depending on the protocol enabled on the computer. For example, enter an IPv6 address for an IPv6-only computer. Enter an IPv4 or IPv6 address for a computer with dual stacks (both IPv4 and IPv6 enabled).

  - To enter an IPv4 address, use the dotted decimal format: *nnn.nnn.nnn.nnn*. For example: 127.0.0.1.

- When entering an IPv6 address you must include brackets [] to separate those addresses from their port numbers when you enter a value in this field. For example: **[2001:DB8::1234:0000:0000:5678:ABCD]:9550**. Use one of the following methods to enter an IPv6 address:
  - 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:DB8:0000:1234:0000:0000:5678:ABCD
  - You can specify an IPv6 address using a shortened format that omits leading zeros: (2001:DB8:0:1234:0:0:5678:ABCD)
  - You can use double colons in place of a series of zeros: (2001:DB8:0000:1234::5678:ABCD)

Setting up a user role and collecting system statistics

If you are logging on to the Tivoli Storage Productivity Center graphical user interface (GUI) for the first time after installing the product, note the following information:

- If you did not set up a Tivoli Storage Productivity Center user role, you might not be able to log on. To set up a Tivoli Storage Productivity Center user role, log on as the superuser and create the roles as described in the IBM Tivoli Storage Productivity Center Information Center .
- After you log on, you will see a panel that prompts you to collect system statistics. Click **Yes**.

Single sign-on

Tivoli Storage Productivity Center can use the single sign-on feature, which enables you to start Tivoli Storage Productivity Center from either the Tivoli Integrated Portal V1.1.1.0 instance that is installed with Tivoli Storage Productivity Center or from an existing Tivoli Integrated Portal V1.1.1.0 instance. If you have a single sign-on environment configured and Tivoli Integrated Portal user authentication is successful, the Tivoli Storage Productivity Center logon window does not open when you start the Tivoli Storage Productivity Center GUI from Tivoli Integrated Portal.

User credential considerations for single sign-on

When you log on to either Tivoli Integrated Portal or Tivoli Storage Productivity Center, a Lightweight Third-Party Authentication (LTPA) token is created and used for single sign-on authentication when you start other applications from within the Tivoli Storage Productivity Center GUI.

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

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

Stopping the Tivoli Storage Productivity Center GUI

This section describes how to stop the IBM Tivoli Storage Productivity Center graphical user interface (GUI).

To stop the Tivoli Storage Productivity Center user interface, click **File** → **Exit**.

Single sign-on

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 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 in context feature to enable you to move smoothly from one application to a functionally-related location in a second application.

Note that single sign-on is not available for the IBM XIV Storage System element manager, XIV Storage Manager.

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's 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's information. The authenticating service uses the corresponding private keys to decrypt the user's 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.

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

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

During the installation 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 graphical user interface (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.

Related concepts

"Changing the user authentication method" on page 157
The IBM Tivoli Storage Productivity Center installation program enables you to select a user authentication method that is used by Tivoli Storage Productivity Center, IBM Tivoli Storage Productivity Center for Replication, and IBM Tivoli Integrated Portal. You can choose to authenticate users against the users defined for the local operating system or to authenticate users against the users defined in a Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory repository. You can change the user authentication method after installation using Tivoli Integrated Portal.

Related tasks

"Changing the LTPA token expiration for single-sign on" on page 166
A single sign-on environment requires a centralized authentication repository that is accessed by all applications within the environment. The user's authentication information is passed between applications using Lightweight Third-Party Authentication (LTPA) tokens. You can change the expiration time for the LTPA tokens using IBM Tivoli Integrated Portal.

User roles

IBM Tivoli Storage Productivity Center provides user roles that can be mapped to user groups contained in the authentication repository, which is either the local operating system or an LDAP-compliant directory, depending on the choice of authentication mechanism made during the Tivoli Storage Productivity Center installation. The roles determine the user's ability to see and act upon specific nodes in the navigation tree of the GUI. Other factors that affect the user's capability are the components that are licensed, installed, and running.

The superuser is the highest level role and can be used to install IBM Tivoli Storage Productivity Center and to initially set up the product. The superuser has authority over all the other roles in Tivoli Storage Productivity Center:

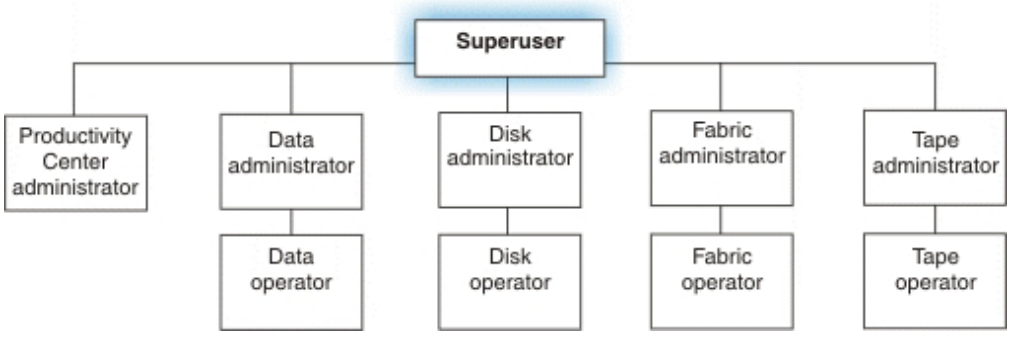

Figure 3. User roles in Tivoli Storage Productivity Center

The functions that are visible for the user roles require that the following components be licensed and installed as shown in Table 4.

Table 4. Functions that require components to be licensed and installed

Function	Component
Data Manager	IBM Tivoli Storage Productivity Center for Data
Disk Manager (requires Device server installation)	IBM Tivoli Storage Productivity Center for Data or Tivoli Storage Productivity Center for Disk or IBM Tivoli Storage Productivity Center for Fabric

Table 4. Functions that require components to be licensed and installed (continued)

Function	Component
Disk performance function (requires Device server installation)	Tivoli Storage Productivity Center for Disk
Fabric Manager (requires Device server installation)	IBM Tivoli Storage Productivity Center for Fabric or IBM Tivoli Storage Productivity Center for Data or Tivoli Storage Productivity Center for Disk
Fabric performance function (requires Device server installation)	IBM Tivoli Storage Productivity Center for Fabric
Tape Manager (requires Device server installation)	IBM Tivoli Storage Productivity Center for Fabric or IBM Tivoli Storage Productivity Center for Data or Tivoli Storage Productivity Center for Disk
Element Manager	IBM Tivoli Storage Productivity Center Basic Edition or IBM Tivoli Storage Productivity Center Standard Edition

Table 5 shows the GUI nodes that require special roles to view, edit, and act upon the services. The Productivity Center administrator or superuser is required for viewing configurations.

Table 5. GUI nodes that require special roles or licenses

GUI node	Special roles
Administrative Services	All roles can see this node but the user must be a Productivity Center administrator or superuser to edit and act upon the services (for example enable or disable tracing, shut down services, delete services, and so forth).
Administrative Services - Configuration - Role-to-Group Mappings	Only the Productivity Center administrator or superuser can see this node.
Administrative Services - Configuration - License Keys	Only the Productivity Center administrator or superuser can see this node.
Tivoli Storage Productivity Center - My Reports - System Reports - Data	Requires the data administrator role to create or edit objects, or to run jobs, (for example, pings and scans). Requires the data operator role to view reports.
Tivoli Storage Productivity Center - My Reports - System Reports - Disk	Requires that the disk performance function be installed and the disk operator role.
Tivoli Storage Productivity Center - My Reports - System Reports - Fabric	Requires the fabric administrator role to create or edit objects, or to run jobs (for example, pings and scans). Requires the fabric operator role to view reports. This includes being able to view the fabric assets, port connections, SAN switch report, switch port errors, and switch ports (packet rate and data rate).

Table 5. GUI nodes that require special roles or licenses (continued)

GUI node	Special roles
Tivoli Storage Productivity Center - My Reports - Batch Reports	Requires any administrator role to create the batch reports.
Tivoli Storage Productivity Center - Monitoring - Probes	Requires any administrator role to create or edit probes.
Tivoli Storage Productivity Center - Alerting - Alert log	Requires any administrator role to delete or clear alerts. The Storage subsystems node appears if any disk array is installed.
Data Manager	Requires the data administrator role to create or edit objects, or to run jobs (for example, pings and scans). Requires the data operator role to view reports.
Data Manager for Databases	Requires the data administrator role to create or edit objects, or to run jobs (for example, pings and scans). Requires the data operator role to view reports.
Data Manager for Chargeback	Requires the data administrator or data operator role.
Disk Manager	Requires the disk administrator role to create or edit objects, or to run jobs (for example, pings and scans). Requires the disk operator role to view reports. The Storage Subsystem Performance Monitors node requires that the Disk performance function is available.
Disk Manager - Monitoring - Subsystem Performance Monitors - Profile Management - Reporting - Storage Subsystem Performance	Disk performance function must be available.
Fabric Manager	Requires the fabric administrator role to create or edit objects, or to run jobs (for example, pings and scans). Requires the fabric operator role to view reports.
Fabric Manager - Monitoring - Switch Performance Monitors - Reporting - Switch Performance - Constraint Violations	Fabric performance function must be available.
Tape Manager	Requires the tape administrator role to create or edit objects, or to run jobs (for example, pings and scans). Requires the tape operator role to view reports.
Element Manager	All roles can view and edit objects and run jobs.

User tasks

This section provides an overview of the key IBM Tivoli Storage Productivity Center tasks that you can perform. It includes tasks that are common to all configurations and tasks that require that specific components be installed and licensed.

Common tasks

These tasks can be performed on any IBM Tivoli Storage Productivity Center configuration.

- Authorize users by associating user groups with roles (requires superuser role).
- View and manage Data server services.
- Perform agent probes
- Perform discoveries
- Perform configuration tasks under the **Administrative Services** node in the graphical user interface.
- Check the alert log to see detailed information about alerts.
- View the topology to determine the status of the resources and detailed information about them.
- Customize topology views.
- Launch element managers from the topology view.
- Use Tape Manager to view tape libraries, create tape library groups and run probes against those groups, and view tape library alerts.

Data Manager tasks

Data Manager tasks are tasks specific to the Data Manager component.

The Data Manager tasks include the following:

- Perform agent discovery.
- Perform CIMOM discovery.
- Perform out-of-band discovery.
- Deploy Data agents or Storage Resource agents on the hosts that you want to monitor.
- Create profiles to define file systems to be scanned and statistics to be gathered.
- Create monitoring groups for computers, file systems, directories, users, OS user groups, database tablespaces, and tables.
- Create and work with alerts for computers, file systems, and directories, and database instances, tablespaces, and tables.
- Establish quotas and constraints to define storage use limits for users and user groups.
- Automatically extend local file systems of managed hosts.
- Collect information about the availability of storage resources.
- Collect storage usage and trending information.
- Create user and file system groups for generating predefined reports.
- Generate and view data system reports.
- Generate invoices that charge back for storage usage.

Disk Manager tasks

The Disk Manager tasks are tasks specific to theDisk Manager component.

The Disk Manager tasks include the following:

- View and manage Device server services.
- Performing CIMOM discovery.
- Define and launch an element manager for storage subsystems.
- Create and manage virtual disks and volumes.
- Use the SAN Planner facility to analyze Tivoli Storage Enterprise Storage Server subsystems performance data.
- Create workload profiles to be used by the Volume Planner to describe characteristics of the volumes to be created.
- Create provisioning profiles to be used by the SAN Planner to define the actual capacity (and security - zoning) characteristics of whatever the profile happens to be associated with.
- Create storage subsystem groups.
- Create and work with storage subsystem device alerts.
- Monitor storage subsystem performance.
- Generate and view disk system reports and storage subsystem performance reports.

Fabric Manager tasks

The Fabric Manager tasks are tasks specific to the Fabric Manager component.

The Fabric Manager tasks include the following:

- View and manage Device server services
- Install, configure, and add Fabric agents
- Perform Fabric agent in-band discovery
- Perform CIMOM discovery
- Perform out-of-band discovery.
- Create and work with fabric, switch, and endpoint device alerts
- Monitor switch performance
- Generate and view fabric system reports and switch performance reports
- Create and work with fabric zones, zone aliases, and zone sets

Data collection

Determining what information you want to gather about the storage resources within your enterprise is critical to helping you implement a storage management strategy. Once you have determined the types of information you want to collect about your storage resources, define and schedule monitoring jobs to collect that information.

The data and statistics gathered by monitoring jobs are stored in the database repository. You must define and run monitoring jobs before you can perform the other tasks within IBM Tivoli Storage Productivity Center, such as viewing reports, enforcing storage policies, setting alerts, using the topology viewer, and managing storage subsystems.

The following types of data collection jobs are available in Tivoli Storage Productivity Center:

Availability Monitoring
Use pings to collect information about the availability of the storage assets in your enterprise. You can run pings against servers and computers only. Pings are not available for storage subsystems and other devices.

Asset Collection
Use probes to collect statistics about the storage assets in your enterprise, such as computers, disk controllers, hard disks, filesystems, and RDBMS objects.

Usage Monitoring
Use scans to collect statistics about the usage and trends of your actual storage consumption. Scans perform the majority of the work for Data Manager by providing all the data for usage reporting, as well as for quota and constraint analysis. They gather information at the file and directory level. You can run scans against servers and computers only. Scans are not available for storage subsystems and other devices.

Switch Performance Monitoring
Use fabric probes and switch performance monitoring jobs to collect information about switch performance.

Tivoli Storage Productivity Center Server Monitoring
Use TPC Server probes to collect the storage asset and capacity information that has been gathered by other Tivoli Storage Productivity Center servers in your environment. Creating master and subordinate relationships between among servers enables you to use a single interface to generate reports based on data and metrics collected by multiple servers in a storage environment.

You can create groups and profiles to use with monitoring jobs to indicate more efficiently the storage resources to monitor and the types of statistics to collect.

Monitoring Groups
Use monitoring groups to ping, probe, or scan multiple storage resources using the same job. You can create named groups of computers, storage subsystems, filesystems, directories, databases, tablespaces, tables, users, and OS user groups.

For example, to gather information about all of the computers in a payroll department, create a computer monitoring group named Payroll and populate that group with the computers from the payroll department. Then, to collect information about those computers, simply select the computer monitoring group named Payroll when defining a monitoring job. As computers are added or removed from the payroll department, simply update the Payroll computer monitoring group with those changes. Any monitoring job that uses Payroll will automatically collect information about the new computers in the group during its next run.

Profiles
Use profiles to fine tune and control the files to be scanned during a file system scan and to specify the statistical information to be gathered.

Default profiles are supplied with Data Manager with Data Manager for Databases. With the default profiles you can select a specific statistic that you want to gather. If you want to scan a computer or file system and gather all the available statistical information, select all of the default profiles when you create your scan.

You can also use profiles with **Reporting** → **Usage** reports to determine what information is displayed in those reports.

Alerts

You can set up IBM Tivoli Storage Productivity Center so that it examines the data that it collects about your storage infrastructure and writes an alert to a log when an event occurs. You also can specify that an action be initiated, such as sending an SMNP trap, sending an e-mail, or running a script when the event occurs.

Alerts are triggered based on the data collected by monitoring jobs (pings, scans, and probes), so the alerts must be defined before the monitoring jobs are run. For each alert, you select a condition that triggers the alert and (optionally) an action to be performed when that condition occurs.

You can define an alert in the following ways:

- As part of a data collection job
- As a separate, alerting job

When an event occurs and triggers an alert, the alert is written to a log. You can also choose one or more other ways to be notified of the event. These alert notifications include SNMP traps, IBM Tivoli Enterprise Console events, Tivoli Storage Productivity Center login notifications, operating-system event logs, or e-mail.

Note: Alerts are not generated in a Tivoli Storage Productivity Center instance for actions that you perform from that instance. For example, if you start Tivoli Storage Productivity Center and use Disk Manager to assign or unassign volumes for a subsystem, you will not receive alerts for those volume changes in that instance. However, if you assign and unassign volumes outside of that Tivoli Storage Productivity Center instance, an alert is generated.

Triggering conditions for alerts

In general, the following types of conditions can trigger alerts:

- A data collection job failed to complete.
- A change occurred in the storage infrastructure.
- A performance threshold was violated.

The specific conditions that can trigger events vary, depending on the type of storage resource that is being monitored. See Appendix E, “Triggering conditions for alerts,” on page 791 for more information on triggering conditions.

Alert notifications and triggered actions

Notification methods, or triggered actions, define the method by which you should be notified of an alert: SNMP Trap, Tivoli Enterprise Console Event, Login Notification, Windows Event Log, UNIX Syslog, and e-mail. Alerts are always written to the error log. Additionally, if the managed system is running the Data agent, you can run a script or start an IBM Tivoli Storage Manager job in response to the alert. See Appendix F, “Triggered actions for alerts,” on page 801 for more information on triggered actions.

Prerequisites for successfully using alerts

The following conditions must be met in order to successfully use alerts:

- Data collection jobs are configured and scheduled to run on a regular basis.

- If you want to be notified about an alert in some way other than an entry in the log file, such as using SNMP traps, Tivoli Enterprise Console events, or e-mail, alert notification instructions must be configured prior to using the alert. See "Configuring alert destinations" on page 256).

Default alerts

The following types of alerts are set for your system by default:
- New entity found
- Job failed

You can change the triggered actions and enable or disable these alerts, but you cannot delete them.

Provisioning planner

Use the IBM Tivoli Storage Productivity Center *host planner*, the *storage planner*, and the *security planner* to provide guidance when you provision your storage system. The collective term for all three individual planners working together is the *provisioning planner*.

For example, to add more storage to a host, the three planners work together to create the storage, zone the new storage to the host, and configure the multipath settings on the host. After using the planners to perform these tasks, you can continue to perform typical Tivoli Storage Productivity Center jobs with the storage on the host, such as creating the file system.

You can call each planner individually. For example, you might modify an existing multipath or zone setting or create volumes on the subsystem for later use with the host planner.

Use the planners to plan an integrated set of high-level policies for performance, availability, and security, based on the current resource utilization. The set of selected policies determines the best action responses to performance, availability, and security events and constitutes the output plan, which is the record of the output of the planners. You can either accept plan advice to trigger an associated action (like creating a volume or performing zoning), or reject the advice and change the plan.

Modes of use

You can use the planners in the following ways:

Planning mode (Host planning advice)
Planning mode uses each of the planners: storage planner, provisioning planner, and security planner, to configure an end-to-end plan for storage connectivity. This results in a recommendation for zoning, volume assignment, and multipath software configuration.

Assist mode
Assist mode provides zoning recommendation that meet the connectivity requirements, as established in current policies.

Validate mode
Validate mode evaluates the current configuration against the connectivity requirements, using the configuration checker utility.

Host planner

Use the IBM Tivoli Storage Productivity Center *host planner* to provide guidance when you configure host to subsystem connectivity. The host planner accesses default requirements or requirements that you set, and helps identify configurations that do not meet those requirements.

You can perform common network and storage provisioning tasks all at once by using the host planner with the other two planners, which are the *security planner* and the *storage planner*. The collective term for the three individual planners working together is *provisioning planner*.

For example, to add more storage to a host, the three planners work together to create the storage, zone the new storage to the host, and configure the multipath settings on the host. After using the planners to perform these tasks, you can continue to perform typical Tivoli Storage Productivity Center jobs with the storage on the host, such as creating the file system.

Host planning is the task of determining the connectivity between a host and the storage subsystems to obtain pathing and performance requirements. Inputs to the planning process are:

- The host
- Host ports and storage controller ports (and port performance utilization)
- Existing multipathing and zoning information
- The subsystem and identifiers
- The application profile (I/O, performance, and availability requirements)
- Pathing requirements (whether to use multipathing if it is available)

The planning process uses storage planning, configuration checking, and security planning for portions of the overall process.

Collection of information

The planner retrieves information from the Tivoli Storage Productivity Center database as its default input. You can provide explicit values to override default input. If you explicitly specify the number of paths, for example, the host planner verifies whether the value satisfies performance requirements. If not, the host planner calculates the number of paths, and which storage controller and host ports satisfy the requirements.

The fabric, disk and data discovery and probe operations must run to collect the information needed for the planner. If the information is not collected, the planner output identifies the limitation and presents a recommendation based on the subset of information that is available.

Required components

The host planner also uses the following components:

Data agent

The Tivoli Storage Productivity Center data agent is enhanced to deal with an SD multipathing driver on the host.

The enhanced data agent also collects SDD information, including SDD present, version, HBAs and ports, HBA info, vdisk and vpath configuration, and multipathing settings per vpath.

Multipathing control action application programming interfaces (APIs)
Unlike the storage planner or the security planner, control action APIs for multipathing do not exist.

Overview of the planning process

Host planning requires the following steps:

1. Validating that the host and subsystem are compatible by type and level of the host, the operating system, the subsystem, and the SDD software, for example
2. Determining the number of paths needed for the pathing policy and the expected I/O performance
3. Determining the connectivity required between the host and the subsystem
4. Determining the configuration for multipathing software
5. Configuring the zoning and SDD multipathing configuration

Policies related to the host planner

The host planner uses the following policies that you configure with the configuration checker utility:

Multipath
Specifies the multipath policy requirement (multiple or single path)

Preferred number of paths
Specifies required number of paths between host and subsystems

Multipath driver mode
Specifies multipath driver mode (fail-over, load balancing, round robin)

Redundant fabric
Specifies routing through isolated fabrics (must have at least 2 routes in separate fabrics, only 1 route required)

Planning guidance outputs

The host planner provides guidance that is based on the pathing and performance requirement inputs that it accesses.

Outputs from the planning process are:

- Validation that a plan is possible
- Zoning
- Subsystem assignment
- Multipathing (SDD) configuration information
- Configuration of zoning and SDD

Indication of violations

The planner notifies you of current policy violations before running a new plan that compensates for the violations. Violations are logged in the job log.

Command-line interface

There is no command-line interface for invoking the planning wizard.

Storage planner

Use the IBM Tivoli Storage Productivity Center *storage planner* to perform storage administration tasks, such as creating storage.

You can perform common network and storage provisioning tasks all at once by using the storage planner with the other two planners, which are the *host planner* and the *security planner*. The collective term for the three individual planners working together is *provisioning planner*.

For example, to add more storage to a host, the three planners work together to create the storage, zone the new storage to the host, and configure the multipath settings on the host. After using the planners to perform these tasks, you can continue to perform typical Tivoli Storage Productivity Center jobs with the storage on the host, such as creating the file system.

Collection of information

The storage planner collects and evaluates capacity, storage controller preferences, storage controller types, the number of volumes and volume size, performance requirements, the use of unassigned volumes, the RAID level and performance utilization and capacity utilization input from the Tivoli Storage Productivity Center database and from explicit values that you define.

Planning guidance outputs

The storage planner generates a plan that presents optimal storage controller and storage pool configurations.

Control action application programming interfaces (APIs)

The control action APIs are not implemented by the storage planner but are implemented by the device manager (DM).

Overview of the planning process

If you explicitly specify the storage pool and controller information, the storage planner checks to see whether the input performance and capacity requirements can be satisfied.

Supported devices

The storage planner provides support for IBM System Storage DS8000, IBM System Storage DS6000, and IBM Enterprise Storage (ES800 controllers). The storage planner is capable of working with any combination of these specific controllers.

Security planner

Use the IBM Tivoli Storage Productivity Center *security planner* to perform security administration tasks.

You can perform common network and storage provisioning tasks all at once by using the security planner with the other two planners, which are the *host planner* and the *storage planner*. The collective term for the three individual planners working together is *provisioning planner*.

For example, to add more storage to a host, the three planners work together to create the storage, zone the new storage to the host, and configure the multipath

settings on the host. After using the planners to perform these tasks, you can continue to perform typical Tivoli Storage Productivity Center jobs with the storage on the host, such as creating the file system.

Collection of information

The security planner collects and evaluates existing host port information, storage controller ports, zones, zone sets, and switches from the Tivoli Storage Productivity Center database and from explicit values that you define as zoning input and volume masking input.

Planning guidance outputs

The security planner generates a plan that presents optimal zoning and volume masking.

Control action application programming interfaces (APIs)

The security planner control APIs are currently already implemented by the fabric and DM APIs.

Overview of the planning process

If you specify exact zoning and volume masking information details, the planner informs you whether your input is different than the existing configuration, or different than the optimal configuration that the planner calculates based on policies and the existing security parameters. If the optimal configuration differs from your selections, the planner proposes zoning and volume masking settings.

Reporting

This section contains conceptual information about reporting.

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

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

Fabric and zones

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

Zones, Zone Aliases, and Zone sets

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

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

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

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

Zone membership can be specified by:

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

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

There are two types of zoning:

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

Software zoning
: Software zoning uses the Simple Name Server (SNS) that runs inside a fabric switch. It is based on the node WWN or port WWN of the zone members to be included. Software zoning lets you create symbolic names for the zones and zone members.

For zone control of Brocade switches, Fabric Manager preserves the FC switch traffic through the use of *temporary zone sets*. A temporary zone set is used when a change is made to the zones of the active zone set. The active zone set of the current zone definition is copied to a temporary zone set, and this temporary zone set is activated. Then the previously active zone set is deleted, and the desired active zone set is created and activated. Finally, the temporary zone set is deleted. If creation or activation of the desired active zone set fails, the temporary zone set

remains in place to preserve FC switch traffic. The user can identify the temporary zone set by the name 'tmp*xxxx*', where '*xxxx*' is a randomly generated four-digit number.

Note: Temporary zone sets are not used if the active name of the zone set changes.

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

Note:

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

You can view zones in SANs using the topology viewer console. Fabric Manager lets you create, update, and delete zones, zone aliases, and zone sets in a SAN. In addition, you can do the following:

- Assign zones to zone sets
- Activate and deactivate zone sets
- View zone membership
- Add and remove zone members to and from a zone
- Add and remove zone members to and from a zone alias
- Assign zone aliases to zones

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

Zone configuration

When you configure zones, you can use Fabric Manager or the management application for the devices. The advantages to using Fabric Manager rather than the management applications of the devices are that Fabric Manager lets you perform zoning from a single interface, and you can use consistent, familiar methods to work with devices from multiple vendors.

Consider the following guidelines if you use Fabric Manager with zones:

- For Brocade and McDATA fabrics, topology information is collected through CIM agents. If no CIM agents are configured, or if the fabric is a QLogic or Cisco fabric, have at least one managed host per zone for the complete topology display. In this way, you can manage your switches and a string is returned, provided by the vendor, that can be the worldwide name (WWN) or a vendor, model, and level type designation.

- If CIM agents cannot be used to collect zoning, the Fabric Manager can retrieve the zone information from IBM 2005 and 2109 switches and from Brocade Silkworm Fibre Channel Switches. SNMP agents must log into the switch with administrative rights. To see zone information, specify the login ID for the agents you define.

  For zone discovery and zone control for any fabrics with Brocade switches or vendor Brocade switches, you must have a SMI-S agent configured for the fabric. If a SMI-S agent is not configured, you must enable advanced Brocade discovery. To do this, enter the administrator user ID and password for at least one switch in a fabric. However, you should not enable advanced discovery for all Brocade switches that have been configured as SNMP agents. The enabled switch acts as a proxy and gathers zone information for the entire fabric. You can enable a second switch for redundancy. Enabling all Brocade switches puts an unnecessary load on the switches and fabric. You should enable newer, more powerful switches (such as director class) and those running the highest level of firmware.
- A zone set can be deleted while it contains zones. With QLogic switches, those zones are put into an orphan zone set. The zones can later be moved from an orphan set into other zone sets. The Fabric Manager zone configuration panels cannot display zones in an orphan zone set. However, the zones exist on the switch, and they can be configured using the switch vendor's tool.
- For McData switches, you can add empty zones to inactive zone sets. You cannot add empty zones to inactive zone sets for other vendors' switches.

  Note: Activating an inactive zone set that contains empty zones will fail if a switch does not support empty zones in active zone definitions.

Note: In-band discovery does not detect orphan zones (that is, zones that do not belong to any zone set). These zones are not listed in the Zone Configuration window. However, the zones exist on the switch, and they can be configured using the switch vendor's tool. However, Brocade orphan zones, which are discovered through SMI-S agents and out-of-band discovery, are listed in the Zone Configuration window and can be configured by the Fabric Manager. Brocade orphan zone aliases, which are discovered through SMI-S agents, are listed in the Zone Configuration window and can be configured by the Fabric Manager. Similarly, for McDATA fabrics managed where a CIM agent is configured, orphan zones are listed in the Zone Configuration window and can be configured by the Fabric Manager.

Switch zoning capabilities

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

To see a list of supported switches and their capabilities, default values, ranges, and possible effects, go to IBM Tivoli Storage Productivity Center for Fabric and select the Interoperability link.

Zoning for Cisco MDS 9000 devices:

Cisco switches support virtual SANs (VSANs), which is the logical partitioning of a fabric into multiple fabrics. The overall network is referred to as the physical infrastructure, and the logical fabrics are the VSANs. Fabric Manager provides

basic zone discovery and zone control for the Cisco MDS 9000 series of devices using the FC-GS-3 protocol. This allows in-band zone control and defaults to the Cisco VSANs.

The following tasks are supported:

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

Fabric Manager supports these zone member types:

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

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

Brocade switches in zones:

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

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

Note:

- A zone set displayed on the zone control interface might have no zones. This could mean that the zone set contains a non-standard zone on a Brocade switch.
- A zone set might have non-standard zones on a Brocade switch that are not displayed on the zone control interface.
- When creating a new zone using the same zone name as that of a non-standard zone that already exists on a Brocade switch, you might get a error that the zone already exists. Because the non-standard zone is not displayed on the zone control interface, you must verify this by using the switch management application.
- For zone control of Brocade switches, Fabric Manager preserves the FC switch traffic through the use of temporary zone sets.
- Tivoli Storage Productivity Center allows you to add empty zones to an inactive zone definition. Activation of a zone set containing empty zones may still fail if the switch does not support empty zones in active zone definitions.

Tivoli Common Agent Services

IBM Tivoli Storage Productivity Center uses Tivoli Common Agent Services for software distribution and desired state management. An *agent* is a program that automatically performs some service, such as data collection. To take advantage of some Tivoli Storage Productivity Center software management features, the Common agent must be installed on all managed endpoints.

The Common agent environment consists of the following:

Common agent
The Common agent is a common container for all the subagents to run within. It enables multiple management applications to share resources when managing a system. The Common agent is installed on each endpoint.

Agent Manager
The Agent Manager is the server component of the Common agent that provides functions that allow clients to get information about agents and resource managers. It enables secure connections between managed endpoints, maintains the database information about the endpoints and the software running on those endpoints, and processes queries against that database from resource managers. It also includes a registry and an agent recovery service.

Resource manager
A resource manager is the server component of a management application product that manages the agents. Examples of resource managers are Tivoli Storage Productivity Center (server component) and IBM Tivoli Storage Productivity Center for Data for Data (server component).

The Common agent

The Common agent is a common infrastructure in which product-specific agent code can be deployed and run. Examples of product–specific agents are Fabric Manager agent and Data Manager agent. A product-specific subagent consists of one or more *OSGi bundles*. A bundle is an application that is packaged in a format defined by the Open Services Gateway Initiative (OSGi) Service Platform specification, which is implemented in a lightweight infrastructure based on WebSphere Everywhere Deployment technology. The Common agent code will only be installed once on a managed host. For example, if you have two management applications on the same managed host (application A and application B), the Common agent code will only be installed once. However, there will be two product–specific agents: one for application A and one for application B. The term *agent* or *Common agent* is used to refer to both the Common agent and product–specific agent unless specifically stated otherwise.

The Common agent provides these services:

- Deployment and life cycle management of product agent bundles
- Agent registration and security certificate management
- Agent Manager queries
- Common agent health monitoring and configuration monitoring services

  The Common agent provides configuration information to the Agent Manager. Updates are initiated by the Common agent or product agent bundle events. There are times when management applications might require periodic updates of the configuration of product agents. The Common agent allows any product agent to participate and provide application specific status. The Common agent provides a registration interface that allows resource managers to receive agent configuration updates. The Agent Manager distributes the configuration monitoring reports to all interested parties.

  The monitoring service has a ″heartbeat″ function that allows for periodic updates of status reports. This function will force the monitoring service to send a status report after a period of time since the last status report delivery was attempted. The frequency of this update can be configured or completely disabled (the default frequency is 24 hours). You can change this frequency by editing the **status.heartbeat.frequency** property in the **endpoint.properties** file.

To disable this function, set the `status.heartbeat.frequency` property to zero. The **endpoint.properties** file is located in the `config` directory where you installed the agent.

The Common agent contacts the registration server and reports its status and any configuration changes at these times:

- After initial registration
- When a Common agent starts or stops
- After a configurable period of time
- Any time a bundle is installed, upgraded, or removed

Agent Manager

The Agent Manager is a network service that provides authentication and authorization using X.509 certificates and the Secure Sockets Layer (SSL) protocol. It also processes queries about its registry of configuration information about the agents and management applications (which are also called *resource managers*).

A resource manager is the server component of a management application product that manages the agents. Examples of resource managers are IBM Tivoli Storage Area Network Manager (server component) and IBM Tivoli Storage Area Network Manager (server component).

Resource managers and agents must register with the Agent Manager before they can use its services to communicate with each other. Registration is password protected, with separate passwords for the registration of agents and resource managers.

Related concepts

"Agent Manager registry"
The registry is a database that contains the current configurations of all known agents and resource managers. Resource managers and agents must register with the Agent Manager before they can use its services to communicate with each other. Registration is password protected, with separate passwords for the registration of agents and resource managers.

"Agent recovery service" on page 74
The Agent Manager provides an agent recovery service that is a network service for error logging for agents that cannot communicate with other Agent Manager services. Agents use an unsecured HTTP connection to communicate with the agent recovery service. Because the connection is unsecured, an agent can always communicate with the agent recovery service, even if the agent is incorrectly configured or has expired or revoked certificates. The agent recovery service is a WebSphere servlet container that runs on the Agent Manager server.

Agent Manager registry:

The registry is a database that contains the current configurations of all known agents and resource managers. Resource managers and agents must register with the Agent Manager before they can use its services to communicate with each other. Registration is password protected, with separate passwords for the registration of agents and resource managers.

The registry contains the identity, certificates, and communication information for each resource manager, and the following information about agents:

- The identity of every known agent and its computer system
- The certificate issued to each agent

- Basic configuration information about each agent, including information about the type and version of the hardware and operating system
- The configuration of each agent (updated by the agent at a configurable interval)
- The errors reported by each agent (updated by the agent at a configurable interval)
- Current communication parameters for the agent, including the IP address, the port or ports for which the agent is configured, and the supported protocol
- The agents on which each bundle is installed

The information in the registry is updated by asynchronous events, such as the registration of agents and resource managers, and by updates from the agent. The agent provides a configuration update when it starts, when a bundle is installed or uninstalled, and at a configurable interval (by default, daily). By default, the registry contains only the most recent configuration update and error information about each agent. However, the retention period for these records is configurable. For all other information, the registry contains the complete history of your agents and resource managers.

The registry can be placed in DB2 Universal Database or Oracle9i Database.

Agent recovery service:

The Agent Manager provides an agent recovery service that is a network service for error logging for agents that cannot communicate with other Agent Manager services. Agents use an unsecured HTTP connection to communicate with the agent recovery service. Because the connection is unsecured, an agent can always communicate with the agent recovery service, even if the agent is incorrectly configured or has expired or revoked certificates. The agent recovery service is a WebSphere servlet container that runs on the Agent Manager server.

Agents locate the agent recovery service using the unqualified host name `TivoliAgentRecovery`. Your Domain Name System (DNS) server must map the host name `TivoliAgentRecovery` to the computer system where you installed the Agent Manager. The normal DNS lookup mechanism iterates through the domain search list for the agent, appends each domain in the list to the unqualified host name, and then performs a DNS lookup to attempt to resolve the name. The agent recovery service listens for recovery requests on two ports: port 80 and a configurable port (by default, 9513). Using port 80 makes the request more likely to pass through a firewall between the agent and the agent recovery service. However, if the Agent Manager is on the same system as the HTTP server, port 80 is not available. The configurable second port provides an alternate communication port, in case you need to disable the use of port 80 by the agent recovery service.

Agents

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

Common Information Model agents

Common Information Model (CIM) agents enables communication between the storage device (storage subsystem, fabric switch, or tape library) and IBM Tivoli Storage Productivity Center. Each storage device that you want to manage must have a CIM agent either installed or embedded.

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

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

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

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

Note:

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

Data and Storage Resource agents

Data agents and Storage Resource agents collect information from the systems (for example, servers, computers, and workstations) in the environment. You must install one of these agents on each system that you want IBM Tivoli Storage Productivity Center to monitor.

Data agents gather the following information:

- Asset information
- File and file system attributes
- Database application information
- Novell Directory Services (NDS) tree information
- Network attached storage (NAS) device information

Storage Resource agents gather the following information:

- Asset information
- File and file system attributes

Using Data agents or Storage Resource agents:

Use a Data agent or a Storage Resource agent to collect storage data about a computer. The storage entity that you want to monitor and the type of information you want to collect can help determine which of these agents to install on a computer.

The following table shows the different features in IBM Tivoli Storage Productivity Center that are available for computers on which Storage Resource agents or Data agents are installed. Use this table to help determine which type of agent best meets your storage monitoring needs.

Table 6. Available features for Data agents and Storage Resource agents

Feature	Data agent	Storage Resource agent
Is Agent Manager required?	Yes	No
How do you install these agents?	Use the Tivoli Storage Productivity Center installation program to install Data agents on computers. This program is included on the disks provided with the product. See the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for more information about how to run the installation program.	To install Storage Resource agents, expand **Administrative Services** > **Configuration**. Right-click **Storage Resource Agent Deployments**. Click **Create Storage Resource Agent Deployments**. You do not need to run the installation program to install Storage Resource agents.
Is the Java Runtime Environment (JRE) required on the computer where the agent is installed?	Yes	No
Is there a list of hardware requirements and operating systems on which these agents can run?	Yes See the *Planning > Planning for the IBM Tivoli Storage Productivity Center family > Hardware and software requirements for the IBM Tivoli Storage Productivity Center family* topic in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* to view the hardware requirements and a list of operating systems supported for Data agents.	Yes See the *Planning > Planning for the IBM Tivoli Storage Productivity Center family > Hardware and software requirements for the IBM Tivoli Storage Productivity Center family* topic in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* to view the hardware requirements and a list of operating systems supported for Storage Resource agents.
What data collection jobs can you run against the computers on which the agents are installed?	Probes, pings, scans	Probes (host systems only)
Can you include computers and file systems that are monitored by agents in monitoring groups?	Yes	No, you cannot include computers and file systems that are monitored by Storage Resource agents in monitoring groups.
What reports show the data that is collected by the agents?	See "Reports for Data agents and Storage Resource agents" to view a list of the reports that show data collected by Data agents.	See "Reports for Data agents and Storage Resource agents" to view a list of the reports that show data collected by Storage Resource agents.

Reports for Data agents and Storage Resource agents:

You can collect data about the computers on which Data agents and Storage Resource agents are installed. The type of agent determines what kind of information you can collect and which reports display that information.

Use this table to learn:

- The reports that reflect the information collected by data collection jobs

- The data collection jobs that you can run against each agent type

Table 7. Available data collection jobs and reports for Data agents and Storage Resource agents

Reports	Data collection job	Data agent	Storage Resource agent
Batch reports	Probes, Scans, Pings	Yes	No
System Reports: • Data • Fabric • Disk	 Probes, Scans Probes Probes	 Yes No Yes	 Yes, but not all [1] No No
Asset Reports: • By Cluster • By Computer • By Hypervisor • By OS Type • By Storage Sybsystem • System-wide	 Probes Probes Probes Probes Probes Probes, Scans	 Yes Yes Yes Yes Yes Yes, all system-wide reports	 Yes Yes No Yes No Yes, but not all [2]
Availability Reports: • Ping reports • Computer Uptime reports	 Pings Probes	 Yes Yes	 No Yes [3]
Capacity reports	Probes	Yes	Yes
Usage reports	Scans	Yes	No
Usage Violation reports	Scans	Yes	No
Backup reports	Scans	Yes	No
Data Manager for Databases reports	Probes, Scans	Yes	No
Storage Subsystem reports	Probes	Yes	Yes
Storage Subsystem Performance reports	Subsystem Performance Monitors	Yes	No
Rollup Reports: • Asset • Database Asset • Capacity • Database Capacity	TPC server probes [4]	Yes Yes Yes Yes Yes	Yes Yes No Yes No

Notes:

1. Only the following Data system reports are available for Storage Resource agents: Disk Capacity Summary, Disk Defects, Storage Capacity, Total Freespace
2. Only the following System-wide Asset reports are available for Storage Resource agents: Agents, Computers, Disk/Volume Groups, Disks, File Systems or Logical Volumes, Volumes, Exports or Shares.
3. Computer Uptime reports are not available for Storage Resource agents that are invoked using the non-daemon protocol for data collection. These reports do not contain data for computers on which non-daemon based Storage Resource agents are deployed. To collect computer uptime information using a Storage Resource agent, that agent must use a daemon service for runtime operation.
4. Run IBM Tivoli Storage Productivity Center server probes to collect information from subordinate servers for display in the master server's rollup reports. You cannot use Storage Resource agents to collect information about relational databases, so the Database Asset and Database Capacity reports will not display data for those agent types.

Deployment considerations for Storage Resource agents:

There are a number of guidelines you must take into consideration when deploying and working with Storage Resource agents in your environment.

Required authority for deploying Storage Resource agents

You must be logged in to IBM Tivoli Storage Productivity Center with a user ID that has the superuser role to schedule Storage Resource agent deployments. See the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for more information about user roles.

Supported operating systems

Storage Resource agents are not deployable on all the operating systems where you can deploy Data agents. See "Using Data agents or Storage Resource agents" on page 75 to view the operating systems that are supported by Storage Resource agents.

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 probe job 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 or 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 an on-demand service (non-daemon service) runs as a stand-alone executable file on the monitored computer. Connectivity between the server and agent uses the same protocol that was used during the deployment of the agent.

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 running 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 install a Storage Resource agent on a computer where a Data agent is already installed *and* pointing to the same Data server as that Storage Resource agent. For example, if you install a Storage Resource agent on a computer, and then later install a Data agent that points to the same Data server as that Storage Resource agent, the Storage Resource agent is automatically uninstalled after the Data agent runs a successful probe job.

You can install a Storage Resource agent and a Data agent on the same computer if those agents communicate with different Data servers.

Storage Resource agents on computers running Windows 2008
Before you can deploy a Storage Resource agent on a computer that is running the Windows 2008 operating system, 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.

Data collection and Storage Resource agents
You can only run probe jobs against Storage Resource agents. You cannot run other data collection jobs against them such as scans, pings, and performance monitors. See "Reports for Data agents and Storage Resource agents" on page 76 for a list of reports that show the data collected by Storage Resource agents.

Upgrades and job logs
If you deploy a Storage Resource agent and later decide to upgrade it to a Data agent, the job logs for the probes that are run by the Storage Resource agent are removed from the computer and you are unable to access those job logs through the user interface. The job logs for probes that are run by the Data agent are retained.

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

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

Additional information about Storage Resource agents
See the *Planning > General Planning > Planning for Storage Resource agents* section in the *IBM Tivoli Storage Productivity Center Installation and Configuration Guide* for additional planning, configuration, and certificate information about Storage Resource agents.

In-band fabric agents

Fabric agents collect information about the Storage Attached Network (SAN) fabrics. You must install in-band fabric agents on systems with host bus adapters (HBAs) that are connected to the SAN fabrics that you want to manage.

The fabric agents communicate in-band through the HBA and collect the following information:

- Attribute and topology information
- Host-level information, such as file system and volume mappings
- HBA information, including make, model, and driver versions
- Port state information
- Zoning information (for QLogic and Cisco fabrics, and for McDATA fabrics if SMI-S agents are not installed)

In-band fabric agents are discovered during the installation process.

You can install the in-band fabric agent locally or you can deploy it remotely from the IBM Tivoli Storage Productivity Center server. Install an in-band fabric agent on *at least* one system that has an HBA that is connected to each switch that you want to manage; consider installing two in-band fabric agents for each switch to ensure redundancy.

Out-of-band fabric agents

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

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

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

IBM Tivoli Storage Productivity Center only supports SNMP V1.

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

Security

The IBM Tivoli Storage Productivity Center environment supports role-based security. Tivoli Storage Productivity Center also has several applications that use certificates to ensure server to server security.

For role-based authorization, user groups contained in the authentication repository, which is either the local operating system or an LDAP-compliant directory depending on the choice of authentication mechanism made during the Tivoli Storage Productivity Center installation, are associated with predefined roles. When a user ID is authenticated to Tivoli Storage Productivity Center through the GUI, CLI, or APIs, membership in a user group determines the authorization level.

For Agent Manager, the program comes with demonstration certificates you can use for a test environment, but you can also create new generated certificates during the installation of the Agent Manager. These generated certificates are used when the Tivoli Storage Productivity Center server or agent registers with the Agent Manager.

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 user's authorization level.

Table 8 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 8. 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. This includes reports on tape devices.
Fabric administrator	Has full access to Tivoli Storage Productivity Center for Fabric functions.
Fabric operator	Has access to reports only for Tivoli Storage Productivity Center for Fabric functions.

Table 8. Roles and authorization levels (continued)

Role	Authorization level
Data administrator	Has full access to Tivoli Storage Productivity Center for Data functions.
Data operator	Has access to reports only Tivoli Storage Productivity Center for Data 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:
   - 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 occurs, stop and restart the Tivoli Integrated Portal server.

Agent Manager certificates

The Agent Manager provides a demonstration certificate to facilitate the rapid deployment of a test environment or demonstration environment. The demonstration certificate and its password are publicly available. Anyone who has a demonstration certificate can participate in your deployment. Using the demonstration certificate does not provide the level of security required by a typical production environment. In a production environment, let the Agent Manager installation program generate a certificate that is unique to Agent Manager installation. When you install Agent Manager, you can specify to use the demonstration certificates or create a new certificate on the *Specify Certificates* page.

SAN Volume Controller topics

This section contains information about the following SAN Volume Controller topics: managed disks (MDisks), managed-disk states, managed-disk groups, and virtual disks (VDisks).

IBM Tivoli Storage Productivity Center SAN Volume Controller

The IBM Tivoli Storage Productivity Center SAN Volume Controller is a SAN appliance that attaches open-systems storage devices to supported open-systems hosts.

The SAN Volume Controller provides symmetric virtualization by creating a pool of managed disks from the attached storage subsystems, which are then mapped to a set of virtual disks for use by attached host computer systems. System administrators can view and access a common pool of storage on the SAN, which enables them to use storage resources more efficiently and provides a common base for advanced functions. The SAN Volume ControllerSAN Volume Controller performs the following functions for the SAN storage that it controls:

- Creates a single pool of storage
- Manages logical volumes
- Provides advanced functions for the SAN, such as large scalable cache, copy services, and space management

A *node* is a single storage engine. The storage engines are always installed in pairs with one or as many as four pairs of nodes constituting a *cluster*. Each node in a pair is configured to back up the other.

Each pair of nodes is known as an *I/O group*. All I/O operations are cached on both nodes for resilience. To eliminate any single point of failure, each of the two nodes are protected by different uninterruptible power supplies. A node can be in only one I/O group.

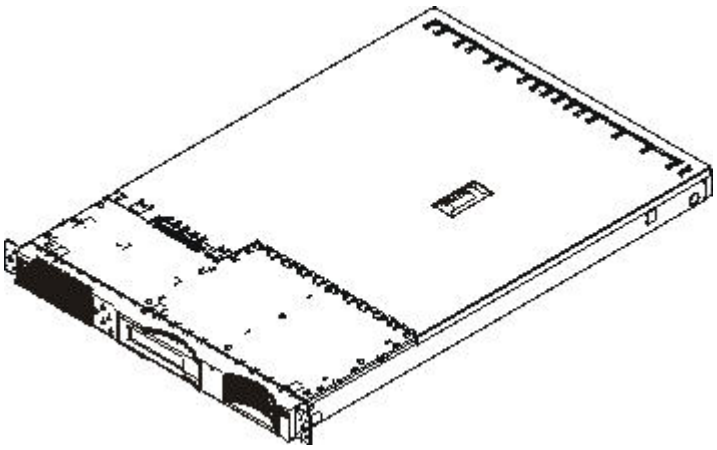

Figure 4. A SAN Volume Controller node

The SAN Volume Controller I/O groups see the storage presented to the SAN by the back-end controllers as a number of disks, known as *managed disks*. The application services do not see these managed disks. Instead, they see a number of logical disks, known as *virtual disks*, that are presented to the SAN by the SAN Volume Controller. Each node must only be in one I/O group and provides access to the virtual disks in the I/O group.

The fabric contains two distinct zones: a host zone and a disk zone. In the host zone, the host systems can identify and address the nodes. You can have more than one host zone. Generally, you create one host zone per operating system type. In the disk zone, the nodes can identify the disk drives (storage subsystems). Host systems cannot operate on the disk drives directly; all data transfer occurs through

the nodes. As shown in Figure 5, several host systems can be connected to a SAN fabric. A cluster of SAN Volume Controllers is connected to the same fabric and presents virtual disks to the host systems. You configure these virtual disks using the disks located on the RAID controllers.

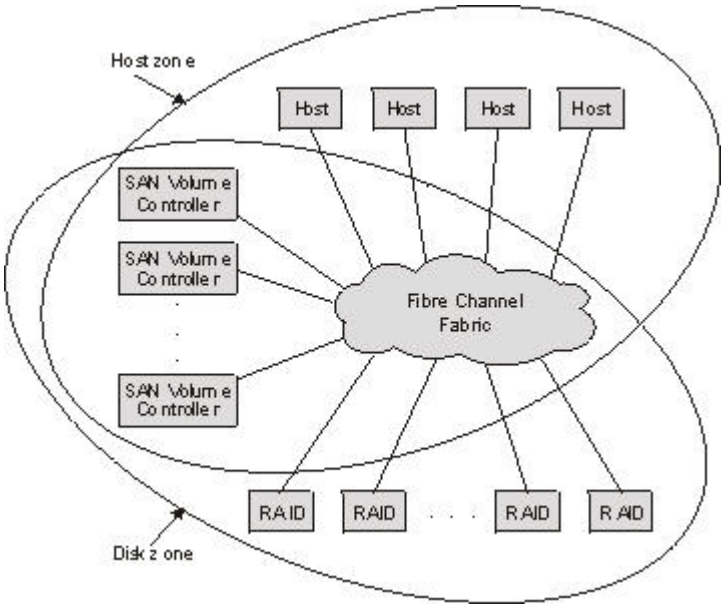

Figure 5. Example of a SAN Volume Controller in a fabric

Managed disks

A managed disk (MDisk) is a logical disk that a SAN Volume Controller has exported to the SAN fabric to which the nodes in the cluster are attached.

A managed disk might consist of multiple physical disks that are presented to the SAN as a single logical disk. A managed disk always provides usable blocks of physical storage to the cluster, even if it does not have a one-to-one correspondence with a physical disk.

Figure 6 shows physical disks and managed disks.

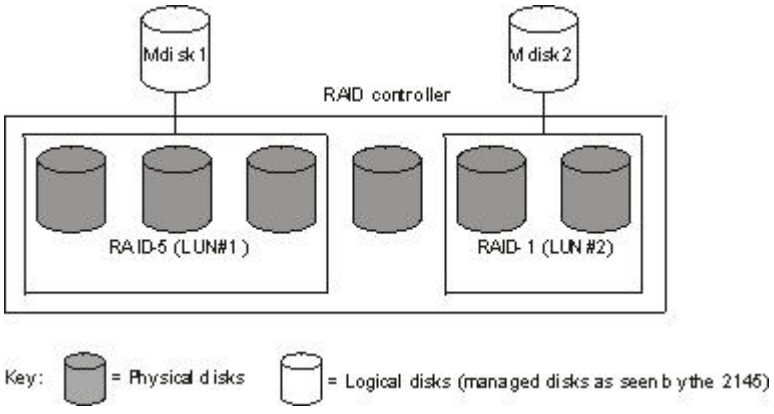

Figure 6. Physical disks and managed disks (MDisks)

Extents

Each MDisk is divided into chunks of equal size called *extents*. Extents are numbered sequentially (beginning at 0) from the start to the end of the managed disk. Extents are a unit of mapping the data between MDisks and virtual disks (VDisks).

The extent size is a property of managed-disk groups (MDisk groups). When a managed disk is added to a managed-disk group, the size of the extents that the MDisk will be broken into depends on the attribute of the managed-disk group to which it has been added.

Access modes

The access mode determines how the SAN Volume Controller cluster uses the MDisk. The possible modes are the following:

Unmanaged
The MDisk is not used by the cluster.

Managed
The MDisk is assigned to an MDisk group and provides extents that virtual disks (VDisks) can use.

Image The MDisk is assigned directly to a VDisk with a one-to-one mapping of extents between the MDisk and the VDisk.

Managed disk states

This topic describes the operational states of managed disks (MDisks).

Managed disks can be in the following states:

Online
The MDisk can be accessed by all online nodes. That is, all the nodes that are currently working members of the cluster can access this MDisk. The MDisk is online when the following conditions are met:

- All timeout error recovery procedures complete and report the disk as online.
- Volume inventory of the target ports correctly reported the MDisk.
- Discovery of this volume completed successfully.
- All of the managed disk target ports report this volume as available with no fault conditions.

Degraded
The MDisk cannot be accessed by all the online nodes. That is, one or more (but not all) of the nodes that are currently working members of the cluster cannot access this MDisk. The MDisk might be partially excluded; that is, some of the paths to the MDisk (but not all) have been excluded.

Excluded
The MDisk has been excluded from use by the cluster after repeated access errors.

Offline
The MDisk cannot be accessed by any of the online nodes. That is, all of the nodes that are currently working members of the cluster cannot access this MDisk. This state can be caused by a failure in the SAN, the storage

subsystem, or one or more physical disks connected to the storage subsystem. The MDisk will only be reported as offline if all paths to the disk fail.

Attention: If your fabric is undergoing transient link breaks or you have been replacing cables or connections in your fabric, you might see one or more MDisks change to the degraded status. If an I/O operation was attempted during the link breaks and the same I/O failed several times, the MDisk will be partially excluded and change to a status of degraded. You must include the MDisk to resolve the problem.

Managed disk path Each managed disk will have an online path count, which is the number of nodes that have access to that managed disk; this represents a summary of the I/O path status between the cluster nodes and the particular storage device. The maximum path count is the maximum number of paths that have been detected by the cluster at any point in the past. Thus if the current path count is not equal to the maximum path count then the particular managed disk may be degraded. That is, one or more nodes may not see the managed disk on the fabric.

Managed disk groups

A SAN Volume Controller *MDisk group* is a collection of MDisks that jointly contain all the data for a specified set of virtual disks (VDisks).

All MDisks in a group are split into extents of the same size. VDisks are created from the extents that are available in the group. You can add MDisks to an MDisk group at any time in order to increase the number of extents that are available for new VDisks or to expand existing VDisks. You can add only MDisks that are in unmanaged mode. When MDisks are added to a group, their mode changes from unmanaged to managed.

Figure 7 shows an MDisk group containing four MDisks.

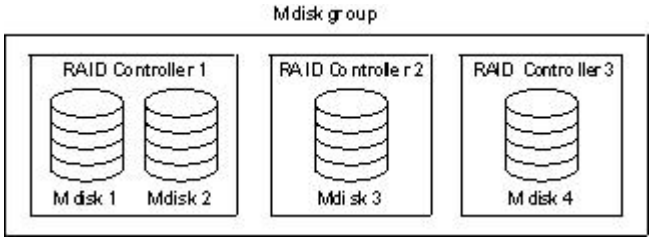

Figure 7. MDisk group

You can delete MDisks from a group under the following conditions:

- VDisks are not using any of the extents that are on the MDisk.
- Enough free extents are available elsewhere in the group to move any extents that are in use from this MDisk.

Attention: If you delete an MDisk group, you destroy all the VDisks that are made from the extents that are in the group. If the group is deleted, you cannot recover the mapping that existed between extents that are in the group and the extents that VDisks use. The MDisks that were in the group are returned to unmanaged mode and can be added to other groups. Because the deletion of a group can cause a loss of data, you must force the deletion if VDisks are associated with it.

Table 9 describes the operational states of an MDisk group.

Table 9. Managed disk group states

Status	Description
Online	The MDisk group is online and available. All the MDisks in the group are available.
Degraded	The MDisk group is available; however, one or more nodes cannot access all the MDisks in the group.
Offline	The MDisk group is offline and unavailable. No nodes in the cluster can access the MDisks. The most likely cause is that one or more MDisks are offline or excluded.

Attention: If a single MDisk in an MDisk group is offline and therefore cannot be seen by any of the online nodes in the cluster, then the MDisk group of which this MDisk is a member goes offline. This causes *all* the VDisks that are being presented by this MDisk group to go offline. Care should be taken when creating MDisk groups to ensure an optimal configuration.

Consider the following guidelines when you create MDisk groups:

- If you are creating image-mode VDisks, do not put all of these VDisks into one MDisk group because a single MDisk failure results in all of these VDisks going offline. Allocate your image-mode VDisks between your MDisk groups.
- Ensure that all MDisks that are allocated to a single MDisk group are the same RAID type. This ensures that a single failure of a physical disk in the storage subsystem does not take the entire group offline. For example, if you have three RAID-5 arrays in one group and add a non-RAID disk to this group, then you lose access to all the data striped across the group if the non-RAID disk fails. Similarly, for performance reasons you should not mix RAID types. The performance of all VDisks will be reduced to the lowest performer in the group.
- If you intend to keep the virtual disk allocation within the storage exported by a storage subsystem, ensure that the MDisk group that corresponds with a single subsystem is presented by that subsystem. This also enables nondisruptive migration of data from one subsystem to another subsystem and simplifies the decommissioning process if you want to decommission a controller at a later time.

Extents

To track the space that is available, the SAN Volume Controller divides each MDisk in an MDisk group into chunks of equal size. These chunks are called *extents*, and are indexed internally. Extent sizes can be 16, 32, 64, 128, 256, or 512 MB.

You must specify the extent size when you create a new MDisk group. You cannot change the extent size later; it must remain constant throughout the lifetime of the MDisk group. MDisk groups can have different extent sizes. However, different extent sizes can place restrictions on the use of data migration. The choice of extent size affects the total amount of storage that can be managed by a cluster. Table 10 on page 88 shows the maximum amount of storage that can be managed by a cluster for each extent size. Because the SAN Volume Controller allocates a whole number of extents to each virtual disk that is created, using a larger extent size can increase the amount of wasted storage at the end of each virtual disk. Larger

extent sizes also reduce the ability of the SAN Volume Controller to distribute sequential I/O workloads across many managed disks. Therefore, larger extent sizes might reduce the performance benefits of virtualization.

Table 10. Capacities of the cluster given extent size

Extent size	Maximum storage capacity of cluster
16 MB	64 TB
32 MB	128 TB
64 MB	256 TB
128 MB	512 TB
256 MB	1 PB
512 MB	2 PB

Virtual disks

An SAN Volume Controller *VDisk* is a logical disk that the cluster presents to the storage area network (SAN).

Application servers on the SAN access VDisks, not managed disks (MDisks). VDisks are created from a set of extents in an MDisk group. There are three types of VDisks: striped, sequential, and image.

Types

You can create the following types of VDisks:

Striped

The striping is at extent level. One extent is allocated, in turn, from each managed disk that is in the group. For example, a managed disk group that has 10 MDisks takes one extent from each managed disk. The 11th extent is taken from the first managed disk, and so on. This procedure, known as a round-robin, is similar to RAID-0 striping.

You can also supply a list of MDisks to use as the stripe set. This list can contain two or more MDisks from the managed disk group. The round-robin procedure is used across the specified stripe set.

Attention: Care should be taken when specifying a stripe set if your MDisk group contains MDisks of unequal size. By default striped VDisks are striped across all MDisks in the group. If some of the MDisks are smaller than others, the extents on the smaller MDisks will be used up before the larger MDisks run out of extents. Manually specifying the stripe set in this case might result in the VDisk not being created.

Figure 8 on page 89 shows an example of a managed disk group containing three MDisks. This figure also shows a striped virtual disk created from the extents available in the group.

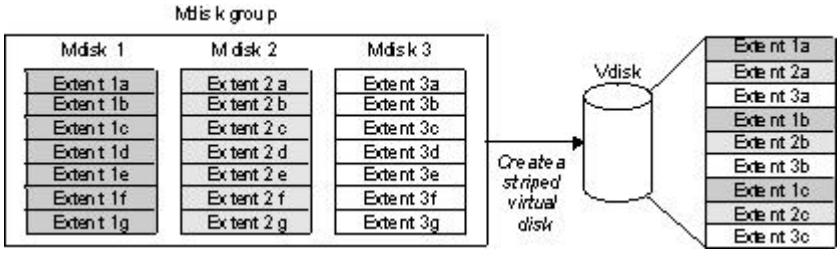

Figure 8. Managed disk groups and VDisks

Sequential
When selected, extents are allocated sequentially on one managed disk to create the virtual disk if enough consecutive free extents are available on the chosen managed disk.

Image Image-mode VDisks are special VDisks that have a direct relationship with one managed disk. If you have a managed disk that contains data that you want to merge into the cluster, you can create an image-mode virtual disk. When you create an image-mode virtual disk, a direct mapping is made between extents that are on the managed disk and extents that are on the virtual disk. The managed disk is not virtualized. In other words, the logical block address (LBA) x on the managed disk is the same as LBA x on the virtual disk.

When you create an image-mode VDisk, you must assign it to a managed disk group. An image-mode VDisk must be at least one extent in size. In other words, the minimum size of an image-mode VDisk is the extent size of the MDisk group to which it is assigned.

The extents are managed in the same way as other VDisks. When the extents have been created, you can move the data onto other MDisks that are in the group without losing access to the data. After you move one or more extents, the virtual disk becomes a real virtualized disk, and the mode of the managed disk changes from image to managed.

Attention: If you add an MDisk to an MDisk group as a managed disk, any data on the MDisk will be lost. Ensure that you create image-mode VDisks from the MDisks that contain data before you start adding any MDisks to groups.

MDisks that contain existing data have an initial mode of unmanaged, and the cluster cannot determine whether it contains partitions or data.

A virtual disk can have one of three states. Table 11 describes the different states of a virtual disk:

Table 11. Virtual disk states

Status	Description
Online	The virtual disk is online and available if both nodes in the I/O group can access the virtual disk. A single node will only be able to access a VDisk if it can access all the MDisks in the MDisk group associated with the VDisk.
Offline	The VDisk is offline and unavailable if both nodes in the I/O group are missing or none of the nodes in the I/O group that are present can access the VDisk.

Table 11. Virtual disk states (continued)

Degraded	The status of the virtual disk is degraded if one node in the I/O group is online and the other node is either missing or cannot access the virtual disk.

If you delete a virtual disk, you destroy access to the data that is on the virtual disk. The extents that were used in the virtual disk are returned to the pool of free extents that is in the managed disk group. The deletion might fail if the virtual disk is still mapped to hosts. The deletion might also fail if the virtual disk is still part of a FlashCopy or a Metro Mirror mapping. If the deletion fails, you can specify the force-delete flag to delete both the virtual disk and the associated mappings to hosts. Forcing the deletion will also delete the copy services relationship and mappings.

HACMP environment

This topic provides information about IBM Tivoli Storage Productivity Center support in an HACMP environment.

IBM Tivoli Storage Productivity Center supports Data agents and Fabric agents installed on High-Availability Cluster Multi-Processing (HACMP) nodes. The Fabric agent can be installed on any number of nodes in the cluster, but the Data agent must be installed on every node in the cluster. Neither of the agents can be clustered, so they cannot be configured to fail over from one node to another. The Data agent collects information from the cluster when the node is probed and 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 and 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 will produce two computer entities in the Data Manager and Disk Manager reports and the topology viewer. Volume groups, NFS shares, and service IP labels that are associated with a clustered resource group can only be reported under the clustered resource group and not the node. Physical volumes, logical volumes, and file systems for clustered volume groups will also be reported under the cluster resource group. Scan and ping jobs can be created for a cluster resource group and these jobs will 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.

The following information lists the support requirements for an HACMP cluster environment.

Agent requirements

Each node in an HACMP cluster must have a Data agent installed. The Data Manager can only monitor and report on nodes that have a Data agent installed. All agents installed in a cluster must use the same port number to communicate with the IBM Tivoli Storage Productivity Center server. Agents in other clusters can be configured with a different port address, but all agents in the cluster must use the port address used by the other agents in that cluster.

Fabric agents are optional in an HACMP cluster. Fabric agents can be installed on any number of nodes in the HACMP cluster.

Probe requirements

Note that probes are not automatically executed in response to cluster events. Probes should be scheduled to run as appropriate for the needs of the environment.

Scan requirements

In order to perform scans on a cluster resource group, the cluster resource group must be configured with an IP address that the IBM Tivoli Storage Productivity Center server can communicate with. This is the IP address that is displayed in reports associated with HACMP clusters. If the server cannot contact this address once it has been configured, the server will try 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.

Scenarios

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

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

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

Monitoring and reporting on I/O utilization

This scenario involves monitoring a DS8000 and a DS4000 to determine what data can be moved to the less inexpensive storage media. Monitoring over time collects the necessary information for analysis. Trending reports, produced with IBM Tivoli Storage Productivity Center Reporter, display utilization information and can be used by management to make decisions about future storage needs.

You have a DS8000 and a DS4000. The DS8000 is nearly full but the DS4000 is only half full. You need to move workload onto the DS4000 as the demand for the DS8000 grows. You will monitor the utilization of volumes on each system to identify potential migration candidates. You will also gather this information over time to indicate trends. This trending information, in report format, will be used for future planning and will be presented to management who will make decisions based on the report.

To perform this scenario, complete the following steps:

Perform performance monitoring

1. **Discover the CIMOMs** This is a two-part process. Run a first discovery to locate the CIMOMs that are visible to the Device server with SLP. After configuring the CIMOM login information, run a second discovery to gather information. Run a CIMOM discovery every four hours.

   Run the first discovery

   a. Expand **Administrative Services** → **Data Sources** → **CIMOM Agents** .
   b. For each CIMOM you use, click on the magnifying glass to the left of the CIMOM name to open the View and Edit window. Enter the name and password information. Check the **Test CIMOM connectivity before editing**.
   c. Click the **Save** icon. This saves your information and performs a connectivity check. You will need to use this process for each SLP-discovered CIMOM you choose to use.

   Run the second discovery

   a. Expand **Administrative Services** → **Discovery**, right-click on **CIMOM**, and click **Run Now**.
   b. In the **CIMOM Discovery job submitted** window, click OK.
   c. When the job has finished, you can view the list of discovered storage subsystems by expanding **Disk Manager** → **Storage Subsystems**.

2. **Create a probe** Create the probe to run during off-hours.
   a. Expand **Tivoli Storage Productivity Center** → **Monitoring** → **Probes**. Right-click **Probes** and select **Create Probe**.
   b. In the **What to Probe** tab, move the DS6000 and the DS4000 from the **Available** list box to the **Current Selections** box.
   c. In the **When to run** tab you will schedule when you want the probe to run. We strongly recommend that you run probes when nothing else is running.
      1) Choose to run the job repeatedly according to a schedule you specify.
      2) Select the current day, month, and year as the date when you want data collection to begin. Select 1:00 AM as the time the job should begin running.
      3) Select the days the job should run. Select all of them. A check mark will appear next to each day you have chosen.
   d. In the **Alerts** tab accept the default (error log) as the means of notification if the job fails.
   e. Click **Save**. The probe will run daily at 1:00 AM.

3. **Create a storage-subsystem performance monitor** Create a storage-subsystem performance monitor to gather performance data and view the results in the log file.
   a. Expand **Disk Manager** → **Monitoring**, right-click **Subsystem Performance Monitor** and select **Create Performance Monitor**.
   b. In the Create Storage Subsystem Performance Monitor window, click the Storage Subsystems tab. Select the DS8000 and DS4000 you want to report on and move them to the Current Selections field. Make sure these storage subsystems have been probed. If they have not been, data collection will fail.
   c. Click the Sampling and Scheduling tab. In the Sampling area, specify the length of the sampling window and how frequently the performance data is to be gathered, using the following values:
      - Interval length = 5 minutes
      - Duration = Indefinite
      - Scheduling = Begin immediately

d. In the Alert tab, **Monitor Failed** is the only triggering condition. In the Triggered Actions area accept the default (error log) means of notification.

e. Click **Save**. When prompted, type a name for the threshold and click **OK**. The name of the storage subsystem performance monitor is displayed under the Subsystem Performance Monitors node. The performance monitor will take 3 hours to run, as set in step c. The status icon will indicate when the job has completed successfully.

f. When the job has completed, a table is displayed listing the storage subsystems. To view the log file for each storage subsystem, select the magnifying glass to the left of the storage subsystems.

4. **Tivoli Storage Productivity Center Reporter for Disk**

   This scenario uses Tivoli Storage Productivity Center Reporter for Disk to prepare a finished report automatically. Tivoli Storage Productivity Center Reporter is a free Java 2 Platform application that extracts storage subsystem information and hourly performance statistics. The statistics are transcribed into a paper-style PDF file containing your storage server utilization, which is saved on the local machine.

   If you do not have Tivoli Storage Productivity Center Reporter for Disk, complete the following steps. If you have installed Tivoli Storage Productivity Center Reporter for Disk, go to step 5.

   a. Go to the Tivoli Storage Productivity Center Reporter for Disk web page, http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS2618.

   b. Select **Download Tivoli Storage Productivity Center Reporter for Disk for Windows** (60MB)

   c. Follow steps to download Tivoli Storage Productivity Center Reporter.

   d. Run the Installer

5. **Create your report with Tivoli Storage Productivity Center Reporter**

   a. To start the program go to **Start** → **Programs** → **Tivoli Storage Productivity Center Reporter for Disk** → **Tivoli Storage Productivity Center Reporter for Disk**. Select **Start**. The Confirmation window will open.

   b. Enter the Hostname (or IP address), Port, UserID, and Password of the system that you want the report on. Click **Continue**. The Select Serial Number window will open. Tivoli Storage Productivity Center Reporter for Disk saves a list of the commands and the status of the commands you are using in the GUI. The commands and status of the commands are displayed in the Status window.

   c. If the probe of the selected subsystems was successful, the DS8000 and DS4000 storage subsystems will be listed in the Select Serial Number window. Select the DS8000 storage subsystem. Click **Select**. You can only create a report for one storage subsystem at a time.

   d. The Select Date Range window opens. Select the **Last 7 days** radio button. Click **Select**,

   e. The Customer Information window opens. This information will be included in the title information of the finished report. The first text field, **Company**, is required. This is the company name that will appear on the title page of the report. All other fields are optional. When finished, click **Continue**.

   f. The IBM/Partner Information window opens. All fields are optional. Click **Continue**.

g. Select a file name and directory in which to save the generated report. The default directory is ″My Documents″ and the default file name is TPCReport.pdf. It is strongly recommended that you save the PDF file under a different name. Click **Save**.

h. The steps Tivoli Storage Productivity Center Reporter takes to create the report are shown in the status window. Once the report has been generated, the **Exit** button becomes available. Click **Exit**.

i. Go to ″My Documents″, find and double-click the PDF report.

The report includes the following information for storage subsystems, ports, arrays, and volumes:

- General information
- Performance summary
- Statistics
- Definitions

6. Using the DS8000 report, look at the Volume Performance Summary for volumes with low levels of I/O and throughput in Megabytes. These might be candidates for moving to the DS4000.

Printing

The report will have been saved as a PDF file using the latest Acrobat Reader on your system. If you only have the Reader of that version of Acrobat, you will be able to view the report, but you will be unable to print it. In this case take the following steps to print the report:

1. Open your complete version of Adobe Acrobat.
2. Select **File** → **Open** and navigate to the report.
3. Select **File** → **Print** to print the report.
4. Close the report. You will not be able to save the report as an earlier version of an Acrobat PDF file.

The following documents were referenced for this scenario:

- *IBM Tivoli Storage Productivity Center User's Guide* GC32-1775
- *Tivoli Storage Productivity Center Reporter for Disk Version 1.01 - Process and Report Overview*

Storage provisioning using Workflows

IBM Tivoli Storage Productivity Center and Tivoli Provisioning Manager Workflow automation can reduce provisioning time from hours to minutes and significantly simplify the task of provisioning new storage capacity. This automation is especially valuable when adding storage provisioning to existing automated server provisioning tasks, such as deploying operating systems, software and network resources. This scenario shows you how to use Tivoli Provisioning Manager and Tivoli Storage Productivity Center Workflows to automate provisioning of SAN attached storage capacity to a server in a Microsoft Windows environment.

Tivoli Storage Productivity Center and Tivoli Provisioning Manager work together to provide automated storage provisioning for the data center.

Tivoli Provisioning Manager provides the means to automate common tasks in the data center that surround each of the four major pillars of provisioning:

- OS provisioning (Windows, Linux, AIX, Solaris, HP-UX, and so forth)

- Software provisioning (DB2, Web Sphere, Oracle, and many more)
- Network provisioning (provisioning VLANs, ACLs, IP addresses, activating and deactivating switch ports, and so forth)
- Storage provisioning (attaching SAN disks for use by the server)

Tivoli Provisioning Manager Workflows automate tasks, especially those that are complex, repetitive, and prone to error by applying best practice rules for zoning, device configuration, and path selection automatically.

Tivoli Storage Productivity Center provides heterogeneous management functions for the SAN storage environment. Tivoli Storage Productivity Center Workflows provide current status information on its SAN storage environment to Tivoli Provisioning Manager. They also allow Tivoli Provisioning Manager to make changes to the SAN through Tivoli Storage Productivity Center.

The following illustration shows the movement of workflow information between Tivoli Provisioning Manager and Tivoli Storage Productivity Center.

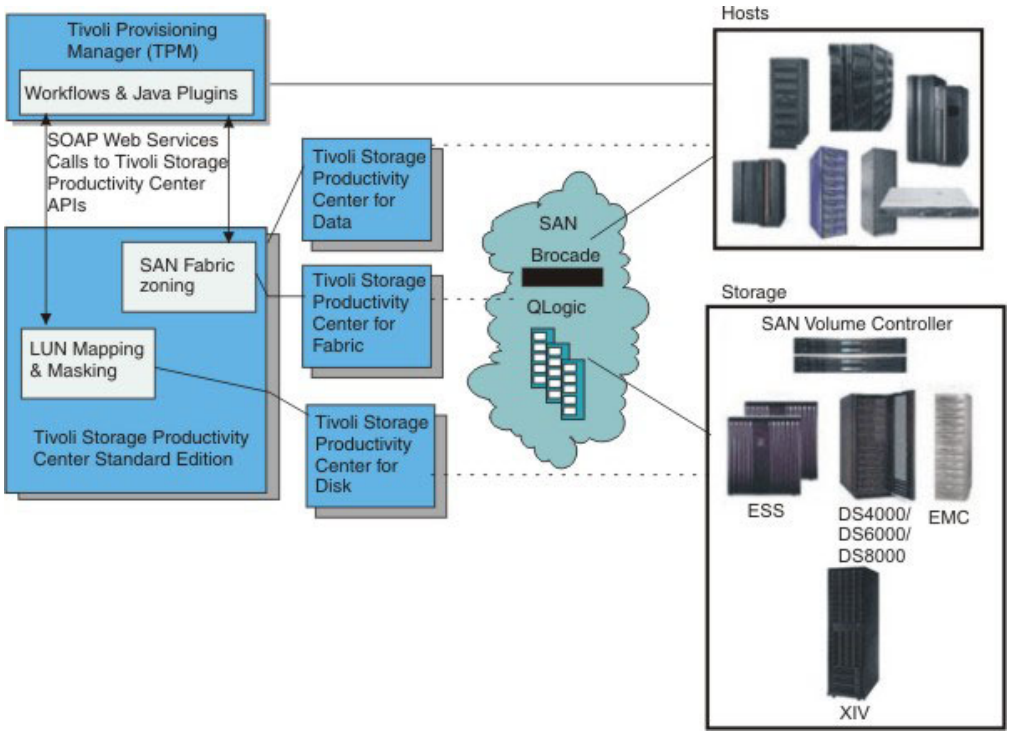

Figure 9. High-level architecture

End-to-end provisioning of SAN storage is difficult due to the complexity of modern environments and inputs that need to be considered. This scenario creates a Tivoli Provisioning Manager storage template used for end-to-end storage provisioning that can automate the process. Once the storage template has been created, it can be re-used to automate end-to-end provisioning of SAN storage with a few mouse clicks.

Without workflows you would need to perform the following individual steps in Tivoli Storage Productivity Center to provision storage:

- Add a volume (storage subsystem)
  - Select storage subsystem
  - Select or create new volume

  - Select host HBA ports (WWNs)
  - Select subsystem controller ports (WWNs)
  - Map volume to controller ports
  - Map volume to host HBA ports
- Set paths (SAN fabric switches)
  - Determine if multiple paths are required
  - Create or update zone(s)
  - Get active zone set
  - Add zone to zone set
  - Activate zone set
- Map the HBA Volumes to the operating system and file system (Host server)
  - Rediscover attached volumes
  - Configure multipathing drivers
  - Map volumes to OS devices
- Create volume and file system (Host server)
  - Create Windows volume
  - Create or extend NTFS file system

With workflows, once the template has been created, these tasks are automated and condensed to the following:

- Select the host that needs additional storage
  - Select the template name from the template list
  - Click OK.

In order to use workflows, the following prerequisites must be met:

- Tivoli Provisioning Manager v5.1 has been installed and configured.
- Tivoli Storage Productivity Center v4.1 has been installed and configured.
- The following two automation packages have been downloaded, installed, and configured from the Open Process Automation Library (OPAL):
  - IBMTPC.tcdriver
  - WindowsNTFS.tcdriver

For more information on the prerequisites needed to use workflows, see the IBM White Paper,*Tivoli Provisioning Manager and IBM Tivoli Storage Productivity Center-Storage Provisioning Configuration and Exercises*.

To create and use a storage template, complete the following steps:

1. To create the Tivoli Provisioning Manager workflow template, complete the following steps:
   a. Select **Inventory** → **Manage Templates** → **Storage Templates** from the Tivoli Provisioning Manager navigation tree.
   b. Select the **Edit** button and click **Add Storage Template**. **Create Storage Template** window opens.
   c. Enter **NTFS_SVC** as the name for the storage template. Click **Next**.
   d. Define the volume container settings. In the **Configure Volume Container Settings** window, enter the following information:
      - Type **NTFS** for the **Volume Container Name**.
      - Select **NTFS** as the **Storage Manager**.
      - Click the **Add Volume Container Settings** button;
      - Click **Next**.
   e. Define the physical volume settings

- Scroll down to the **SAN Physical Volume Settings** and specify a SAN volume.
- Type **SAN** in the **Name** field.
- In the **Volume Container Settings** field, select **NTFS**.
- Select **mdiskgrp0** as the appropriate storage pool from **Storage Pool/Storage Subsystem**.
- Type **5G** as the value for **Consumable Size**.
- Set the **Function Type** to **Generic**.
- For the **RAID Level**, select **UNKNOWN**.
- Click **Add SAN Setting** and click **Next**.

f. Define the data path settings:
- In the **Configure Data Path Settings** window, select **SAN** from the SAN **Physical Volume Setting** list box.
- Click **Add Data Path Settings** and click **Next**.

g. Define the logical volume settings:
- In the **Configure Logical Volume Settings** window, type **TPMVolume** in the **Name** field.
- Select **Simple** as the **Logical Volume Type**.
- Select **NTFS** for the **Volume Container Settings**.
- Specify **5G** for **Consumable Size**.
- Set the **Function Type** to **Generic**.
- Set the **RAID Level** to **UNKNOWN**.
- Click **Add Logical Volume Settings** and click **Next**.

h. Define file system settings:
- In the **Configure File System Settings** window, select **TPMVolume** from the **Logical Volume Settings** list.
- Type **H:** as the **Label**.
- For the **File System Type**, select **NTFS** from the list.
- Specify **5G** for **File System Size**.
- Click **Add File System Settings** and click **Next**.

i. File system mount settings are not used. In the **Configure Filesystem Mount Settings** window, click **Next**.

j. Define disk partition settings:
- In the **Configure Disk Partition Settings** window, set the **Partition Size** to **5G**.
- From the **Physical Volume Settings** list, select **SAN**
- From the **Logical Volume Settings** list, select **TPMVolume**.
- Click **Add Disk Partition Settings** and click **Next**.

k. Verify the summary information. When you are done, click **Finish**.

l. Verify that the template, **NTFS_SVC**, appears in the Tivoli Provisioning Manager GUI at **Inventory** → **Manage Templates** → **Storage Templates**.

This template will allow any system administrator to perform end to end provisioning.

2. To perform provisioning, complete the following steps

a. From the **General** tab of a Windows computer in Tivoli Provisioning Manager, click the **Edit** button and select **Add Storage to Host**.

b. Select **NTFS_SVC** from the **Storage Template** list

c. Click **OK**.

d. Navigate to **Automation → Workflow Status** and track the status of the deployment request in the workflow status log until it completes successfully. It should take approximately 10 to 15 minutes to complete.

This procedure will result in a new 5G volume created on the SAN Volume Controller storage pool **mdiskgrp0** that is assigned to your windows host and mounted and formatted as the drive letter **H** within a few minutes with consistent results and no user errors.

For more information about workflows, see the following publications:

- *IBM Tivoli Storage Productivity Center Workflow User's Guide*
- Redpapers:
  - *Automated Storage Provisioning for Windows using Tivoli Storage Productivity Center and Tivoli Provisioning Manager*
  - *An Introduction to Storage Provisioning with Tivoli Provisioning Manager and Tivoli Storage Productivity Center*
- White Paper: *Tivoli Provisioning Manager and Tivoli Storage Productivity Center - Storage Provisioning Configuration and Exercises*

Creating a performance threshold within an alert

You can monitor the performance of your enterprise by creating alerts with performance thresholds for storage subsystems and SAN switches. By defining alerts that are triggered by performance thresholds, you can get timely information about performance issues in your enterprise.

To properly monitor the performance of your storage subsystem you will need to create the appropriate data collection jobs and alert definitions. The performance monitors collect data from the devices they are monitoring and at each sample interval the data is compared to the performance thresholds you have set. If the thresholds are exceeded, a notification is generated and the actions you defined in an alert are performed.

Note: The term Alert can be used one of three ways. An IBM Tivoli Storage Productivity Center *Alert definition* is defined at **[Data | Disk | Tape] Manager → Alerting → [Type of Alert]**. Within the Alert definition there is an *Alert tab*. Within the Alert tab there is a *Triggering Condition* area which gives the condition that will result in an *Alert message*. Most Triggering Condition areas in the Alert tab have only one Triggering Condition, such as **Probe Failed**. However, in Disk Manager some of the triggering conditions have thresholds, which are condition boundaries you set. For example, **Total I/O Rate Threshold** requires two upper and two lower boundaries or values. Only if these boundaries are matched or exceeded will an Alert message be generated and sent by the method you chose in *Triggered Actions*.

In this scenario, you are defining a performance threshold alert to monitor a disk subsystem (**DS8000-21-7-1300271-IBM**), for the **Total I/O Rate Threshold** condition, and generating, viewing, and printing the results.

To complete this scenario, you need to perform the following tasks:

1. Create a storage subsystem performance monitor
2. Create a storage subsystem alert definition
3. Generate and view a report showing the collected storage subsystem data

Note: You must run probes against storage subsystems prior to successful performance monitoring of those subsystems.

1. **Create a subsystem performance monitor**
   a. Expand **Disk Manager** → **Monitoring**, right-click **Subsystem Performance Monitor** and select **Create Subsystem Performance Monitor**. The **Create Storage Subsystem Monitor** page opens to the **Storage Subsystem** tab.
   b. Select the disk you want to monitor from the **Available subsystems** column and move it to the **Selected subsystems** column. Although we do not do it here, you can select more than one subsystem at a time.
   c. Click the **Sampling and Scheduling** tab. In the **Sampling** area, set the interval for gathering data to **1** hour. Select **Continuing gathering data for** and select a value of **1** month. In the **Scheduling** area, select **Begin immediately**.

      Note: You do not want to sample too frequently as it increases the size of the database repository significantly. Monitors should be set to run for long periods, such as weeks or indefinitely.
   d. Click the **Alert** tab. **Monitor Failed** is the only **Triggering Condition** available. In **Triggered Actions**, select nothing to accept the default. If the job fails, a message will be sent to the log file automatically.
   e. Click the **Save** icon. Name the monitor **SS Monitor**. It will be listed at **Disk Manager** → **Monitoring** → **Subsystem Performance Monitors** → **TPCUser.SS Monitor**. The job will begin running immediately.
2. **Create a storage subsystem performance alert**
   a. Expand **Disk Manager** → **Alerting**, right-click **Storage Subsystem Alerts** and select **Create Storage Subsystem Alert**. The **Create Storage Subsystem Alert** page opens.
   b. Click the **Alert** tab. Select **Total IO Rate Threshold** from the **Condition** drop down list in the **Triggering-Condition** area. Text fields will open for **Critical Stress**, **Warning Stress**, **Warning Idle**, and **Critical Idle**.

      Note: The threshold values given in this scenario are examples and are **not** recommended values. They are for the purpose of illustration only. Your values will undoubtedly be different. Performance behavior varies tremendously between different devices and even different applications for different devices. We recommend monitoring your performance for several weeks and use this historical data to determine reasonable values for each threshold.
      - In the **Critical Stress** field, type **3000** as the upper boundary for critical stress.
      - In the **Warning Stress** field, type **2200** as the upper boundary for warning stress.
      - Leave the **Warning Idle** field, the lower boundary for warning idle, blank.
      - Leave the **Critical Idle** field, the lower boundary for critical idle, blank.
   c. In the **Triggered-Actions** area, select **Email** to send an e-mail to a mail address you specify.
      1) When you select **Email** the **Email Recipients** text box becomes active.
      2) Click **Add** and type an e-mail address in the text box. You can type multiple e-mail addresses.

3) Click **Edit e-mail** to customize the e-mail message. Note that errors sent to the error log are very generic, so you might want to be more descriptive in the text of the e-mail.

d. In the **Storage Subsystems** tab, select **Storage Subsystems** → **DS8000-21-7-1300271-IBM** in the **Available** column and move it to the **Current Selections** column.

e. To save the alert, click the **Save** icon. Name the alert **SS Alert**. It will be listed at **Disk Manager** → **Alerting** → **Storage Subsystem Alerts** → **TPCUser.SS Alert**. The conditions defined in the alert will be checked when the next probe runs against the subsystem.

3. **Generate and view a report**

a. Expand **Disk Manager** → **Reporting** → **Storage Subsystem Performance**, and click **By Storage Subsystem**. The **Storage Subsystem Performance: By Storage Subsystem** window opens.

b. In the **Selection** tab, all performance metrics fields are selected by default. To remove a field from the report, click on the metric name in **Included Columns** and move it to **Available Columns**. In the **Report Filter Specifications** area enter the number of storage subsystems upon which you want a report in the **Return maximum of** text field.

c. Click the **Generate Report** button. A **Storage Subsystem** tab will open containing the report. There will be one line in the report for each storage subsystem.

d. To save the report definition, click the **Save** icon. The **Save As** window opens. Name the report **SS Report** and click **OK**. The report will be saved at **IBM Tivoli Storage Productivity Center** → **My Reports** → **TPCUser's Reports** → **TPCUser.SS Report**.

e. Click the **Print** icon to print the current page on the local machine. You can only print one screen at a time.

Note: You might want to monitor the **Total IO Rate Threshold** condition for the storage subsystem by using constraints instead of thresholds. The Triggering Condition for a constraint would be **Violating Files Consume More Than**. This will alert you if any file goes over the condition value you set. In the **Value** field, type **2200**, which is the value of Warning Stress. In **Value Units**, select **Megabytes** from the pull-down list. See "Archive and delete files" on page 101 for more information on creating constraints.

You can use the above steps, with some minor changes, to create performance monitors and alerts for certain types of switches that have been discovered by Tivoli Storage Productivity Center. Switch performance monitors, for example, are created in the Fabric Manager node:

- To create a switch performance monitor, expand **Fabric Manager** → **Monitoring** → **Switch Performance Monitors**
- To create a switch performance alert, expand **Fabric Manager** → **Alerting** → **Switch Performance Alerts**.

Related tasks

"Creating storage subsystem performance monitor" on page 242
Create a storage-subsystem performance monitor to gather performance data and view the results in the log file.

"Creating a storage subsystem alerts" on page 253
You can use these alerts to monitor storage subsystems for changes in configuration or status. When the condition is detected during a probe, an alert is generated. These alerts can be applied to storage subsystem groups and storage subsystems.

Related reference

"Disk Manager reports" on page 642
Use Disk Manager reports to view the relationships among the components of a storage subsystem and the storage capacity for a computer, filesystem, storage subsystem, volume, Array Sites, Ranks, and disk.

Archive and delete files

This scenario shows you how to reclaim storage space by archiving and deleting files that meet criteria you define.

You are running out of available storage capacity and need to find stale data which can be archived and deleted in order to reclaim space and make it available to more pertinent data. Specifically, you are **TPCUser** and you want to find files on the **C:** file system of **computer1** that are larger than **5MB**, and have not been accessed in **365** days.

Steps include the following tasks:

- Running a probe (if needed)
- Defining the constraint
- Running a scan
- Generating a constraint violation report
- Running an archive/backup on selected files in the report

Before you can use the IBM Tivoli Storage Manager archive or backup functionality you must ensure the prerequisites listed in "Requesting an archive or backup" on page 333 have been met.

To create an archive job, complete the following steps:

1. **Run a Probe**

   You may not need to complete this step if a successful probe has been recently completed. Check the most recent day and time a probe was run successfully at **IBM Tivoli Storage Productivity Center → Monitoring → Probes**. To do this, expand the probe name to see the date the last successful probe was run. In most cases the default probe is sufficient.

   If you do need to run a probe, complete the following steps:

   a. Expand **IBM Tivoli Storage Productivity Center → Monitoring → Probes**, and select **Create Probe**. The **Create Probe** window opens.

   b. Enter a description for the probe.

   c. In the **What to Probe** tab, select the resources you want to probe, in this case **Computer Groups → TPCUser.Default Computer Group → computer1**.

   d. In the **When to Run** tab, accept the defaults: **Run Now** and **Local time in each time zone**.

e. In the **Alerts** tab, accept the default. If the probe fails, a notification will be sent to the error log automatically.

f. Click **Save** to save the probe. Enter **Archive Backup Probe** as the probe name and click **OK**. An entry for the probe appears under **IBM Tivoli Storage Productivity Center → Monitoring → Probes → TPCUser.Archive Backup Probe**.

g. The probe will begin running immediately. The number of times this job has been run and the date and time will be listed under the probe name. To see the status of the job, right-click on the probe name and select **Update Job Status**.

2. **Define Constraints**

   Define the criteria you are looking for with a constraint.

   a. Expand **Data Manager → Policy Management → Constraints**, right-click and select **Create Constraint**.

   b. In the **Description** field, enter a description of this constraint.

   c. In the **Filesystems** tab **Available** column, select **Filesystems → computer1 → C:** and move it to **Current Selections**.

   d. In the **File Types** tab, check **Forbid files matching these patterns**. All file types should be in the **Choose a file type** column. This will apply the constraint to all file types.

   e. In the **Users** tab, Check the **Forbid files owned by selected users**. All users should be listed in the **Available Users** column. This will apply the constraint to all users.

   f. In the **Options** tab, leave the default, **200**, in the text box in the **Violating File Limits** section. In the **Alerts on Files** section, select **number of days since last access greater than** and type **365**. Select **bigger than** and enter **5 MB**. Below the **Alerts on Files** section, **File Filter Text:** will list the values you picked in the **File Types**, **Users**, and **Options** tabs.

   g. In the **Alert** tab under **Triggering Conditions**, the only triggering condition is **Violating files consume more than**. Type **5** in the **Value** field. Select **Megabytes** from the **Value Units:** pull-down. Accept the default in **Triggered Actions**, which will send a message to the error log if the constraints you have defined are violated.

   h. Click the **Save** icon. Name the constraint **Archive File Constraint**. The constraint will be saved as **Data Manager → Policy Management → Constraints → TPCUser.Archive File Constraint**.

3. **Run a Scan**

   Scan with the constraint violation you have created.

   a. Expand **Data Manager → Monitoring**, right-click **Scans** and select **Create Scan**. The **Create Scan** window opens.

   b. In the **Filesystems** tab, select **Filesystems → computer1 → C:/** from the **Available** column and move it to the **Current Selections** column.

   c. Do not select anything in the **Directory Groups** tab.

   d. In the **Profiles** tab, select all the profiles listed in the Available Profiles column and move them to both **Profiles to Apply to Filesystems** and **Profiles to Apply to Directories**.

   e. In the **When to Run** tab, select **Run Now** and **Local Time in each time zone**.

   f. In the **Alert** tab, **Scan Failed** is the only **Triggering Condition**. In the **Triggered Actions** area, select nothing to accept the default. If the job fails, a message will be sent to the log file automatically.

g. Click the **Save** icon. Save the scan as **Archive Scan**. The scan is saved under **Data Manager** → **Monitoring** → **Scans** → **TPCUser.Archive Scan**. The scan starts to run immediately. Scans can take some time to complete depending on the number of storage assets scanned and the profiles used.

4. **Generate a Constraint Violation Report**

   Generate a report listing files that met your constraint violation criteria.

   a. Expand **Data Manager** → **Reporting** → **Usage Violations** → **Constraint Violations** and click **By Constraint**. The **Constraint Violations: By Constraint** window opens to the **Selection** tab.

   b. In the **Report Filter Specifications** area, type **100** in the **Return maximum of** text field.

   c. At the top right of the **Report Filter Specifications** area, click **Selection**. The **Select Resources** pop-up opens listing the available constraints.

   d. Select **Archive File Constraint**. (You may have to de-select the other constraints.) Click **OK**.

   e. Select the columns you want to appear in the report. By default, all columns will be listed in **Included Columns**. If you want to remove columns from the report, select them and move them to **Available Columns**.

   f. Click **Generate Report**. The Constraint Violation Report is displayed on the monitor.

   g. The page will stay the same with the addition of a **Constraints** tab which shows the **Constraint Violations: By Constraint** report. Click the magnifying glass to drill down to the **Constraint Violations: By Filesystem** report. Click the magnifying glass again and the **Largest Files: By Filesystem** report opens. This report, the **Constraint Violation Report**, will be used to create an Archive/Backup job.

5. **Run Archive/Backup on selected files in the report**

   Archive and backup files are automatically sent to the Tivoli Storage Manager server and are stored on that machine. Files are put into storage groups that are controlled through a Tivoli Storage Manager policy that the Tivoli Storage Manager server administrator usually creates.

   a. Select the files in the **Constraint Violation Report** that you want to archive or backup.

   b. Right-click on the selected files and select **Create a new archive\backup job**. The **Create New Archive/Backup Job** window opens.

   c. In the **Options** tab, choose **Archive**. In this scenario the purpose is to archive and delete files that violate the constraints, so below the **Archive** selection, select **Delete After Successful Archival**. If, for some reason, you did **not** want to delete the files after they have been archived, do not select **Delete After Successful Archival**. If the archive is not successful, the files will not be deleted.

   d. The files you have selected for archive/backup are shown in the **Files** tab.

   e. In the **When to Run** tab select **Run Now**. In the **How to handle time zones** section, select **Local time in each time zone**.

   f. In the **Alert** tab, in the **Triggering-Condition** area, the only triggering condition is **File Archive/Backup Failed**. Do not select anything in the **Triggered Actions** area to accept the default (a message sent to the error log).

   g. Click the **Save** icon. Enter the name **Archive Log** and click **OK**. The job log is saved at **Data Manager** → **Monitoring** → **Policy Management** → **Archive/Backup** → **TPCUser.Archive Log**.

h. A notice telling you that the job has been submitted appears. Click **OK**.
i. Expand the job name to see the status of this particular job. You can judge the status by the color coding. When the job has finished, click on the date and time to view the job log file.

Setting an archive/backup from the constraint

You can call archive backup when you set the constraint. However, if you set the delete option, it will automatically archive and delete **all** files that match the constraint criteria, which you probably do not want to do. At the very least you should not set archive/backup in the constraint the first time you run it.

To call for an archive/backup within a constraint, select **Archive/Backup** in the **Triggered Actions** area of the **Alert** tab in the **Create Constraint** window. Click **Define**. An **Archive/Backup Options** pop-up opens. Select **Delete After Successful Archival**. Go to the next tab and continue defining your constraint.

Related concepts

"Requesting an archive or backup" on page 333
Request anIBM Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint. Tivoli Storage Manager protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage.

Related tasks

"Creating probes" on page 203
Create and schedule a probe to job to gather information about your storage resources. You can define any number of probes to gather information about different storage resources at different times.

"Creating a file constraint" on page 327
Learn how to create a constraint that defines the acceptable and unacceptable file types, file sizes, and file owners for a computer or a set of computers in your environment.

"Defining an archive or backup in a constraint" on page 334
Learn how to request an IBM Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint. Request a Tivoli Storage Manager archive and backup on the alert page of a constraint definition.

"Creating scans" on page 211
Create and schedule a scan to gather usage information about your storage resources. You can specify which file systems, computers, clusters, directories, databases, and so on. that you want to be scanned. You can define any number of scans to gather information about different storage resources at different times.

Create a profile with a filter

This scenario scans storage on a selected drive looking for the 100 largest media files, and summarizes space used by filesystem or directory and by owner.

There are a number of media files on computer1 that you suspect are taking up space that could be used for other purposes. You would like to search C:\ drive file systems and directories for the 100 largest files and create a report you can use to decide what action to take (contact the owner, archive and delete the files, and so forth). You will also publish this report to the web, where other people with permission levels greater to or equal to yours (TPCUser) can access it.

With this scenario, you perform the following tasks:

1. Create a profile to look for the largest media files
2. Create a filter within the profile
3. Create a new scan job using the profile
4. Create and run a batch report using the information gathered from the scan job
5. Save this report as an HTML file

For this scenario, complete the following steps:

1. **Create a new profile**
   a. Expand **Data Manager** → **Monitoring**, right-click **Profiles** and select **Create Profile**.
   b. The **Create Profiles** window opens. Enter a description for the profile.
   c. In the **Statistics** tab, do the following:
      1) Under **Summarize space usage by**, select **filesystem or directory** and **owner**.
      2) Under **Accumulate History**, select **Per Scan** and type **10** as the number of days to keep.
      3) Under **Gather statistics by length of time since**, select **creation**.
      4) Under **Gather the information on the**, select **largest files** and type **100** in the text field.
   d. In the **File Filters** tab, create the filter by doing the following:
      1) Highlight **(All Files Selected)** and click **New Condition**. The **Create Condition** window opens.
      2) Select **NAME** and **Matches any of** from the pull-down lists.
      3) Select **In any Directory**.
      4) In the **match** field, enter the types of files you are looking for in either a comma-separated list or by pressing return after each entry. As you press return, the file type will be listed in the box underneath the text field. For this scenario the file types are ***.avi**, ***.au**, ***.mpeg3**, and ***.gif**. Click **OK**. The **Create Profile** window re-opens, reflecting your choices.
   e. Click the **Save** icon. Name the profile **Media Files Profile** and click **OK**. The profile will be saved as **Data Manager** → **Monitoring** → **Profiles** → **TPCUser.Media Files Profile**.
2. **Create the scan**
   a. Expand **Data Manager** → **Monitoring**, right-click **Scans** and select **Create Scan**.
   b. The **Create Scans** window opens. Type a description of the scan.
   c. In the **Filesystems** tab, select **C:/** as the file system to scan.
   d. Do not select anything in the **Directory Groups** tab.
   e. In the **Profiles** tab, select **TPCUser.Media Files Profile** from Available Profiles and move it to **Profiles to apply to Filesystems**.
   f. In the **When to run** tab, choose **Run Now** and select **Local time in each time zone**.
   g. In the **Alerts** tab, the only **Triggering Condition** is **Scan Fails**. In the **Triggered Actions** area, select nothing to accept the default. If the job fails, a notification will be sent to the error log automatically.
   h. Click the **Save** icon. Name the scan **Media Files Scan** and click **OK**. The scan will be saved as **Data Manager** → **Monitoring** → **Scans** → **TPCUser.Media Files Scan** and the job will begin running immediately. This scan could take approximately 20 minutes to complete.

3. **Create the batch job**

   Once you have run a scan with the new profile, the data is in the database but you have to go get it with a custom-designed report or batch job.

   a. Expand **IBM Tivoli Storage Productivity Center → My Reports**, right-click **Batch Reports** and select **Create Batch Report**.

   b. The **Create Batch Report** window opens. Enter a description for the report.

   c. In the **Available** column in the **Report** tab, select **Usage → Files → Largest Files → by Computer** and move that to **Current Selections**.

   d. In the **Selection** tab, select **TPCUser.Media Files Profile** from the available profiles. By default, all columns are included in **Included Columns**. If you do not want a column in your report, select it and move it to **Available Columns**.

   e. In the **Options** tab, make the following selections:

      - Under **Agent Computer Specification**, select **computer1**. This is the computer where the report will be run and the data stored. In most cases this is the Agent running on the Tivoli Storage Productivity Center server.
      - Under **Report Type Specification**, select **HTML File**.
      - Select nothing under **Script**.
      - Under **Output File Specification**, enter the mask that will determine the output file name. It is recommended that you name the report something other than the default so that it can be easily identified by you and others. **{Report name}.{Report run date}** is recommended.

   f. In the **When to run** tab, choose **Run Now** and select **Local time in each time zone**.

   g. In the **Alerts** tab, the only **Triggering Condition** is **Report Failed**. In the **Triggered Actions** area, select nothing to accept the default. If the job fails, a notification will be sent to the error log automatically.

   h. Click the **Save** icon. Name the batch report **Largest Media Files** and click **OK**. The batch report criteria will be saved under **IBM Tivoli Storage Productivity Center → Reports → Batch Reports → TPCUser.Largest Media Files**. The job will begin running immediately. The job number, date, and time will be listed under **TPCUser.Largest Media Files**.

4. These reports can be saved to any server running IBM Tivoli Storage Productivity Center for Data. By default, the server from which you run the report will have an active browser. The HTML file will be saved to `C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\log\computer1\reports\Largest Media Files.<date>.HTML`, and can be opened in the browser by anyone with your permission level.

Related tasks

"Creating a profile" on page 220
Create a profile for use within a scan.

"Creating scans" on page 211
Create and schedule a scan to gather usage information about your storage resources. You can specify which file systems, computers, clusters, directories, databases, and so on. that you want to be scanned. You can define any number of scans to gather information about different storage resources at different times.

"Creating batch reports" on page 700
Create a batch report.

Determine application performance problems

This scenario provides a method for finding out why an application has stopped performing as expected.

A previously normal application has stopped performing well. You must find out what is impacting the application so that you can take steps to correct the situation. You suspect a storage subsystem problem.

With this scenario you perform the following tasks:

- Probe the Data agents
- Probe the Storage Subsystems and start the Performance Monitor running
- Create a **Storage Subsystem: By Filesystem/Logical Volume** report
- Create a **Storage Subsystem Performance: By Volume** report
- Create a **Storage Subsystem Performance: By Array** report

Tip: Ordinarily, you run probes and performance monitors on both storage subsystems and switches, because the problem might be in either place. However, in this scenario we are only gathering information on storage subsystems.

Prerequisites:

- The Data Agents must be installed on the managed servers before you can complete the tasks in this scenario. These are usually installed as part of the IBM Tivoli Storage Productivity Center installation process.
- Data Agents and Arrays must have already been probed.
- CIMOMs for storage subsystems must be added and discovered.

Substitutable values that are used in this scenario:

- Device path = M:\ (from application owner or server administrator)
- Computer = ODCCL163.odcclass.ibm.com (from Filesystem/Logical Volume report)
- Volume name = testvol-100 (id.1000)(from Filesystem/Logical Volume report)
- Storage subsystem name = DS6000–1750–13AB44A-IBM (from Filesystem/Logical Volume report)
- Array = 1750.13AB24A-1 and A-2 (from the By Volume report, Drill Up to By Array report)

To track down application performance problems, collect information by running probes, performance monitors, and reports. Then use that information to determine actual performance and use values.

To perform this scenario, complete the following steps:

1. When the problem is reported, get computer information, device path information, and the application name from the application owner or the server administrator.
2. Probe the Data agents.
   a. Expand **IBM Tivoli Storage Productivity Center → Monitoring**. Right-click **Probes** and select **Create Probes**.
   b. In the description field, type `Probe for Data agents`.
   c. In the **What to Probe** tab, expand 'Computers' in the **Available** box, select All Computers, and move that to **Current Selections**. This will include all Data agents in the probe,
   d. Click the **When to Run** tab and choose **Run Repeatedly**, and set a regular time the probe will run. The following settings will set up data collection for 2:00 AM every morning.
      1) Set the day in **Beginning at** a day ahead of the current day. If the day is February 21st, 2008, enter February 22, 2008, at 2:00 AM.
      2) Select **Repeat Every** and set **23 Hour(s)** to specify how often the probe should run. This allows time for the probe to shut down and start up again.

      In the **How to handle time zones** area, select **Use the timezone that the server runs in**, which uses the time zone where the Data server resides.
   e. On the **Alert** tab, **Probe Failed** is the only triggering condition that triggers an alert. In the **Triggered Actions** area, select **Windows Event Log** and set the **Event Type** to **Warning**.
   f. Click the **Save** icon. Type `DataAgentsProbe_1` as the probe name and click the **Save** button. The probe is listed in the navigation tree as **IBM Tivoli Storage Productivity Center → Monitoring → Probes → Administrator.DataAgentsProbe_1**.
   g. To begin collecting data immediately, expand **IBM Tivoli Storage Productivity Center → Monitoring → Probes**, right-click **Administrator.DataAgentsProbe_1** and select **Run Now**. The status line for the probe is listed under **Administrator.DataAgentsProbe_1** by date and time.

      A Data Agent probe typically takes under a minute.
3. **Probe the Storage Subsystems**. Creating the Storage Subsystem probe is very similar to creating the probe for the Data Agents.
   a. Expand **IBM Tivoli Storage Productivity Center → Monitoring**, right-click **Probes** and select **Create Probe**. In the **Description** field, type `Probe for Storage Subsystems`.
   b. In the **What to Probe** tab Available area, expand **Storage Subsystems**. Select **All Storage Subsystems** and move it to Current Selections.
   c. In the **When to Run** tab, select **Run Repeatedly** beginning a day later than the current day at 3:00 AM.
   d. Click the **Save** icon. Type `StorageSubsystemsProbe_1` as the probe name and click the **Save** button. The listing for the probe appears in the navigation tree as **IBM Tivoli Storage Productivity Center → Monitoring → Probes → Adminitrator.StorageSubsystemProbe_1**.
   e. To collect initial data, right click **Administrator.StorageSubsystemProbe_1** and select **Run Now**. Depending on the number and size of the storage subsystems to be probed, this might take up to 15 minutes.

4. **Begin running the storage subsystem Performance Manager**. If you have a current Performance Monitor running and you have added another storage subsystem, you must start a new Performance Monitor.
   a. Expand **Disk Manager** → **Monitoring**, right-click **Subsystem Performance Monitor** and select **Create Performance Monitor**. In the **Description** field, type `Subsystem performance monitor`.
   b. Click the **Storage Subsystems** tab. All storage subsystems that have had probes created and completed are listed under **Available subsystems**. Select all the available subsystems and move them to **Selected subsystems**.
   c. Click the **Sampling and Scheduling** tab. In the Sampling area specify how frequently data is to be gathered and the length of the sampling window.
      1) For the Interval Length, type `5` in the minutes field.
      2) For Duration, select **Continue indefinitely**.
      3) Select **Begin Immediately** in the **Scheduling** area.
   d. Click the **Alert** tab.
      1) In the **Condition** area, **Monitor Failed** is the only triggering condition.
      2) In the **Triggered Actions** area, select **Windows Event Log** with a **Warning** event type.
   e. Click the **Save** icon and enter `SSPerformanceMonitor_1` as the performance monitor name. Click **OK**. The name of the storage subsystem performance monitor is displayed at **Disk Manager** → **Monitoring** → **Subsystem Performance Monitor** → **Administrator.SSPerformanceMonitor_1**.

      Tip: Setting performance alerts is not necessary for this scenario, but at this point you can review the historical performance data for your environment. Then you can define performance alerts with values that you consider out of the ordinary for your environment. You can set performance alerts for both storage subsystems and switches. After a performance alert is defined, future deviations will generate an alert and create an entry in the Performance Manager Constraint Violation report.
5. Gather information for at least two cycles (in this case, gather the information for at least ten minutes). When information has been gathered, run the following reports to collect the data you will use to check performance.
   - **Disk Manager** → **Reporting** → **Storage Subsystems** → **Computer Views** → **By Filesystem/Logical Volume**
   - **Disk Manager** → **Reporting** → **Storage Subsystem Performance** → **By Volume**
   - **Disk Manager** → **Reporting** → **Storage Subsystem Performance** → **By Array**

   The data from the reports is saved by date and time, but you might want to print a copies of the reports for easier comparison.
   a. The **Storage Subsystem: By Filesystem/Logical Volume** report identifies the back-end storage volumes associated with the file system. Variable values that are gathered from this report include those for Computer, Storage Subsystem, and Volume Name.
      1) Run the report by expanding **Disk Manager** → **Reporting** → **Storage Subsystems** → **Computer Views** and click **By Filesystem/Logical Volume**.
      2) Click the **Selection** tab, and in the **Report Filter Specifications** area, select **Volumes** to **Relate Filesystems/Logical Volumes to** in the list.
      3) Click the **Selection** button, clear the listed computers, and select your Device Path (M:\). Click **OK** to return to the **Selection** tab. When you run the report, the information is limited to your device path.

4) By default all columns are listed as **Included Columns**. If there are columns you do not want in the report, select them and move them to **Available Columns**. Do **not** to move the ″volume name″ columns to **Available Columns**.

5) Click **Generate Report**.

b. The **Storage Subsystem Performance: By Volume** report shows you the actual performance metrics that are associated with the volumes the problem file system resides on.

1) Expand **Disk Manager** → **Reporting** → **Storage subsystem Performance** and click **By Volume**. By default all columns are listed as **Included Columns**. Do not move any of them to **Available Columns**.

2) Click the **Selection** button. Click **Deselect All**, then find your subsystem name (DS6000–1750–13AB24A-IBM) and component (the Volume Name - testvol-1000 (ID.1000)), and select that line. The information in the report will be limited to that subsystem and volume.

3) Click **OK**. Click **Generate Report**.

c. The **Storage Subsystem Performance: By Array** report shows the other volumes on the same array as the problem volume. It indicates whether the problem is specific to that one volume or if all of the volumes on the same array have a similar problem.

Note: Arrays are only supported on Tivoli Storage Enterprise Storage Server, DS8000, and DS6000 subsystems.

1) Expand **Disk Manager** → **Reporting** → **Storage subsystem Performance** and click **By Array**.

2) In the **Report Filter Specifications** area, select **Display latest performance data (by sample)**.

3) Click the Selection button. Click **Deselect All** then find your arrays (Component) (1750.13AB45A-1 and A-2) and select those lines. The information in the report will be limited to those arrays.

4) By default all columns are listed as **Included Columns**. Do not move any of them to **Available Columns**.

5) Click **Generate Report**.

6. Parts of the information that you use to determine the problem is in all of the reports. Look for the following things:
   - Look at the I/O rate and the response time together for a complete picture of the performance. Ideally, you want a high I/O rate and low response time.
   - If the I/O rate is high but the response time is slowing over time, look at the cache hits to see how well cache is performing and, therefore, how often the I/O has to go to disk.
   - If the I/O rate is low, look at the data rate. When there is a large number of I/O requests with high transfer sizes, the data rates become more important than I/O rates. In this case, look at the data rates and transfer sizes together.
   - Look at all the volumes in the same array to determine the I/O rate and response times. To determine this, drill up from a selected volume in the 'By Volume' report to view the arrays that the volume is a part of. Then drill down from the array to get a filtered 'By Volume' report of just the volumes on that array. Now compare the performance metrics against all the volumes in the same array. If all of them show the same performance problem, the performance bottleneck is the array. If only one volume has a performance problem, the performance bottleneck is somewhere else.

7. In this scenario, based on the I/O rates and the response times for the array over a period of time, you determine that the array is being saturated with I/O and I/O should be off-loaded to other storage to improve performance.

   Tip: You can also use Tivoli Storage Productivity Center to determine and project when more storage space is needed by monitoring I/O rates and response times over a period of time.

Related information:

- *IBM Tivoli Storage Productivity Center User's Guide* or IBM Tivoli Storage Productivity Center Information Center
  - Adding CIMOMs, see "Adding a CIM agent manually" on page 124.
  - Discovering CIMOMs, see "CIMOM discovery" on page 190.
  - Setting performance alerts, see "Creating alerts" on page 247.
- *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* or IBM Tivoli Storage Productivity Center Information Center
  - Installing agents, see Installing the agents.
  - Planning performance metrics, see Planning for performance metrics.

Monitor space on a remote server

Monitor a remote storage subsystem and take action when available space on a storage subsystem is less than 50GB.

A IBM Tivoli Storage Productivity Center server at a remote location (doremon) has storage on a DS4800 which is close to full. Monitoring has been set so that when the available storage space is less than 50GB an alert is generated to the server on which the Universal Agent is running. A system administrator can then use the Tivoli Enterprise Portal console to Launch in Context to the Tivoli Storage Productivity Center Wasted Space Report. Based on this report, orphan files and files that have not been accessed in a year on the DS4800 can be identified.

This scenario is designed for the Windows operating system. The following IBM products are used:

- **IBM Tivoli Monitoring** optimizes IT infrastructure performance and availability through a customizable workplace portal.
- **Tivoli Enterprise Portal** is the user interface for IBM Tivoli Monitoring. It is a portal-like application server that enables retrieval, manipulation, and analysis of data from the database server.
- A **Universal Agent** solution that reports Tivoli Storage Productivity Center asset information to IBM Tivoli Monitoring. This data is available for display in the Tivoli Enterprise Portal for reporting, charting and establishing situations in Tivoli Monitoring. A Universal Agent monitoring solution is provided by Tivoli Storage Productivity Center.
- The **Launch in Context** feature that allows you to pass parameters to the Tivoli Storage Productivity Center GUI. Based on the parameters you pass to the GUI, you can automatically log into the specified Tivoli Storage Productivity Center server, navigate to a specific panel, and perform tasks on the specified server. Launch in Context is a feature of the Tivoli Storage Productivity Center user interface.
- **Tivoli Storage Productivity Center** provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance and assets in your enterprise environment.

The following applications must be running:

- Tivoli Storage Productivity Center
- IBM Tivoli Monitoring, which includes Tivoli Enterprise Portal.
- The Tivoli Storage Productivity Center Universal Agent monitoring solution must be installed and configured per the readme. The readme is located in the TPCUA.zip file, which is on install disk1 in the tool folder.

Substitutable values used in this scenario:

- Tivoli Storage Productivity Center server = doremon
- doremon password = J24875
- IP address for doremon = 9.52.172.127
- User = Administrator
- <ITMHOME> = C:\ibm\ITM
- Situation name = DS4800_full
- Sampling interval = 15 minutes
- Wasted Space report name = WastedSpaceReport_Doremon

To complete this scenario, follow these steps:

1. Set up the IBM Tivoli Monitoring Universal Agent to monitor doremon.

   This procedure is fully documented in the readme. Note that in real-life situations the Universal Agent would monitor much more than one server.

   a. Start a command line session and change directories to `<ITMHOME>`.

   b. Add doremon as a server to be monitored using the `-config add` command.

      Enter the following command:

      ```
      <ITMHOME>\tmaitm6\scripts\tpcua TPCUA.bat -config add -tpcDisplay
      "TPC Server Doremon" -tpcDBIP 9.52.172.127 -tpcDBPort 50000
      -tpcDBSchema TPC -tpcDBName TPCDB -tpcDBUser Administrator -tpcDBPass
      J24875
      ```

   c. Schedule information collection using the Windows Scheduler.

      1) Enter the following command:

         ```
         <ITMHOME>\tmaitm6\scripts\tpcua\TPCUA.bat -collect
         ```

      2) Go to **Start** → **All Programs** → **Accessories** → **System Tools** and click **Scheduled Tasks**. Double-click **Add Scheduled Task**. This starts the Scheduled Task Wizard. Click **Next**.

      3) Select **Browse**. Go to `<ITMHOME>\tmaitm6\scripts\tpcua` and select `TPCUA.bat`. Click **Open**.

      4) In the text field, type **TPCUA**. Select the **Daily** radio button. Click **Next**.

      5) In the **Start time** field, accept the default (current time). Select the **Every Day** radio button. In the **Start date** field, accept the default (current date). Click **Next**.

      6) Enter **Administrator** as the user name. Enter and confirm the password. Click **Next**.

      7) Verify your selections. Click **Finish**. You will set the "repeat" time in the next steps.

      8) Go to **Start** → **All Programs** → **Accessories** → **System Tools** → **Scheduled Tasks**, right-click **Add Scheduled Task** and select **Properties**.

      9) In the "Schedule" tab, click **Advanced**. Select the **Repeat task** check box and specify 10 minutes as the time to repeat the collection. Click **OK**.

      10) Click **OK**.

d. Import the metafile. On the command line, go to `<ITMHOME>\tmaitm6` and enter the following command:

```
kumpcon import TPC_Network.mdl
```

Once the metadata file has been imported, the Universal Agent collects Tivoli Storage Productivity Center data and provides it to IBM Tivoli Monitoring. It may be necessary to restart the Universal Agent after importing the .mdl file.

2. From IBM Tivoli Monitoring, set a situation for doremon that will alert you when the available storage space on the DS4800 falls below 50GB.
   a. Using the Tivoli Enterprise Portal, open the Navigator to the Physical view. Expand **Enterprise** → **Windows Systems** → **DOREMON** → **Universal Agent** → **doremon:TPC_NETWORK00**. and right-click **Select Situations**.
   b. Select **STORAGE**. Right-click **STORAGE** and select **Situations**. The ″Situations for - Storage″ window will open.
   c. Right-click **STORAGE** in the navigation tree and select **Create New**. The ″Create Situation″ window opens for you to enter the situation name and description.
   d. Enter `DS4800_full` for the name and `″DS4800 storage subsystem nearly full on doremon″` as the description. Click **OK**.
   e. The ″Select Condition″ window opens.
      1) Select **Attribute Comparison** as the **Condition Type**
      2) Choose **STORAGE** in the Attribute Group area
      3) Choose **Available Capacity** in the Attribute Item area
      4) Click **OK**.
   f. The ″Situations for - STORAGE″ window reopens to the **Formula** tab. The description for DS4800_full may or may not be listed in the **Description** text box. If it is not, re-type it.
      1) In the Formula area, **Available Capacity** is listed with three cells below it. Click in the first cell. A formula area opens. Select a value of less than (<) 50.
      2) Leave **15 min** as the default **Sampling interval**.
      3) To play a sound to alert you, select **Enable** in the Sound area.
   g. Click the **Action** tab.
      1) For **Action Selected**, select **System Command** and type `″net send DS4800 managed by Doremon is nearly full″` in the **System Command** text box.
      2) For **If the condition is true for more than one monitored item**, select **Only take action on the first item**.
      3) For **Where should the Action be executed (performed)**, select **Execute the Action at the Managed System (Agent)**. This will send the command window with message and the sound to the system on which the Universal Agent monitoring Doremon is running.
      4) For **If the condition stays true over multiple intervals**: select **Do not take action twice in a row (wait until situation goes false then true again)**.
   h. Click the **Distribution** tab. Verify that **doremon:TPC_NETWORK00** is the correct assigned agent. If there is more than one assigned agent, move any incorrect agents to the Available Managed Systems area.
   i. Click **OK**. The ″Situations for - STORAGE″ navigation tree will list **STORAGE** → **DS4800_full**. If you should be returned to the ″Welcome SYSADMIN″ window, expand **Enterprise** → **Windows Systems** →

DOREMON → Universal Agent → doremon:TPC_NETWORK00 and right-click **Situations**. The ″Situations for - STORAGE″ window will open and **DS4800_full** will be listed.

3. IBM Tivoli Monitoring checks the data provided by the Universal Agent every 15 minutes. When storage on the DS4800 falls below 50GB, you will be notified by a sound and the message you entered in the Action tab in a command window on the system on which the Universal Agent monitoring Doremon is running.
4. Run a Wasted Space report. When you are alerted that storage space is less than 50GB on the DS4800, use Tivoli Enterprise Portal to create a definition for Launch in Context that will launch a Wasted Space report.
   a. In the Tivoli Enterprise Portal window, expand **Enterprise → Windows Systems → DOREMON → Windows OS**, right-click **Disk** and select **Launch**.
   b. If you **have not** previously created a Launch in Context definition for the Wasted Space report task, complete the following steps:
      1) In the Tivoli Enterprise Portal window, expand **Enterprise → Windows Systems → doremon → Windows OS** Right-click **Disk** and select **Launch**.
      2) The ″Create or Edit Launch Definitions″ window opens. Click **Create New**.
      3) In the **Name** text field, type `WastedSpaceReport_Doremon`,
      4) For the **Target** text box, select **Browse** and go to C:\Program Files\IBM\TPC\gui\tpc.bat.
      5) In the **Arguments** text field, enter the parameters and values you will use to run the Wasted Space report:

         ```
         -user admin -passwd J24875 -server doremon:5000 -function
         TPC.reports.data.wasted_space
         ```

      6) For the **Start in** location, click **Browse** and select C:\Program Files\IBM\TPC\gui.
      7) Click **Apply** to save the report definition as **WastedSpaceReport_Doremon** in the **Existing Launch Definitions** panel.
      8) Click **Launch** to run the report.
   c. If you **have** previously created a definition for **WastedSpaceReport_Doremon**, select the definition from the **Existing Launch Definitions** panel and click **Launch**.
5. The Wasted Space report is displayed within the Tivoli Storage Productivity Center GUI, showing the totals for orphaned and stale files that can probably be removed to free up space on the DS4800.

The following documents were referenced for this scenario:

- A Universal Agent Solution readme
- IBM Tivoli Monitoring User's Guide or IBM Tivoli Storage Productivity Center Information Center:
  - Creating a situation
- IBM Tivoli Storage Productivity Center User's Guide or IBM Tivoli Storage Productivity Center Information Center:
  - Launch in Context
  - Wasted Space report

Identifying performance hot spots

This scenario uses the Storage Optimizer to analyze storage subsystems to determine if there are performance hot spots. This scenario also illustrates how to use the Storage Optimizer to get recommendations for improving storage subsystem performance.

You want to analyze your two DS8000 storage subsystems to determine if there are performance hot spots. You want to identify the hot spots and get recommendations on how to improve storage subsystem performance.

Prerequisites for using Storage Optimizer

- IBM Tivoli Storage Productivity Center Standard Edition license
- One or more of these supported IBM subsystems: DS8000, DS6000, DS4000, Tivoli Storage Enterprise Storage Server, SAN Volume Controller.
- At least one week's worth of performance monitoring data gathered in advance of using Storage Optimizer. Providing a longer time interval for data collection will increase the accuracy of the analysis and recommendations.

Monitor the performance of your DS8000 subsystems

At least one week in advance, perform performance monitoring for your two DS8000 storage subsystems. The general steps are as follows:

1. Run a CIMOM discovery for storage subsystem CIMOMs. See "Discover storage subsystem, tape library, and fabric information" on page 197 for more information.
2. Create a probe. See "Creating probes" on page 203 for more information.
3. Create a storage subsystem performance monitor and let it run for at least one week. See "Creating storage subsystem performance monitor" on page 242 for more information.

To perform this scenario, complete the following steps:

Part 1: Create an analysis report that includes both DS8000 storage subsystems.

1. At least one week in advance, run a storage subsystem performance monitor to collect performance data for both DS8000 subsystems.
2. Navigate to **IBM Tivoli Storage Productivity Center > Analytics**.
3. Right-click **Storage Optimizer** and select **Create Analysis**. The Create Analysis panel lets you specify the storage subsystems that you want to analyze.
4. Under Create Analysis, enter ″Analyze DS8000 subsystems″ in the Description field.
5. Click **Add** in the Select Storage pane. The Optimizer Selection panel opens. This provides a topology view of the current storage network.
6. Select both of your DS8000 storage subsystems.
7. Click **OK** to return to the Storage Optimizer Create Analysis panel with the DS8000 subsystems displayed in the Select Storage pane.
8. Set the Performance Time Interval to start one week prior to the current date.
9. For How often to run, select **Run Now**.
10. Choose **File > Save** to save and submit the analysis report job.
11. When prompted to enter a name for the report job, enter ″Analyze DS8000 subsystems″ and click **OK**.

12. View the analysis report job status for ″Analyze DS8000 subsystems″ under the **Storage Optimizer** node in the navigation tree.

Part 2: View the analysis report and identify hot spots

When the analysis report job ″Analyze DS8000 subsystems″ is complete (status is green), select it to view the analysis report for your DS8000 subsystems. The heat maps display color-coded cells that represent the projected utilization of both DS8000 subsystems, as measured against the desired performance threshold of 80%. Each cell represents a storage pool.

1. From the list of components for **Heat Map based on**, choose Hard Disk to update the heat maps to display the analysis data for that component.
2. Move the performance threshold slider to see its affect on the heat maps.
3. Click the **Performance Tables** tab to display a tabular view of the data, if desired.
4. In the Hard Disk heat map for one of your DS8000 subsystems, you notice a red cell, or ″hot spot″. Move your cursor over that cell to display the details for the storage pool named Pool 1. The details for Pool 1 indicate that it is exceeding the desired performance threshold of 80%.

   Just to the left of the red cell is a green cell for Pool 2. The green color indicates that Pool 2 is performing at less than or equal to 25% of the established performance threshold.

You think that you might be able to improve the performance of Pool 1 by moving some of its volumes to Pool 2, and you know that Pool 2 has some amount of available capacity.

Part 3: Create an optimization report that includes recommendations on how to improve the performance of Pool 1.

1. In the analysis report, select the red cell which represents Pool 1.
2. Click **>>** next to **Source Entities** to move Pool 1 into the list of source entities.
3. Select the green cell which represents Pool 2.
4. Click **>>** next to **Target Entities** to move Pool 2 into the list of target entities.
5. Choose **File > Save** to save and submit the optimization report job.
6. When prompted to enter a name for the optimization report job, enter ″Optimize DS8000 subsystems″ and click **OK**.
7. Under **View Previously Run Optimization Reports**, click **Refresh Job Status** to update the report job status.

Part 4: View the optimization report and recommendations.

When the optimization report job ″Optimize DS8000 subsystems″ is complete (job status is ″Success″), click  to display the optimization report. The report includes a recommendation for how to improve the performance of Pool 1.

The heat maps in the optimization report display the performance utilization of Pool 1 and Pool 2 before and after implementing the recommendation. The heat maps show you that the red cell representing Pool 1 changes to green after implementing the recommendation. The heat maps also show you that the green cell representing Pool 2 changes to yellow after implementing the recommendation. This means that Pool 2 will be still be performing at less than or equal to 75% of the performance threshold after you implement the recommendation.

To view the effects on the before and after heat maps of implementing only some of the recommendations:

- Remove the check mark next to **Select all recommendations** and select a subset of the recommendations from the recommendations table.

To print the optimization report:

- Click the Print button in the optimization report.

The printed report includes additional information such as SVC sample scripts to help you implement the recommendations using the SAN Volume Controller command-line interface.

Retiring underutilized storage subsystems

This scenario illustrates how you can use the Storage Optimizer to create recommendations about retiring underutilized storage subsystems and migrating their data to other subsystems in your data center.

Your data center has seven DS8000 storage subsystems in use. Using various Tivoli Storage Productivity Center capacity reports, you have determined that two of the subsystems are underutilized. You want to create a report that provides recommendations to help guide you in retiring the underutilized subsystems and migrating the data into your remaining five subsystems.

Prerequisites for using Storage Optimizer

- IBM Tivoli Storage Productivity Center Standard Edition license
- One or more of these supported IBM subsystems: DS8000, DS6000, DS4000, Tivoli Storage Enterprise Storage Server, SAN Volume Controller.
- At least one week's worth of performance monitoring data gathered in advance of using Storage Optimizer. Providing a longer time interval for data collection will increase the accuracy of the analysis and recommendations.

Monitor the performance of your DS8000 subsystems

At least one week in advance, perform performance monitoring for all your DS8000 storage subsystems. The general steps are as follows:

1. Run a CIMOM discovery for storage subsystem CIMOMs. See "Discover storage subsystem, tape library, and fabric information" on page 197 for more information.
2. Create a probe. See "Creating probes" on page 203 for more information.
3. Create a storage subsystem performance monitor and let it run for at least one week. See "Creating storage subsystem performance monitor" on page 242 for more information.

To perform this scenario, complete the following steps:

Part 1: Create an analysis report for the seven DS8000 storage subsystems in your data center.

1. At least one week in advance, run a storage subsystem performance monitor to collect performance data for all DS8000 subsystems.
2. Navigate to **IBM Tivoli Storage Productivity Center > Analytics**.
3. Right-click **Storage Optimizer** and select **Create Analysis**. The Create Analysis panel lets you specify the storage subsystems that you want to analyze.

4. Under Create Analysis, enter ″Analyze DS8000 subsystems″ in the Description field.
5. Click **Add** in the Select Storage pane. The Optimizer Selection panel opens. This provides a topology view of the current storage network.
6. Select all your DS8000 storage subsystems.
7. Click **OK** to return to the Storage Optimizer Create Analysis panel with the DS8000 subsystems displayed in the Select Storage pane.
8. Set the Performance Time Interval to start one week prior to the current date.
9. For How often to run, select **Run Now**.
10. Choose **File > Save** to save and submit the analysis report job.
11. When prompted to enter a name for the report job, enter ″Analyze DS8000 subsystems″ and click **OK**.
12. View the analysis report job status for ″Analyze DS8000 subsystems″ under the Storage Optimizer node in the navigation tree.

Part 2: View the analysis report.

When the analysis report job ″Analyze DS8000 subsystems″ is complete (status is green), select it to view the analysis report for your DS8000 subsystems.

- For **Heat Map based on**, choose the **Space** option to view the heat maps in the analysis report based on disk space utilization. The heat maps confirm that the majority of pools in subsystems 6 and 7 are performing at less than 75% of capacity.

Part 3: Create an optimization report that includes recommendations on how to retire the underutilized subsystems 6 and 7 and migrate the data to your remaining five subsystems.

1. In the heat maps, select all pools for storage subsystems 6 and 7 using the **Select all pools for storage subsystem** checkbox.
2. Click **>>** next to **Source Entities** to move subsystems 6 and 7 into the list of source entities.
3. Click **Make retirement recommendations for the selected entity**.
4. In the heat maps, select all pools for storage subsystems 1 through 5 using the **Select all pools for storage subsystem** checkbox.
5. Click **>>** next to **Target Entities** to move subsystems 1 through 5 into the list of target entities.
6. Choose **File > Save** to save and submit the optimization report job.
7. When prompted to enter a name for the report job, enter ″Retire DS8K″ and click **OK**.
8. Under **View Previously Run Optimization Reports**, click **Refresh Job Status** to update the report job status.

Part 4: View the optimization report and recommendations.

When the optimization report job ″Retire DS8K″ is complete (job status is ″Success″), click  to display the optimization report. If it is possible to retire subsystems 6 and 7, the report will include recommendations for retiring subsystems 6 and 7 and migrating the data into subsystems 1 through 5.

- For **Heat Map based on**, choose **Space** to view the before and after optimization heat maps based on disk space utilization.

- To view the effects on the before and after heat maps of implementing only some of the recommendations, remove the check mark next to **Select all recommendations** and select a subset of the recommendations from the recommendations table.

The before and after heat maps illustrate the effects of retiring subsystems 6 and 7 and migrating the data into subsystems 1 through 5. The heat maps also confirm that even after retiring subsystems 6 and 7, the remaining subsystems will still perform within the limits of the desired performance threshold.

To print the optimization report:

- Click the Print button in the optimization report.

The printed report includes additional information such as SVC sample scripts to help you implement the recommendations using the SAN Volume Controller command-line interface.

Identifying and investigating a suspended Metro Mirror session

IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication work together to help meet your storage management needs and automate your replication management tasks. This scenario shows you how to use IBM Tivoli Storage Productivity Center to view and investigate a suspended Metro Mirror between two IBM DS8000 storage systems that was detected by IBM Tivoli Storage Productivity Center for Replication.

You want to use Tivoli Storage Productivity Center to analyze alerts generated by events detected by Tivoli Storage Productivity Center for Replication. You want to open a page in Tivoli Storage Productivity Center for Replication from the Tivoli Storage Productivity Center user interface.

Prerequisites

- Two DS8000 storage systems.
- Use Tivoli Storage Productivity Center for Replication to mirror a database application between the DS8000 storage systems.
- Set up Tivoli Storage Productivity Center to monitor DS8000 storage systems. The general steps are as follows:
  1. Run a CIMOM discovery to identify the storage systems managed by CIMOMs in your environment. See *Discovering Storage Resources* in the *IBM Tivoli Storage Productivity Center User's Guide* for more information about how to discover storage systems.
  2. Create and schedule a probe to run against the storage systems. See *Collecting storage statistics (probes)* in the *IBM Tivoli Storage Productivity Center User's Guide* for more information about how to define probes.

Tivoli Storage Productivity Center includes a predefined set of alerts that correspond to events detected by Tivoli Storage Productivity Center for Replication. In this scenario, Tivoli Storage Productivity Center for Replication detects that a Metro Mirror session just went into a suspended state. That event is communicated to Tivoli Storage Productivity Center and recorded in the alert log. Use the alert log to view information about that alert and start the Tivoli Storage Productivity Center for Replication user interface. Use Tivoli Storage Productivity Center for Replication to determine that the paths between your DS8000 devices in a Metro Mirror session were removed and resolve the problem.

The following tasks are described in this scenario:

1. Tivoli Storage Productivity Center: Set up notification methods for replication alerts.
2. Tivoli Storage Productivity Center: Use the alert log to view information about a suspended Metro Mirror session that was detected by Tivoli Storage Productivity Center for Replication.
3. Tivoli Storage Productivity Center for Replication: Start Tivoli Storage Productivity Center for Replication from the Tivoli Storage Productivity Center user interface to access the Session Details page for the alert and resolve the problem.

To use Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication to analyze storage events, complete the following steps:

1. **Tivoli Storage Productivity Center: Set up notification methods for replication alerts.**
   a. Start Tivoli Storage Productivity Center.
   b. Expand **Replication Manager → Alerting → Replication Alerts** in the navigation tree.
   c. Click **TPCUser.Replication Suspending Event**. The Edit Replication Alerts panel is displayed. Use this panel to determine how you are notified if Tivoli Storage Productivity Center for Replication detects a suspending event in your storage environment.
   d. In the **Triggered Actions** section, select the notification methods for the alert. By default, all suspending events that are detected by Tivoli Storage Productivity Center for Replication are recorded as entries in the Alert Log panel under **IBM Tivoli Storage Productivity Center → Alerting**. If you want to be notified using additional methods when the alert log is updated for a suspending event, select any of the following options:
      - Select **SNMP Trap** to generate an SNMP trap message to any NMS, console, or terminal to indicate the triggering of a suspending alert.
      - Select **TEC/OMNIbus Event** to send an alert to the Tivoli Event Console or OMNIBus.
      - Select **Login Notification** to be notified of a suspending event when you log in to Tivoli Storage Productivity Center.
      - Select **Windows Event Log** or **UNIX SYSLOG** to write the suspending alert to the operating system log file.
      - Select **Run Script** to run a script when a suspending alert is triggered.
      - Select **E-mail** to indicate the e-mail addresses that receive notification when a suspending alert is triggered.

      You are not required to select any of these options. Triggered alerts from Tivoli Storage Productivity Center for Replication are automatically recorded in the alert log. These options are meant to provide you with additional methods for receiving alerts.
   e. Click the **Save** icon to save the alert definition. When Tivoli Storage Productivity Center for Replication detects that a session has suspended, this alert is triggered, an entry is added to the Tivoli Storage Productivity Center alert log, and any additional triggered actions are run.
2. **Tivoli Storage Productivity Center: Use the alert log to view information about a suspended Metro Mirror session.** This step occurs after Tivoli Storage Productivity Center for Replication has detected a suspended session and triggered an alert in the Tivoli Storage Productivity Center user interface.

a. Expand **IBM Tivoli Storage Productivity Center → Alerting → Alerting Log** in the navigation tree. A visual indicator is displayed next to each alert type under this node. In this scenario, a red, filled-in circle is displayed next to the **Replication** node to indicate that Tivoli Storage Productivity Center has received an alert from Tivoli Storage Productivity Center for Replication that is unresolved.

b. Click the **Replication** node. The content pane displays a table that contains a row for each alert received from Tivoli Storage Productivity Center for Replication. The **Alert Type** column in this table indicates the type of alert that was triggered.

c. Click the magnifying glass icon next to the alert that displays **Suspending Event** in the **Alert Type** column. The Detail for Alert panel displays a description of the event that triggered the alert:

```
IWNR1958W [2009/03/04 15:11:49.250 EST] Session [name] has
changed to the SEVERE status due to an unexpected error.
```

where *[name]* represents the name of the session in Tivoli Storage Productivity Center for Replication.

d. Click the back arrow in the tool bar to return to the Alert Log panel.

3. **Start Tivoli Storage Productivity Center for Replication from a suspended event alert in the alert log.** The Session Details page indicates that the session went suspended because the paths between the devices were dropped.

a. On the **IBM Tivoli Storage Productivity Center → Alerting → Alerting Log** page, right-click the row that represents the **Suspending Event** alert type. The menu that is displayed contains the actions that you can perform against the alert log entry. One of the actions is to start the Tivoli Storage Productivity Center for Replication user interface.

b. Click **Launch Replication Session Details** on the menu. The Tivoli Storage Productivity Center for Replication user interface starts in a Web browser and displays the Session Details page. Use this page to view detailed information about the suspended event.

Single sign-on is an authentication process that enables you to enter one user ID and password to access multiple applications. Single sign-on integrates with the launch in context feature to enable you to move smoothly from the Tivoli Storage Productivity Center user interface to a specific location in the Tivoli Storage Productivity Center for Replication user interface without re-entering a user ID and password.

For example, this scenario shows how you can use single sign-on and launch in context to open a page in the Tivoli Storage Productivity Center for Replication user interface based on a triggered alert in Tivoli Storage Productivity Center.

c. Use Tivoli Storage Productivity Center for Replication to resolve the problem that triggered the alert.

You can use steps similar to the ones above to use Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication together to investigate and resolve different conditions and events that occur within your storage environment.

Chapter 3. Administering IBM Tivoli Storage Productivity Center

This topic provides information on how to administer IBM Tivoli Storage Productivity Center.

Administering data sources

Use the following topics for information about how to administer CIMOM agents, Data agents, Storage Resource agents, Inband Fabric agents, Out of Band Fabric agents, IBM Tivoli Storage Productivity Center servers, and VMware VI data sources.

Checking the health of your agents

This topic provides information about checking the health of your agents.

Expand **Administrative Services** → **Data Sources** and left-click on the following nodes:

- **CIMOM Agents**
- **Data/Storage Resource Agents**
- **Inband Fabric Agents**
- **Out of Band Fabric Agents**

A list of agents is displayed in the right pane. If these agents are up and running, there is a green health status icon in the State column. For the CIMOM, a green health status icon is displayed in the Connection Status column.

Manually changing the Windows service logon for the agent

This topic describes how to change the Windows service logon for the Common agent.

To change the Windows service logon for the Common agent, complete the following steps:

1. In Windows, open the Services panel: **Start** → **Settings** → **Control Panel** → **Administrative Tools** → **Services**.
2. On the Services panel, right-click **IBM Tivoli Common agent - 'C:\Program Files\IBM\TPC\ca'**. Select **Properties**.
3. Select the Log On tab. Change **This account** field with your changed logon ID. If your Tivoli Storage Productivity Center server is part of a Windows domain, change this logon to `<domain>\<account>`. For example, **mydomain\myaccount**. Enter the password if you have changed the password. Click **Enable** and then **OK**. The Common agent requires that the domain account have local administrator privileges as well as ″Log on as a service″ and ″Act as part of the operating system″ user rights.

CIM agents

Common Information Model (CIM) agents provide a CIM interface for management applications. These include IBM Tivoli Storage Enterprise Storage Server (Tivoli Storage Enterprise Storage Server), Storage Area Network SAN Volume Controller, IBM System Storage DS4000, IBM System Storage DS6000, IBM

System Storage DS8000, IBM XIV Storage System, tape, switch and any SMI-S certified Common Information Model Object Manager (CIMOM).

Important: XIV Storage System will have an embedded CIM agent that Tivoli Storage Productivity Center will use to run discovery and probe jobs. Tivoli Storage Productivity Center support for this CIM agent is targeted for a future XIV Storage System software release. The XIV Storage System information provided in the Tivoli Storage Productivity Center 4.1 documentation is only for planning purposes until the supported XIV Storage System software is available. A flash will be issued when Tivoli Storage Productivity Center support for XIV Storage System is available.

Adding a CIM agent manually

CIM agents provide a Common Information Model (CIM) interface for management applications. This topic discusses how to manually add a CIM agent.

To add a CIM agent to the CIMOM node, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services → Data Sources** and click **CIMOM Agents**.
2. In the right pane, click **Add CIMOM**. The Add CIMOM window opens.
3. Specify information about the CIM agent:

   Host The hostname of the machine on which the CIMOM is installed. You can enter a host name, IPv4, or IPv6 address depending on what is supported within your environment:

   - Enter an IPv4 address in the following format: ###.###.###.###. For example: 192.168.1.65
   - Enter an IPv6 address. 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:DB8:0000:1234:0000:0000:5678: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.

     You can enter an IPv6 address for a CIMOM if the machine where IBM Tivoli Storage Productivity Center is installed is IPv6 or dual stack (IPv4 and IPv6) enabled.

   Port The port on which the CIMOM is listening. By default this is 5989 for a secure connection and 5988 for an unsecure connection. See the documentation for the appropriate SMI-S CIM Provider for more information about available ports.

   Username
   The user ID used for authentication, if required by the CIMOM.

   Password and Password Confirm
   The password, if required by the CIMOM.

   Interoperability Namespace
   This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how IBM Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information. The following are the default namespaces for CIMOM agents for switches:

   - IBM: /root/ibm

- Brocade: /interop
- NetApp: /interop
- Cisco: /root/cimv2

  For version 3.2.1 or later: /root/pg_interop
- Engenio: /interop
- EMC: /root/emc
- HDS: For HiCommand 5.6 or higher, use /root/smis/current

  For a HiCommand version that supports SMI-S 1.2, use following namespace to traverse the model as SMI-S 1.1: /root/smis/smis11

  For HiCommand versions lower than HiCommand 5.6, use: /root/hitachi/dmxx, where xx is represents the level of HiCommand.
- HP: /root
- McData: /interop
- SUN (storage subsystem): /root/sun3510 or /interop
- XYRATEX: /root/PG_interop

To view a list of default namespaces for CIMOM agents for storage subsystems that is included with Tivoli Storage Productivity Center, open the following file in a text editor: *install_directory*/data/config/namespace.config, where *install_directory* represents the directory where the product is installed.

Check the documentation of the appropriate SMI-S CIM Provider or contact the storage system vendor to ensure you use the most current namespaces.

Protocol
The version of the cim-xml protocol. Can be http or https.

Truststore Location
The location (path on this computer) of a certificate file for certificate based authentication in the https protocol. This field applies only to certain non-IBM devices. See the documentation for the appropriate SMI-S CIM Provider for more information about the truststore location.

Truststore Passphrase
The passphrase for the truststore for the CIMOM.

Display Name
The name of the CIMOM, as specified by the CIMOM provider, that will appear in the Tivoli Storage Productivity Center interface.

Description
The optional description.

Test CIMOM connectivity before adding
Check box. Check this box to have Tivoli Storage Productivity Center ensure that communication to the CIMOM is working properly before adding information about that CIMOM.

4. Click **Save**.

Note: This procedure does not physically add a CIMOM, but adds the information you have given about the CIMOM to Tivoli Storage Productivity Center so that Tivoli Storage Productivity Center can communicate with the CIMOM.

Discover CIM agents automatically

This topic discusses how to discover storage subsystem CIMOMs that are visible to the Device server using Service Location Protocol (SLP).

In many cases CIM agents can be automatically discovered using SLP. The criteria is that the CIM agent must be visible to the Device server.

To automatically discover all CIM agents that are visible to the Device server, complete the following steps:

1. In the navigation tree pane, expand **Administrative Services** → **Discovery** and click **CIMOM**.
   a. Click the **Options** tab.
   b. Enter the IP address or host name for the SLP directory agents for CIMOM discovery. Or select the **Scan local subnet** check box.
   c. Click **File** → **Save** on the menu bar to save the information.
2. Right-click **CIMOM** and click **Run Now**.
3. After you submit a CIMOM discovery job, a message window is displayed with the text **CIMOM Discovery Job submitted**. Click **OK**. The job name is located below **CIMOM**. While it is running, the CIMOM job has a blue circle outline to the left of the job name.
4. To determine if the job has completed, right-click **CIMOM** and click **Update Job Status**. If the job has completed successfully, a green square is shown in front of the job name. If the job has completed but failed, a red circle is shown.

   Note: The job returns as failed if one or more sub-jobs failed. Some CIMOMs might still have been discovered.
5. After the job completes, the discovered CIM agents are listed in the **CIMOM Agents** table. Expand **Administrative Services** → **Data Sources** and click **CIMOM Agents**.
6. For each CIM agent that you want to use, highlight the CIM agent and click the magnifying glass at the left of the listing. You must select the CIM agents one at a time. The **CIMOM Agents** window is opened.
7. Enter additional information for the CIM agent, such as the user name and password, and select the **Test CIMOM connectivity before updating** check box.
8. Click **File** → **Save** on the menu bar to save the information and perform a connectivity check.

Note: This procedure does not physically add the CIM agent to your environment, but provides IBM Tivoli Storage Productivity Center with the information it needs to communicate with that CIM agent. This includes information that you have defined and CIM agents discovered using SLP.

Viewing information about a CIM agent

Use panels and reports in IBM Tivoli Storage Productivity Center to view detailed information about a CIM agent.

Use the following reports to view information about a CIM agent:

- In the Navigation Tree pane, expand **IBM Tivoli Storage Productivity Center** → **Reporting** → **Data Source Reports** → **CIMOM Agents** and click **By CIMOM Agent** or **By Managed Device**. See "Data source reports - CIMOM agents" on page 766 for more information about these reports.

- In the Navigation Tree pane, expand **Administrative Services → Data Sources** and click **CIMOM Agents**. A list of CIM agents appears in the content pane. Click the magnifying glass next to a CIM agent to view the following information about that agent:

  Service URL
  The service URL of the CIMOM containing the IP address of the CIMOM, the port on which the CIMOM is listening, and the protocol used for communication. This URL has a protocol [http|https], an IP or Hostname, and a port number. This field displays IPv4 and IPv6 addresses as appropriate.

  Display Name
  The name of the CIMOM as specified by the CIMOM provider that will appear in the IBM Tivoli Storage Productivity Center interface.

  Description
  The optional description that was entered on the Add CIMOM window.

  Username
  The CIMOM user ID used for authentication.

  Password and Password Confirm
  The password for the CIMOM.

  Interoperability Namespace
  Enter the interoperability namespace of the CIMOM. This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information.

  Truststore Location
  The location (path on this computer) of a certificate file for certificate based authentication in the https protocol.

  User Interface Description
  The name of the Human Interface Service (if any) supported by this CIMOM.

  Software Level
  The software version level of the CIMOM agent.

  Protocol Version
  The version of the cim-xml protocol.

  Authentication Mechanism
  The authentication mechanism supported by the CIMOM. This field can contain the following values: Unknown, None, Other, Basic, Digest.

  Alias The alias of the CIMOM.

  Service ID
  The service ID for the CIMOM.

  Protocol
  The communication protocol used for the CIMOM. Possible values are http and https.

  SLP Attributes
  The standard set of attributes for this CIMOM. The attributes are retrieved via SLP.

Connection Status
The status of this CIMOM with respect to Tivoli Storage Productivity Center. Possible values are: SUCCESS, UNCONFIGURED, UNKNOWN, INVALID_NAMESPACE, TIMEOUT, REFUSED, LOGIN_FAILED, SSL_HANDSHAKE_ERROR, SSL_REGISTRATION_INVALID, CIMCLIENT_ERROR

Status Timestamp
The date/time when the Connection Status information was last collected.

Test CIMOM connectivity before updating
If this box is checked Tivoli Storage Productivity Center ensures that communication to the CIMOM is working properly before making any changes.

Updating CIM agent user identification

After a CIM agent has been added to the **CIMOM** node, you can update the user ID, password, and other information.

To update the CIM agent user identification, follow this procedure:

1. In the Navigation Tree window, expand **Administrative Services → Data Sources**. Left-click **CIMOM**.
2. In the right pane, click on the icon to the left of the agent for which you want to change information.
3. Information about the CIM agent is displayed in the Information window. The fields which can be updated are:

   Display Name
   The name of the CIM agent, as specified by the CIMOM provider, that will appear in the IBM Tivoli Storage Productivity Center interface.

   Description
   The name of the CIM agent, as specified by the CIMOM provider, that will appear in the Tivoli Storage Productivity Center interface.

   Username
   The user ID used for authentication, if required by the CIM agent.

   Password and Confirm Password
   The password for the CIM agent, if required by the CIM agent.

   Interoperability Namespace
   This namespace within the CIM agent allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIM agent when retrieving information.

   Test CIMOM connectivity before updating
   Check box. Check this box to have Tivoli Storage Productivity Center ensure that communication to the CIM agent is working properly before updating information about that CIM agent.
4. Save your changes by clicking the **Save** icon.
   - If you have selected **Test CIMOM connectivity before updating**, the CIMOM connectivity check will be run when you save the changes.
   - If **Test CIMOM connectivity before updating** is not selected, a discovery will be started immediately after you save the changes.

Testing a CIM agent connection

Manually check to make sure that communication to the CIM agent is working properly.

To test a CIM agent connection, follow this procedure:

1. In the Navigation Tree, expand **Administrative Services** → **Data Sources**. Left-click **CIMOM**.
2. In the right pane, select a CIM agent. Click **Test CIMOM Connection**. a confirmation dialog appears: ″Testing CIMOM connectivity can take up to several minutes in case of an incorrectly entered port number, network problems or an unpassed firewall. Would you like to continue anyway?″ Select **yes** to perform the CIMOM connection test.
3. To close the Test CIMOM Connection window, click **OK**.

Removing a CIM agent

You can remove a CIM agent from the navigation tree and the repository database. Data discovered by the CIM agent is not removed from the repository.

To remove a CIM agent, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services** → **Data Sources**. Left-click **CIMOM Agents**.
2. In the right pane, select a CIM agent. Click **Remove CIMOM**.
3. The Remove CIMOM confirmation window opens. Click **Yes**. The CIM agent is immediately deleted from the list.

Show managed devices for a CIM agent

You can display the managed devices for a CIM agent.

To show the managed devices for a CIM agent, follow these steps:

1. In the Navigation Tree pane, expand **Administrative Services** → **Data Sources**. Left-click **CIMOM Agents**.
2. In the right pane, select a CIM agent. Click **Show Managed Devices**.
3. The CIMOM Managed Devices window opens displaying the managed devices. Click **OK** to close the window.

Collecting CIM agent logs

You can collect logs for certain IBM CIM agents using the command line interface.

1. Change to the directory where the CIM agent is installed.

DS3000, DS4000, DS6000, DS8000	On Linux	/opt/IBM/cimagent/cimom
	On Windows	C:\Program Files\IBM\cimagent\ cimom
SAN Volume Controller	On Linux	`<usr or opt>/IBM/svcconsole/ support`
	On Windows	`C:\Program Files\IBM\svcconsole\ support`

2. Run one of the following commands:

On Linux	**collectLogs.sh**
On Windows	**collectLogs.bat**

A **collectedLogs.zip** file is created.

Attention: this file is overwritten if you run the script again.

Collecting XIV CIM agent logs

The log files for the XIV Storage System CIM agent are bundled with the XIV Storage System system logs.

The log files for the XIV Storage System CIM agent are collected using the XIV Storage System XCLI commands. These commands are available on any host that has the XIV Storage System GUI installed or on the system where Tivoli Storage Productivity Center Element Manager starts the XIV Storage System GUI. You use the XCLI commands to bundle the log files into compressed tar files on the XIV Storage System device and then you use a web browser to retrieve those tar files using HTTP.

The XCLI command utility is included with the XIV Storage System GUI and is located by default at the following location: `C:/Program Files/XIV/GUI10`.

The following command prints the help text for the XCLI command utility:

```
xcli -h
```

To collect the log files, perform the following steps:

1. Enter the following command to list the known XIV Storage System configurations that were set up using the XIV Storage System GUI:

   ```
   xcli -L
   ```

2. Using the information gathered in Step 1, note the IP address of the XIV Storage System for which you want to collect the logs.
3. Enter the following command to collect the XIV Storage System system logs for the XIV Storage System identified in Step 2:

   ```
   xcli -u <admin_user> -p <admin_password> -m <ip_of_the_XIV_system> system_logs_collect
   ```

   where:

   - *admin_user* is the administrator user name.
   - *admin_password* is the administrator password.
   - *ip_of_the_XIV_system* is the IP address of the XIV Storage System to use for collecting the log files.

   For example, the following command collects system logs from the IP address 129.42.58.216:

   ```
   xcli -u admin -p abcabc -m 129.42.58.216 system_logs_collect
   ```

4. Since log collection may take some time, enter the following command to monitor the log collection process:

   ```
   xcli -u <admin_user> -p <admin_password> -m <ip_of_the_XIV_system> system_logs_get_status
   ```

   Log collection is complete when you see status message: `system logs are available over HTTP.`
5. To access the logs, start a web browser and go to the following URL: `http://<ip_of_the_XIV_system>.`
6. When prompted, download the tar files and make a note of the download directory. The tar file will have a name like `system_xray_2810A14unknown_2009-02-03-2319[1].tar.bz2`.
7. Locate the tar file and use an archive manager utility such as WinRAR to extract the log files from the tar archive.
8. Locate the XIV Storage System CIMOM log files by navigating to the following path within the archive: `xray\7\FS\local\cim\log`.

Verifying that a CIM agent is running

You can verify that a CIM agent is running from the command line interface.

To verify that a CIM agent is up and running, run the following command:

```
telnet <IP> <port>
```

Where <IP> is the IP address of the system where the CIM agent is installed, and <port> is the port number. By default, this is 5989 for a secure connection and 5988 for an unsecure connection.

Data agents and Storage Resource agents

This topic provides information about administering, configuring, and viewing information about Data agents and Storage Resource agents.

See "Using Data agents or Storage Resource agents" on page 75 for information about the differences between Data agents and Storage Resource agents.

Viewing information about an agent

From the IBM Tivoli Storage Productivity Center GUI, you can view detailed information about a Data agent or Storage Resource agent.

To view information about a Data agent or 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 right pane, click the magnifying glass icon to the left of a Data agent or Storage Resource agent under the list of agents. In the right pane, a notebook window opens. On the General page, the following information is displayed:

   Agent Displays the name of an agent and indicates whether the agent is enabled or disabled.

   Port Displays the port number on which the agent is listening for requests.

   Host Address
   Displays the host address of the computer on which an agent is running.

   Last update
   Displays the date and time when an agent was last updated.

   Timezone
   Displays the current time zone for an agent.

   Connection Errors
   Displays any connection errors that occurred for an agent.

   Consecutive Errors
   Displays how many errors occurred consecutively.
3. Click the **Details** tab to view more detailed information. The following information is displayed:

   Agent Displays the name of the agent.

   Host Displays the host name of the system on which the agent is running.

   Host ID
   Displays the identifier of the host computer in hexadecimal format.

Tivoli Storage Productivity Center only collects Host ID information for Solaris systems. All other platforms will display GUID information in the Host ID column.

Start Time
Displays the date and time that the agent was started.

Elapsed Time
Displays the time elapsed since the agent was last updated.

VM Size
Displays the size (in MB) of the disk space that is allotted to virtual memory for the agent.

Manufacturer
Displays the manufacturer of the system.

OS Type
Displays the operating system that is running on the system.

CPU Architecture
Displays the architecture of the system.

Job Count
Displays the number of jobs scheduled to run against the agent.

4. Click the **Jobs** tab to view information about the jobs scheduled to run on the managed system.

Viewing an agent log file

The log file for a Data agent or Storage Resource agent contains informational, warning, and error messages for the previous sessions of the agent. You can use the log file to troubleshoot errors that might occur when starting an agent, processing, or shutting down the agent.

The log file is located in the agent installation directory:

Storage Resource agents
For Windows default: `C:\Program Files\IBM\TPC\agent\log\<computer_name>\agent.log`.

For Linux or AIX (default): `/opt/IBM/TPC/agent/log/<computer_name>agent.log`.

Data agents
For Windows (default): `C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\log\<computer_name>`.

For Linux or AIX (default): `/opt/IBM/TPC/ca/subagents/TPC/Data/log/<computer_name>`.

For Data agents, by default, information from the last five sessions of the agent appears in the log window. A session starts when the agent is started and ends when the agent stops. You can increase or decrease how many sessions are kept in the history file by changing the value for the **logFilesKept** parameter in the `agent.config` file. The default location for the `agent.config` file is:

Windows
`C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\config\`

Linux or UNIX
`/opt/IBM/TPC/ca/subagents/TPC/Data/config`

To view the log for a Data agent or 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 right pane, select a Data agent or Storage Resource agent and click **View Log**.

Configuring tracing for an agent

You can configure and enable tracing for a Data agent or Storage Resource agent. This provides extensive logging of exception and trace messages; it can provide useful information for the IBM Support center.

You must have administrator authority to perform this procedure.

To configure tracing for a Data agent or 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 right pane, select a Data agent or Storage Resource agent. Click **Configure Tracing**. The Tracing Configuration window opens.
3. Select the **Enable Trace** check box.
4. Specify the tracing options:
   a. In the **Level** field, select one of the following options:

      DEBUG_MIN
      Minimum level of debugging

      DEBUG_MID
      Medium level of debugging

      DEBUG_MAX
      Maximum level of debugging

   b. In the **Maximum Number of Files** field, specify the maximum number of trace files that are created. When this number is reached, the oldest file is overwritten. By default, this is set to five.
   c. In the **Maximum File Size (kilobytes)** field, specify the maximum size of each trace file.
5. Click **OK**. The settings are saved to the **AgentTraceLog.config** file

Reading the configuration file (Data agent)

If you make changes to the `agent.config` file and want to verify those changes, you can read the configuration file.

The `agent.config` file contains parameters that you can change to customize the behavior of an agent. The default directory for the `agent.config` file is:

Windows

```
C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\config\
```

Linux or UNIX

```
/opt/IBM/TPC/ca/subagents/TPC/Data/config
```

To verify the contents of the configuration file, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources > Data/Storage Resource agents**.
2. In the right pane, select a Data agent. Click **Read Config**.

3. The configuration file is displayed.

Disabling an agent

You can disable a Data agent or a Storage Resource agent. Disabling the agent signifies that the Data agent or Storage Resource agent is unavailable and the Data server should not contact that agent for any job processing.

You might want to disable a Data agent or Storage Resource agent in the following situations:

- The agent computer is undergoing maintenance and will be unavailable. This prevents the server from flagging the agent as ″down″ if it cannot reach the agent. The amount of times that the server tries to contact the agent is defined by the **agentErrorLimit** parameter in the `server.config` file.
- The agent computer is busy with resource-intensive processing and you do not want to add any jobs to that processing load.

While disabled, the agent name in the user interface is outlined with a red circle with a line through it. The server does not attempt to contact the agent.

To disable a Data agent or 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 right pane, select a Data agent or Storage Resource agent and click **Disable**.

Enabling an agent

You can enable an agent that is currently disabled. After it is enabled, the Data server resumes communication with that agent. This action is available for disabled agents only.

If the server cannot contact an agent, the server automatically flags the agent as ″down″. Click **Enable** to reestablish communication between the agent and server. The number of times the server tries to contact the agent is specified in the **agentErrorLimit** parameter in the `server.config` file. The default is 3.

The default directory for the `server.config` file is:

Windows

```
C:\Program Files\IBM\TPC\Data\config
```

Linux or UNIX

```
/opt/IBM/TPC/Data/config
```

Tip: Clicking **Check** also enables an agent that is disabled.

To enable an agent, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the right pane, select an agent and click **Enable**.

Shutting down an agent

You can shut down Data agents andStorage Resource agents (run as daemon processes) from the graphical user interface.

To shut down an agent, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the right pane, select a Data agent or Storage Resource agent and click **Shutdown**. This button is not available for Storage Resource agents that are installed as non-daemons.
3. Select how you want an agent to shut down. You can select from the following options:

   Normal
   Shut down the agent software and let all running processes complete. The agent will continue to accept new jobs that are submitted and will not shut down until:
   - All running processes are complete.
   - There are no new jobs submitted.

   Immediate
   Shut down the agent software when the currently running processes complete. No new jobs will be accepted by the agent and shutdown will occur immediately after the last job completes.

   Abort Shut down the agent software and stop whatever processes are currently running.

   You cannot start a Data agent from the Data/Storage Resource agent panel in the IBM Tivoli Storage Productivity Center user interface. To start a Data agent, you must access the computer on which the agent is installed and start it manually.

   For Storage Resource agents that are run as daemon processes, you can use the **Start** button on the Data/Storage Resource agent panel to start those agents. See "Starting a Storage Resource agent" for more information about how to start Storage Resource agents.

Starting a Storage Resource agent

You can start a Storage Resource agent if the daemon process for that Storage Resource agent is down.

You can only perform this task for Storage Resource agents running as a daemon service.

To start 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 right pane, select a Storage Resource agent and click **Start**.

Checking whether an agent can communicate with the Data Server

You can check whether an agent can communicate with the Data server.

You can only perform this task for Storage Resource agents running as a daemon service.

To start 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 right pane, select an agent and click **Check**.

If a Data agent is selected, this operation verifies that the Data server can communicate with the selected agent.

If a Storage Resource agent is selected, this operation verifies that the Data server can communicate with the host system where the Storage Resource agent is installed with the appropriate authentication information. After verifying that the Data server can communicate with the selected agent, the process checks for the existence of the Storage Resource agent's runtime files.

Changing authentication for a Storage Resource agent

There are times when you might need to change the user ID, password, and certificate location for a Storage Resource agent.

To change the authentication for a Storage Resource agent, follow these steps:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the right pane, select one or more Storage Resource agents for which you want to change authentication information.
3. Click **Change Authentication**.
4. In the Change Authentication window, enter the following information:

   User Enter the user ID that IBM Tivoli Storage Productivity Center should use when logging on to the computer on which the Storage Resource agent is installed.

   Requirement: This ID must have administrative or root privileges when installing a Storage Resource agent on a computer.

   Password
   Enter the password.

   Re-type
   Reenter the password.

   Certificate Location
   Enter the fully qualified path of the certificate on the computer where the Data server is located. This file is used for certificate-based authentication. If you do not enter a value in this field the default location is used. When deploying Storage Resource agents, Tivoli Storage Productivity Center uses the Secure Shell (SSH) protocol to communicate with the target computer.

   Passphrase
   Enter the passphrase for the certificate file. The passphrase was created when the certificate was generated.

   Click **Save**.

Related Tasks:

See **Planning > General Planning > Planning for Storage Resource agents** section in the *IBM Tivoli Storage Productivity Center Installation and Configuration Guide* for additional configuration information about certificates used after deployment.

Adding Storage Resource agents

You can add Storage Resource agents to your system by scheduling a job or by adding an agent immediately.

To add a Storage Resource agent, complete the following steps:

1. In the navigation tree, you have two methods of adding a Storage Resource agent:
   - Click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**. In the right pane, click **Add Storage Resource agents**.
   - Click **Administrative Services > Configuration**. Right-click **Storage Resource Agent Deployments**. On the menu, click **Create Storage Resource Agent Deployments**.
2. The Create Storage Resource agent Deployments window opens displaying the Computers tab. Enter the following information:

   Header fields

   Creator
   Predefined creator name.

   Name Predefined name.

   Description (optional)
   Enter a user-defined description name for the job.

   Enabled
   Check this box to enable Storage Resource agent deployments.

   Add Host List
   Click this button to manually enter names and logon credentials for the computers on which you want to deploy Storage Resource agents.

   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. Enter the domain controller name and login information on the displayed window. On the next window, select the domain and click **Get List of Domain Computers**. The host names of computers that you add using this method are displayed in the **Remote Agent Machines** section of the window. In that section, you must also enter an installation location for each computer and indicate whether you are reinstalling a Storage Resource agent on a computer.
   - Click **Get Agent List From file** to retrieve a list of computer names from a flat text file. The computers that are listed in the file are automatically added to the deployment job and displayed in the **Remote Agent Machines** section of this window. In that section, you must also enter an installation location for each computer and indicate whether you are reinstalling a Storage Resource agent on a computer.

     The file that you retrieve must contain one host name or IP address per line. For example:

     ```
     systemxyz.storage.usa.ibm.com
     9.79.179.179
     systemzyx.storage.usa.ibm.com
     9.89.189.189
     ```

     If Tivoli Storage Productivity Center discovers any syntax problems in the file, none of the host entries in the file are added to the deployment job.

- Enter the computer names or IP addresses in the Remote Agent Machines table. Enter the installation locations. If you do not enter a value in this field the default location is used. The default installation directories are:
  - For Windows: C:\Program Files\IBM\TPC\agent
  - For Linux and UNIX: /opt/IBM/TPC/agent
- Select **Force** under the following circumstances:
  - If an earlier Storage Resource agent installation failed and there are damaged agent files on the computer that cause further installations to fail. If you select this option, Tivoli Storage Productivity Center attempts to overwrite the previous failed deployment on the computer with a new Storage Resource agent.
  - If you want an existing Storage Resource agent to communicate with an additional Tivoli Storage Productivity Center server. To do this, you must create the deployment job from the additional Tivoli Storage Productivity Center server to which you want the Storage Resource agent to communicate.

  Note:
  - You cannot change the communication type for a Storage Resource agent (daemon or non-daemon) when you select the **Force** option. Make sure to select the same communication type as the existing Storage Resource agent when you create a deployment job.
  - If a Storage Resource agent exists on a target computer and you do not select **Force**, an error occurs during validation and the Storage Resource agent is not installed.

Enter the following information:

User Enter the user ID that Tivoli Storage Productivity Center should use when logging in to the host computer to install a Storage Resource agent. The value in the field is applied to all the computers that appear in the **Remote Agent Machines** table.

Requirement: This ID must have administrative or root privileges on the target computer.

Password
Enter the password for the user ID.

Re-type
Reenter the password for the agent machine.

Certificate Location
Enter the fully-qualified path of the certificate on the computer where the Data server is located. This file is used for certificate-based authentication. If you do not enter a value in this field the default location is used.

Passphrase
Enter the passphrase for the certificate file. The passphrase was created when the certificate was generated.

Port Enter the port number on which a Storage Resource agent listens for requests. The default is 9510.

Tip: This value is required if you run the Storage Resource agent as a daemon service.

Use Daemon Service for Runtime Operation
Select this option to run a Storage Resource agent on the monitored computer as a daemon service.

See Storage Resource agent Deployment Considerations for more information about deploying agents as daemon and non-daemon services.

Restriction: Do not select this option to run the Storage Resource agent as a stand-alone executable file on the monitored computer.

Validate before Save
Check this box to indicate that Tivoli Storage Productivity Center should attempt to communicate with the computers that are added to the deployment job when you click the **Save** button on this page. If the validation fails for a computer, the agent is not installed on that computer.

Edit Selected Entries
Click this button to edit settings for the selected Storage Resource agents. This information includes user ID, password, certificate location, passphrase, and port. You must submit the deployment job for the changes to take effect.

Remove
Click this button to remove a Storage Resource agent from the deployment job.

Tip: This button does not remove the agent from Tivoli Storage Productivity Center if it is already installed; it only removes it from the deployment job.

3. Click the **When to run** tab to enter the following information:

How often to run
Specify a time to run:
- **Run now**
- **Run once at** (specify a date and time to run)

How to handle time zones
Specify a time zone to use:
- **Use the time zone that the server runs in**
- **Use this time zone** (select a time zone)

4. Click the **Alert** tab to specify the following information:

Triggering-Condition
The triggering condition you can specify is: Storage Resource agent Deployment Failed

Triggered-Actions
Choose from the following check boxes:
- **SNMP Trap**
- **TEC Event**
- **Login Notification**
- **Windows Event Log**
- **Run Script**
- **Email**

Depending on what action you select, you might have other choices to make. For example, if you select the **Windows Event Log** check box, the **Event Type** field becomes active so that you can specify a severity for the event in the Windows event log.

5. Click **File > Save**.

Related Tasks:

See the *Planning > General Planning > Planning for Storage Resource agents* section in the *Tivoli Storage Productivity Center Installation and Configuration Guide* for additional configuration information about certificates used after deployment.

Check the TPCD_#####.log file in the Tivoli Storage Productivity Center installation directory to view the error messages related to failed validations. The following are the default locations for this log file:

- Windows: C:\Program Files\IBM\TPC\data\log
- UNIX/Linux: /opt/IBM/TPC/data/log

Deleting or uninstalling an agent using the GUI

You can delete or uninstall a Data agent or 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 will no longer be available within IBM 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 Data agent or 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 right pane, select an agent and click **Delete**.

Enabling or disabling automatic upgrade for agents

This topic provides information for enabling or disabling the automatic upgrade function for a Data agent or Storage Resource agent.

For Data agents, the agent will automatically be upgraded to the same version as the Data server. The upgrade occurs when the Data agent registers with the Data server and IBM Tivoli Storage Productivity Center determines that the Data agent is running an older version from the Data server.

For Storage Resource agents, the agent is automatically upgraded to the same version as the Data server. The upgrade occurs when a probe is run against the Storage Resource agent and Tivoli Storage Productivity Center determines that the agent is running an older version from the Data server.

To enable the program so that a Data agent or Storage Resource agent can be automatically upgraded, follow these steps:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the right pane, select one or more agents for which you want an automatic upgrade performed.
3. Click **Enable Auto Upgrade**. Note that if you click this button it shows **Disable Auto Upgrade**.

To disable the automatic upgrade function for agents, select the agents for which you want the automatic upgrade discontinued. Click **Disable Auto Upgrade**. Note that if you click this button, it shows **Enable Auto Upgrade**.

Exporting a Storage Resource agent list

You can export a Storage Resource agent list to a formatted flat file. This list is used to manage user authentication through a script file. This list contains the host name, installation location, user ID, encrypted password, and fully qualified certificate file for the selected Storage Resource agents. The password is in the encrypted format with @ENC@ appended to the encrypted password. You can change the user authentication data such as user ID, password, and fully qualified certificate location.

This option is enabled for Storage Resource agents that use non-daemon based communication only. Information about Storage Resource agents that run as daemon processes on monitored computers is not exported to a file.

Storage Resource agents can either be run as a daemon service or non-daemon service (on-demand service). If the agent is run as a non-daemon service, there is no agent process running on the managed host. Programs on the host system are run as needed.

To export a Storage Resource agent list, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the right pane, click **Export Storage Resource Agent List**.

Updating a Storage Resource agent list

The Storage Resource agent list file is used to manage user authentication through a script file. This list contains the host name, installation location, user ID, encrypted password, and fully qualified certificate file for the selected Storage Resource agents. The password is in encrypted format with @ENC@ appended to the encrypted password. You can change the user authentication data such as user ID, password, and fully-qualified certificate location.

To update a Storage Resource agent list, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the right pane, select the Storage Resource agents you want to update.
3. Click **Export Storage Resource agent List** to export information about those agents to a file.
4. Edit the information in that file to match the changes that you want to make for the Storage Resource agents.
5. Click **Update Storage Resource agent List** and select the file to apply your changes. You can update the logon credentials for the Storage Resource agents that are listed with the information in a file that you specify. The file must use the following format for the host names you want to update:

   ```
   host_name|installation_location|user_ID|password
   ```

   If the Data server cannot find any Storage Resource agent with a host name and location that you specify in the file, that entry is ignored and the error is logged.

   Keep in mind the following considerations when including a password in the update file:

- For non-encrypted passwords: If you use "|" in the line for an agent, you must provide an "escape" character before it. The escape character is "\". For example, if the password you want to include in an update file is "native|agent", use `native\|agent` in the update file.

  If you use "\" in the line for an agent, you must provide an additional "\" as an escape character. For example, if the password is "native\agent", use `native\\agent` in the update file.
- For encrypted passwords:
  - To encrypt a password, start IBM Tivoli Storage Productivity Center's command-line tool (tpctool) from the \cli directory where Tivoli Storage Productivity Center is installed and run the following command: `tpctool encrypt [password]`, where `[password]` represents the password you want to encrypt. This command will generate a text string that represents the encrypted version of the password. For example, output might appear like the following: `iDroqToC07Oubh5i4mxMHQ==`. Copy this encrypted password to the correct location in the update file.
  - To include an encrypted password in an update file, you must add the following after the encrypted text: `@ENC@`. For example, an encrypted password in an update file might appear like this: `iDroqToC07Oubh5i4mxMHQ==@ENC@`

  Note:
  - Do not include `@ENC@` at the end of non-encrypted passwords.
  - See the *IBM Tivoli Storage Productivity Center Command-Line Interface Reference* for more information about **tpctool**.

Registering the Data agent with a different server

This topic provides information on how to register the Data agent with a different server.

Use the `PROBE_ME` file to get an existing Data agent to register with a different IBM Tivoli Storage Productivity Center Data server from the one that it currently uses. This file is also useful to reestablish the connection between the Data server and Data agent (for example, in the case of a Data server reinstallation). To change the registration of the Data agent with a different server, complete the following steps:

1. Change the name of the server in the `agent.config` file by performing the following steps:
   a. In the `<TPC_install>/ca/subagents/TPC/Data/config` directory, open the `agent.config` file.
   b. Change the `serverHost` property to the IP address of the new Tivoli Storage Productivity Center server.
2. Create a blank text file called `PROBE_ME` (all in uppercase letters with no `txt` extension) in the root of the agent's Data agent directory (`<TPC_install_dir>/ca/subagents/TPC/Data/`).
3. Restart the agent.

This will initiate a probe job to run and the agent to register with the server specified in the `agent.config` file. If you put any text in the `PROBE_ME` file, a default scan will also be run.

Deployment considerations for Storage Resource agents

There are a number of guidelines you must take into consideration when deploying and working with Storage Resource agents in your environment.

Required authority for deploying Storage Resource agents

You must be logged in to IBM Tivoli Storage Productivity Center with a user ID that has the superuser role to schedule Storage Resource agent deployments. See the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for more information about user roles.

Supported operating systems

Storage Resource agents are not deployable on all the operating systems where you can deploy Data agents. See "Using Data agents or Storage Resource agents" on page 75 to view the operating systems that are supported by Storage Resource agents.

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 probe job 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 or 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 an on-demand service (non-daemon service) runs as a stand-alone executable file on the monitored computer. Connectivity between the server and agent uses the same protocol that was used during the deployment of the agent.

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 running 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 install a Storage Resource agent on a computer where a Data agent is already installed *and* pointing to the same Data server as that Storage Resource agent. For example, if you install a Storage Resource agent on a computer, and then later install a Data agent that points to the same Data server as that Storage Resource agent, the Storage Resource agent is automatically uninstalled after the Data agent runs a successful probe job.

You can install a Storage Resource agent and a Data agent on the same computer if those agents communicate with different Data servers.

Storage Resource agents on computers running Windows 2008
Before you can deploy a Storage Resource agent on a computer that is running the Windows 2008 operating system, 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.

Data collection and Storage Resource agents
You can only run probe jobs against Storage Resource agents. You cannot run other data collection jobs against them such as scans, pings, and performance monitors. See "Reports for Data agents and Storage Resource agents" on page 76 for a list of reports that show the data collected by Storage Resource agents.

Upgrades and job logs
If you deploy a Storage Resource agent and later decide to upgrade it to a Data agent, the job logs for the probes that are run by the Storage Resource agent are removed from the computer and you are unable to access those job logs through the user interface. The job logs for probes that are run by the Data agent are retained.

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

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

Additional information about Storage Resource agents
See the *Planning > General Planning > Planning for Storage Resource agents* section in the *IBM Tivoli Storage Productivity Center Installation and Configuration Guide* for additional planning, configuration, and certificate information about Storage Resource agents.

In-band Fabric agents

In-band Fabric agents are agents that are used by Fabric Manager for in-band SAN operations.

For information about the supported agent types for switch performance management and fabric zone configuration, see Appendix A, "Supported Fabric Agent Types," on page 771.

Displaying information about an in-band Fabric agent

You can view information about an in-band Fabric agent including the status, IP address, host name, operating system, and Host Bus Adapter data.

To display information about an in-band Fabric agent, follow this procedure:

1. In the Navigation tree pane, expand **Administrative Services → Data Sources**. Left-click **Inband Fabric**.
2. In the right pane, click on the icon to the left of the Fabric agent for which you want information for. In the right pane, a notebook window opens. On the General page, the following information is displayed:

   Status The status of the agent.

   IP Address
   The IP address of the agent.

   Host name
   The host name of the agent.

   Operating System and Version
   The operating system and version of the agent (for example, Windows 9.1: Service Pack 2).

   HBA Data
   The Host Bus Adapter data for the agent.

Checking the in-band agent connection

You can check if an in-band Fabric agent is running and whether the version of the agent is up to date. A check also enables an agent that is disabled.

To check an in-band Fabric agent connection, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services → Data Sources**. Left-click **Inband Fabric**.
2. In the right pane, select an agent and click **Check**.

Note: The Fabric agent status is taken from the database. Because the agent does not send indications when it is shut down, the status may not be current. The `setup.properties` file for the Fabric agent has the following default parameter set: `monitorInterval=10`. This indicates that the monitor will be refreshed every 10 minutes. You can change the monitor setting for a different refresh time.
The status of the agent is displayed in the Agent Status window. If updates are required, they are automatically installed.

Viewing an agent log

Use the content of the log file to troubleshoot any errors that might occur during startup, processing or shutdown. Service alerts, if any, are displayed in red.

To display the log of an in-band Fabric agent, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services** → **Data Sources**. Left-click **Inband Fabric**.
2. In the right pane, select an agent and select **View Log**. In the right pane, the agent log is displayed.

Deleting an in-band Fabric agent

You can delete an in-band Fabric agent from the navigation tree and the database. Data discovered by the agent is removed from the database repository. This action is available for inactive agents only. You must first stop the agent before deleting the information for the agent. The agent is not uninstalled from IBM Tivoli Storage Productivity Center.

To remove information for an in-band Fabric agent, follow this procedure:

1. In the navigation tree pane, expand **Administrative Services** → **Data Sources**. Left-click **Inband Fabric Agents**.
2. In the right pane, select the agent and click **Delete**.

Out-of-band Fabric agents

Out-of-band agents are used by Fabric Manager for out-of-band SAN operations.

For information about the supported agent types for switch performance management and fabric zone configuration, see Appendix A, "Supported Fabric Agent Types," on page 771.

Displaying information about an out-of-band agent

You can view information about an out-of-band agent including the IP address, status and SNMP community. If the agent is a Brocade switch, the admin users ID is also displayed if it has been previously entered.

To display information about an out-of-band agent, follow this procedure:

1. In the navigation tree pane, expand **Administrative Services > Data Sources**. Left-click **Out of Band Fabric Agents**.
2. In the right pane, click on the icon to the left of the agent for which you want information for. In the right pane, a notebook window opens. On the General Information page, the following information is displayed:

   Status The status of the agent.

   Host Name
   The host name of the system on which the agent is installed.

IP Address
The IP address of the system on which the agent is installed. This column displays IPv4 and IPv6 addresses as appropriate.

SNMP Community Read
The name of the SNMP community that the agent belongs to. The SNMP community name acts as a password that is shared by one or more SNMP hosts. The community name is used to authenticate messages being received by this SNMP host. This field is optional and might be blank if the SNMP community has not been set. This default for this field is set for public.

Note: the default SNMP community is public. If this is not the correct community name for your environment, the out-of-band agent might not be able to properly perform scans.

SNMP Community Write
The default for this field is set for private.

Brocade Advanced Properties
The admin user name and password and whether advanced Brocade discovery has been enabled. These properties are optional and might be blank if an admin user name and password has not been set. These properties apply only if the agent is a Brocade switch.

Updating out-of-band agent information

After an out-of-band agent has been added, you can update the SNMP community name and, if the switch is a Brocade switch, the user ID and password and enable advanced Brocade discovery.

Note: You should not add the Brocade user ID and password to all switches; for more information see "Zone configuration" on page 69.

To update information for an out-of-band agent, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services > Data Sources**. Left-click **Out of Band Fabric Agents**.
2. In the right pane, click on the icon to the left of the agent.
3. In the right pane, a notebook window opens. On the General page, update the following information:

   SNMP Community
   The SNMP community the agent belongs to. This field is optional.

   Brocade Advanced Properties
   The admin user name and password. Select the Enable Advanced Brocade discovery to enable advanced Brocade discovery. These properties are optional and apply only if the agent is a Brocade switch.

   The Brocade API is supported on IPv4 systems only. IPv6 systems are not supported.

Checking the out-of-band agent connection

You can check to determine if an out-of-band agent is running.

To check the agent connection, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services > Data Sources**. Left-click **Out of Band Fabric Agents**.

2. In the right pane, select an agent and click **Check**. The status of the agent is displayed in the Agent Status window. If updates are required, they are automatically installed.

Adding an out-of-band agent

This topic provides information about adding an out-of-band agent.

To add an out-of-band agent, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services > Data Sources**. Left-click **Out of Band Fabric Agents**.
2. In the right pane, click **Add**.
3. In the Add Out of Band Fabric Agent window, enter host name, SNMP community (optional) and, if the agent switch is a Brocade switch, enter a user ID and password (optional). Select the Enable Advanced Brocade Discovery check box to enable advanced Brocade discovery. Click **OK**.

   Note: The Brocade API can be used with IPv4 only. IPv6 is not supported.

Deleting an out-of-band agent

You can remove an out-of-band agent from the navigation tree and the database. Data discovered by the agent is not removed from the database repository.

To remove an out-of-band agent, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services > Data Sources**. Left-click **Out of Band Fabric**.
2. In the right pane, select an agent and click **Delete**.

Removing out-of-band Fabric agent data

If you have a McDATA Intrepid 10000 director and have out-of-band Fabric agents as well as in-band or SMI-S agents, and you have run an out-of-band fabric discovery job, you must delete the data collected from the out-of-band Fabric agent because that data is invalid. This topic provides information on how to delete the out-of-band agent data.

To remove out-of-band Fabric agent data, complete the following steps:

1. Open the IBM Tivoli Storage Productivity Center GUI.
2. Expand **Administrative Services → Data Sources**. Left-click **Out of Band Fabric**.
3. In the right pane, select the out-of-band agent and click **Delete**.
4. On the agent machine where you installed the in-band agent, stop the Common agent service.
5. On the server machine, open the Tivoli Storage Productivity Center GUI and display the topology viewer. Wait for the topology viewer to display a status of ″missing″ for the machine where the in-band agent is installed. This can take approximately 5 minutes or more.
6. On the agent machine where you installed the in-band agent, start the Common agent service. This will start the data collection.

IBM Tivoli Storage Productivity Center Servers

Use the **Administrative Services > Data Sources > IBM Tivoli Storage Productivity Center Servers** function to manage the relationships between a *master* IBM Tivoli Storage Productivity Center server and its *subordinate* Tivoli Storage Productivity Center servers.

Creating master and subordinate relationships between Tivoli Storage Productivity Center servers enables you to use a single interface to generate reports based on data and metrics collected by multiple servers in a storage environment:

- A **master** server is a server that performs normal monitoring and reporting of storage resources like a standard server, but also gathers the storage information (using Tivoli Storage Productivity Center server probes) that has been collected by subordinate servers.
- A **subordinate** server is a server that monitors and reports on storage resources like a standard server, but also communicates with the master server during Tivoli Storage Productivity Center server probes. During these probes, the master server collects the storage information gathered by a subordinate server's agents and stores that information in its own database repository.

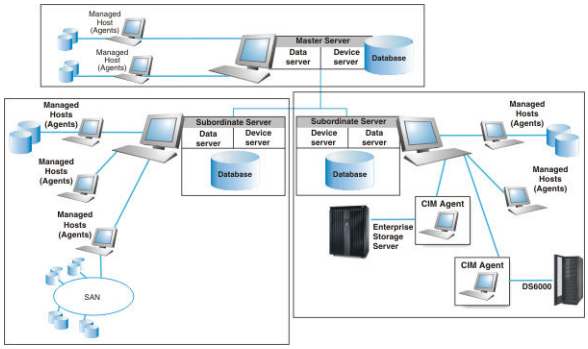

Figure 10. Master and subordinate server architecture

The rollup reports that reflect the storage information collected by the master server from subordinate servers are available in the **Tivoli Storage Productivity Center** > **Rollup Reports** node of the master server's navigation tree.

If the master server is located on an IPv6–only server, it can communicate with existing subordinate servers under the following conditions:

- The subordinate servers are upgraded to Tivoli Storage Productivity Center V4.1 or higher and
- The IPv6 protocol is enabled on the machines where they are located.

Before you can configure and manage subordinate servers, keep in mind the following:

- The master server must be up and running.
- You must be logged in to the user interface as an Tivoli Storage Productivity Center administrator or superuser
- We recommend that the master server should monitor no more than 500 unique data sources. This number includes subordinate servers, Data agents, Fabric agents (Inband and Out of Band), CIMOM agents, and VM servers (VMWare).
- We recommend that each subordinate server monitors no more than 1200 unique data sources. This number includes Data agents, Fabric agents (Inband and Out of Band), 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.
- After upgrading to Tivoli Storage Productivity Center V4.1 or later from a previous version of the application, 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**.

Adding an IBM Tivoli Storage Productivity Center Server

Learn how to add an Tivoli Storage Productivity Center server as a subordinate sever.

Keep in mind the following when adding a Tivoli Storage Productivity Center server as a subordinate server:

- How subordinate servers are configured to monitor storage entities within your environment can determine if fragmentation occurs within rollup reports. See Rollup reports - Fragmentation for more information.
- Tivoli Storage Productivity Center servers must be at version 3.3 or greater to be added as a subordinate server.
- We recommend that the master server should monitor no more than 500 unique data sources. This number includes subordinate servers, Data agents, Fabric agents (Inband and Out of Band), CIMOM agents, and VM servers (VMWare).
- We recommend that each subordinate server monitors no more than 1200 unique data sources. This number includes Data agents, Fabric agents (Inband and Out of Band), 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.
- If the master server is located on an IPv6–only server, it can communicate with existing subordinate servers under the following conditions:
  - The subordinate servers are upgraded to Tivoli Storage Productivity Center V4.1 or higher and
  - The IPv6 protocol is enabled on the machines where they are located.

To add a server as a subordinate server, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services → Data Sources → Tivoli Storage Productivity Center Servers**.
2. In the right pane, click **Add Tivoli Storage Productivity Center Server**.
3. In the Add Tivoli Storage Productivity Center Server window, enter the following information:

   Host Name
   Enter the host name of the machine on which the subordinate server is installed.

   Host Device server Port
   Enter the port with which the subordinate server's Device server is listening for requests. The default port for Tivoli Storage Productivity Center servers is 9550.

   Host Authentication Password
   Enter the password for the subordinate Tivoli Storage Productivity Center server. This host authentication password is entered when installing Tivoli Storage Productivity Center. Contact your system administrator or the person who installed Tivoli Storage Productivity Center to determine this password.

   Display Name (optional)
   Enter a name for the subordinate server.

   Description (optional)
   Enter a description for the subordinate server.

Test Tivoli Storage Productivity Center Server Connectivity before adding
Select this check box to have the master server connect and log into the subordinate server using the information provided on this window. This test is performed when you click **Save**.

4. Click **Save** to add the server as a subordinate server.
   - If you selected **Test Tivoli Storage Productivity Center Server Connectivity before adding**, IBM Tivoli Storage Productivity Center attempts to connect to the subordinate server using the information provided on this window.
   - If you did not select **Test Tivoli Storage Productivity Center Server Connectivity before adding**, the subordinate server is automatically added as a subordinate server but has a connection status of **Unknown** on the **Administrative Services → Data Sources → Tivoli Storage Productivity Center Servers** window.

   Note: If the login information for a subordinate server is incorrect, an error occurs and an associated error message dialog box is displayed.

Viewing subordinate server details

Learn how to view information about the subordinate servers that have been associated with a master server.

After you add a subordinate server to a master server, basic information about that server appears on the **Administrative Services → Data Sources → TPC Servers** window. To view additional information about a subordinate server on the TPC Servers window, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services → Data Sources → TPC Servers**.
2. In the right pane, double-click a server row or click the icon next to a server for more detailed information about that server. Information about a subordinate server will appear on the **Administrative Services → Data Sources → TPC Servers → TPC Server Details** window.

Testing the connection to subordinate servers

Learn how to test the connection from a master server to its associated subordinate servers.

After you add a subordinate server to a master server, basic information about that server appears on the **Administrative Services → Data Sources → TPC Servers** window. To verify that the master server can use the information provided to connect to subordinate server, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services → Data Sources → TPC Servers**.
2. In the right pane, highlight the row for a subordinate server and click **Test TPC Server Connection** to test its connection to the master server. A message window will appear with the results of the tested connection. The information that appears in the **Connection Status** field will be updated to reflect its status.

Removing subordinate servers

Learn how to remove a subordinate server from its association with a master server.

Use the **Administrative Services → Data Sources → TPC Servers** window to remove a subordinate server from a master server. Removing a subordinate server will:

- Remove its information from master server's database repository

- Remove its information from all rollup reports
- Remove it from the **TPC Servers** list on the **Administrative Services → Data Sources → TPC Servers** window

To remove a subordinate server, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services → Data Sources → TPC Servers**.
2. In the right pane, highlight the row of a subordinate server and click **Remove TPC Server**. A confirmation window appears.
3. Click **Yes** to confirm the removal of the subordinate server.

Modifying subordinate server information

Learn how to modify information about the subordinate servers that have been associated with a master server.

After you add a subordinate server to a master server, basic information about that server appears on the **Administrative Services → Data Sources → TPC Servers** window. You can edit the information for a subordinate server by modifying its Display Name, Description, and Host Authentication Password. To view modify information about a subordinate server on the TPC Servers window, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services → Data Sources → TPC Servers**.
2. In the right pane, double-click a server row or click the icon next to a server to edit information about that server. Information about a subordinate server will appear on **Administrative Services → Data Sources → TPC Servers → TPC Server Details**.
3. Edit the information in the following fields:

   Host Authentication Password
   Enter the password for the subordinate IBM Tivoli Storage Productivity Center server.

   Display Name (optional)
   Enter a name for the subordinate server.

   Description (optional)
   Enter a description for the subordinate server.

   Test TPC Server Connectivity before adding
   Select this check box to have the master server connect and log into the subordinate server using the information provided on this window. This test is performed when you click **Save**.
4. Select **File → Save** . If successful, the TPC Servers window appears and any new values for **Display Name** and **Description** are shown for a subordinate server.

Rollup reports - fragmentation

Fragmentation occurs when related storage entities are being monitored by different subordinate IBM Tivoli Storage Productivity Center servers.

The following example shows when fragmentation can occur using subordinate servers and rollup reports:

- **Example:** Host A is assigned a volume from storage subsystem 1. Host A is being monitored only by one Tivoli Storage Productivity Center subordinate

server (Server 1) and storage subsystem 1 is being monitored by another subordinate server (Server 2). Server 1 and Server 2 are subordinate servers to the master server (Server 3).

Reason for fragmentation: In this example, the host and subsystem are fragmented because they are not being monitored by the same Tivoli Storage Productivity Center server. This will cause the **Correlated Volume Space** value in the rollup reports for storage subsystems and volumes to be incorrect.

To prevent fragmentation when using master and subordinate server relationships to generate rollup reports, we recommend the following:

- Do not monitor a storage subsystem and the hosts to which its volumes are assigned with different subordinate servers that are associated with the same master server.
- Do not monitor a fabric and its related switches with different subordinate servers that are associated with the same master server.
- If an agent machine is using storage from a subsystem, then both the agent and subsystem should be monitored by the same subordinate server.
- Configure all the subordinate servers to discover all the storage subsystems within your environment. Note that if a subordinate server is monitoring a subsystem with an agent other than the agent that is using storage on that subsystem, the value for the **Correlated Volume Space** column in rollup reports will be incorrect. For example:
  - Host A has storage from Subsystem 1
  - Subordinate server B is Monitoring Host A and Subsystem 1
  - Subordinate server C is Monitoring Subsystem 1 and probed the subsystem more recently than server B.
  - The master server is monitoring subordinate servers B and server C.

  The master server will report the storage subsystem data from server C. Because the master server does not know about Host A, the value for the **Correlated Volume Space** column will be incorrect.

VMware data source

VMware Virtual Infrastructure data sources are used by IBM Tivoli Storage Productivity Center to collect information from the hypervisors and virtual machines within them.

When you add a VMware data source in IBM Tivoli Storage Productivity Center, you need a user ID that has permission to browse through the data stores on VMware. IBM Tivoli Storage Productivity Center needs to browse through the data stores to collect information from the ESX Servers. However the ″Read Only″ role as defined by VMware does not allow IBM Tivoli Storage Productivity Center to browse the data stores. You can use the ″Virtual Machine Power User″ role as a choice if you do not want to use the Administrator role, or you can create a custom role with the required permissions.

To check what user role you have for VMware, follow these steps:

1. Connect a VI client to the ESX Server or VirtualCenter (whichever is used as the data source). Click on the ESX Server and go to the **Summary** tab. Select a data source, right-click on the data source and select **Browse**, to confirm that the browse permission is working correctly.
2. Connect the VI client to the data source with the admin role, go to the **Permission** tab and determine if the role name is used for the user.

3. Open the **Permissions** view, select that role, then right-click and select the **Edit** role. See what permissions are checked. Expand the check groups to find the specific ones.

For more information about VMware user roles, see http://www.vmware.com/pdf/vi3_vc_roles.pdf.

Adding a VMware VI Data Source

This topic provides information about adding a VMware Virtual Infrastructure data source.

To add a VMware VI data source, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services → Data Sources → VMware VI Data Source**.
2. In the right pane, click **Add VMware VI Data Source**.
3. In the Add VMware VI Data Source window, enter the following information:

   Host Name
   The host name for the VMware data source (ESX Server or VirtualCenter).

   Protocol
   This is http or https. If you use https, you must have first imported the SSL certificate into IBM Tivoli Storage Productivity Center. For information about how to import the SSL certificate, see the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide Version 4 Release 1*. The protocol must match the same protocol that is set for the ESX Server or VirtualCenter.

   Port The port used by the VMware data source.

   Username
   The user name for the VMware data source.

   Password
   The password for the VMware data source.

   Password Confirm
   The password is confirmed for the VMware data source.

   Display Name
   The name displayed in reports for the VMware data source.

   Description
   The description for the VMware data source.

   Test VMware VI Data Source connectivity before adding
   Check this box to check the connectivity for the VMware data source.
4. Click **Save**.

Updating VMware data source information

After a VMware data source has been added, you can update the administrator name or password for the VMware data source.

To update information for a VMware data source, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services → Data Sources → VMware VI Data Source**.

2. In the right pane, click on the icon to the left of the VMware data source you wish to change.
3. In the right pane, a notebook window opens. Make the appropriate changes.

Displaying information about a VMware data source

You can view information about a VMware data source including its host name, protocol, port, and type.

To display information about an out-of-band agent, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services → Data Sources → VMware VI Data Source**.
2. In the right pane, click on the icon to the left of the VMware data source for which you want information for. In the right pane, a notebook window opens. The following information is displayed:

   Host Name
   The IP address of the data source.

   Protocol
   This is http or https.

   Port The port used by the VMware data source.

   Type Virtual Center.

   Administrator Name
   The name of the administrator for the VMware data source.

   Password
   The password for the VMware data source.

   Password Confirm
   Confirm the password for the VMware data source.

   Display Name
   The name displayed for the VMware data source.

   Description
   Description for the VMware data source.

   Software Level
   The software level for the VMware data source.

   ConnectionStatus
   The connection status to the VMware data source.

   Status Timestamp

   Test VMware VI Data Source connectivity before updating
   Check this box to check the connectivity to the VMware data source.

Show managed devices for VMware

You can view information about the managed devices for VMware.

To show the managed devices for VMware, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services → Data Sources → VMware VI Data Source**.
2. In the right pane, click **Show Managed Devices**.
3. A pop-up window displays the VMware managed devices. Click **OK** to close the window.

Deleting a VMware VI Data Source

This topic provides information about deleting a VMware Virtual Infrastructure data source.

To delete a VMware VI data source, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services** → **Data Sources** → **VMware VI Data Source**.
2. In the right pane, click **Remove VMware VI Data Source**.
3. In the confirmation dialog window to remove a VMware data source, click **OK**.

Administering agents on the Virtual I/O Server

Use the **startsvc** and **stopsvc** commands to start and stop the agents on the Virtual I/O Server. Use the **lssvc** command to list the agents installed on the Virtual I/O Server.

Listing agents on Virtual I/O Server

Use the **lssvc** command to list the IBM Tivoli Storage Productivity Center agents on a Virtual I/O Server.

You must log into the Virtual I/O Server as the **padmin** user ID to use this command.

The syntax for the **lssvc** command is:

The **lssvc** command lists all the available agents that are managed by the Virtual I/O Server. If an agent name is passed to the **lssvc** command, then a listing of the attributes with their configured values are displayed.

Examples

To list all the agents installed, enter the following command:

```
lssvc TPC
```

Starting IBM Tivoli Storage Productivity Center agents

Use the **startsvc** command to start Tivoli Storage Productivity Center agents on a Virtual I/O Server.

You must log into a Virtual I/O Server as the **padmin** user ID to use the **startsvc** command.

The parameters for the command are:

The **startsvc** command starts the specified agent. The agent name is case sensitive. For repeated execution, you typically receive a message that the agent has already been started.

The *agent_name* can be:

TPC_data
Starts the Data agent.

TPC_fabric
Starts the Fabric agent.

Note: The agent name is case sensitive.

Examples

To start the Data agent, enter the following command:

```
startsvc TPC_data
```

To start the Fabric agent, enter the following command:

```
startsvc TPC_fabric
```

Stopping IBM Tivoli Storage Productivity Center agents

Use the **stopsvc** command to stop Tivoli Storage Productivity Center agents on the Virtual I/O Server.

You must log into a Virtual I/O Server as the **padmin** user ID to use this command.

The parameters for the **stopsvc** command are:

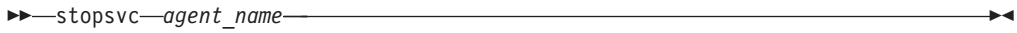

The *agent_name* can be:

TPC_data
Stops the Data agent.

TPC_fabric
Stops the Fabric agent.

Note: The agent name is case sensitive.

Examples

To stop the Data agent, enter the following command:

```
stopsvc TPC_data
```

To stop the Fabric agent, enter the following command:

```
stopsvc TPC_fabric
```

Administering Tivoli Integrated Portal

This topic provides information on how to perform administrative tasks in IBM Tivoli Integrated Portal that are reflected in IBM Tivoli Storage Productivity Center.

Changing the user authentication method

The IBM Tivoli Storage Productivity Center installation program enables you to select a user authentication method that is used by Tivoli Storage Productivity Center, IBM Tivoli Storage Productivity Center for Replication, and IBM Tivoli Integrated Portal. You can choose to authenticate users against the users defined

for the local operating system or to authenticate users against the users defined in a Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory repository. You can change the user authentication method after installation using Tivoli Integrated Portal.

You can select the following user authentication methods in Tivoli Integrated Portal:

- **Federated repositories.** This method authenticates the Tivoli Storage Productivity Center user against a Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory, directory service.
- **Local operating system.** This authentication method authenticates the Tivoli Storage Productivity Center user against the users defined for the local operating system. If you use operating system authentication, 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.

The following topics describe how to select and configure each user authentication method.

Changing the authentication method from local operating system to LDAP federated repositories

You can configure IBM Tivoli Storage Productivity Center to communicate with an external Lightweight Directory Access Protocol (LDAP) repository such as IBM Tivoli Directory Server or Microsoft Active Directory. To change the authentication method from local operating system to LDAP, you must use IBM Tivoli Integrated Portal.

To change the user authentication method from local operating system to LDAP, complete the following steps. Contact your LDAP server administrator for assistance.

1. 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.
2. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator permissions.
3. In the Tivoli Integrated Portal navigation tree, click **Security → Secure administration, applications, and infrastructure**.
4. On the Secure administration, applications, and infrastructure page, select **Federated Repositories** from the **Available Realm Definitions** list.
5. Click **Configure**. The Federated repositories page is displayed.
6. Under **Related Items**, click **Manage repositories**.
7. On the Manage repositories page, add the LDAP repository that you want to use for authentication as follows:
   a. Click **Add** to add a new repository.
   b. Enter the values for the following fields:

- **Repository identifier**. A unique identifier for the LDAP repository. This identifier uniquely identifies the repository within the cell, for example: LDAP1.
- **Directory type**. The type of LDAP server to which you want to connect.
- **Primary host name**. The host name of the primary LDAP server. This host name is either an IP address or a domain name service (DNS) name.
- **Port**. The LDAP server port. The default value is 389, which is not a Secure Sockets Layer (SSL) connection. You can use port 636 for a Secure Sockets Layer (SSL) connection. For some LDAP servers, you can specify a different port for a non-SSL or SSL connection. If you do not know the port to use, contact your LDAP server administrator.
- **Bind distinguished name**. The distinguished name (DN) for the application server to use when binding to the LDAP repository. If no name is specified, the application server binds anonymously. In most cases, bind DN and bind password are needed. However, when anonymous bind can satisfy all of the required functions, bind DN and bind password are not needed.
- **Bind password**. The password for the application server to use when binding to the LDAP repository.

c. Click **OK**.

d. In the **Messages** box on the Manage repositories page, click the **Save** link in **Save to the master configuration**.

8. On the Manage repositories page, click the identifier for the repository that you want to use in the **Repository identifier** column.
9. On the configuration page for the repository, configure the following items:
   a. Click **LDAP entity types** under **Additional Properties**. The LDAP entities page is displayed.
   b. In the Entity type column, click the link for **Group**, **OrgContainer**, and **PersonAccount** and complete the **Search bases** field. This field specifies the search bases that are used to search this entity type. The search bases specified must be subtrees of the base entry in the repository. The following are example search bases, where `o=ibm,c=us` is the base entry in the repository:

      ```
      o=ibm,c=us or cn=users,o=ibm,c=us or ou=austin,o=ibm,c=us
      ```

      Delimit multiple search bases with a semicolon (;). For example:

      ```
      ou=austin,o=ibm,c=us;ou=raleigh,o=ibm,c=us
      ```

   c. Click **OK** and then click the **Save** link in **Save to the master configuration** each time that you update the **Search bases** field.
10. Return to the Federated repositories page and click **Supported Entity Types** under **Additional Properties**.
11. On the Supported entity types page, configure the following items:
    a. In the Entity type column, click the link for **Group**, **OrgContainer**, and **PersonAccount** and complete the **Base entry for the default parent** and **Relative Distinguished Name properties** fields.
       - In the **Base entry for the default parent** field, enter the same value that you entered in the **Search bases** field in step 9.

- In the **Relative Distinguished Name properties** field, enter the appropriate LDAP attribute name. In most cases, the values for this field will be `cn` for **Group**, `o;ou;dc;cn` for **OrgContainer**, and `uid` for **PersonAccount**.

b. Click **OK** and then click the **Save** link in **Save to the master configuration** each time that you update the **Base entry for the default parent** field.

12. Return to the Federated repositories page and click **Apply** and then click the **Save** link in **Save to the master configuration**.
13. Under **Repositories in the realm**, click **Add base entry to Realm**.
14. On the Repository reference page, configure the following items:
    a. In the **Repository** list, select the repository that you created in step 7.
    b. In the **Distinguished name of a base entry that uniquely identifies this set of entries in the realm** field, enter the distinguished name of a base entry that uniquely identifies the repository in the realm. In most instances, this value will be the same value that you entered in the **Search bases** field in step 9.
    c. In the **Distinguished name of a base entry in this repository** field, enter the distinguished name of the base entry within the repository. In most instances, this value will be the same value that you entered in the **Distinguished name of a base entry that uniquely identifies this set of entries in the realm** field.
    d. Click **OK** and then click the **Save** link in **Save to the master configuration** each time that you update the **Distinguished name of a base entry that uniquely identifies this set of entries in the realm** field.
15. On the Federated repositories page, there are now two repositories that are displayed under **Repositories in the realm**: the repository that you have added and a default repository that shows **File** in the Repository type column. Configure the following items:
    a. Leave the value in the **Realm name** field as is or change the name of the realm name.
    b. In the **Primary administrative user name** field, enter the name of a user in the repository that you added. This user will be granted administrative privileges in the Tivoli Integrated Portal server, the Tivoli Storage Productivity Center Device server, and the IBM Tivoli Storage Productivity Center for Replication server.
    c. Click **Server identity that is stored in the repository**.
    d. In the **Server user ID or administrative user on a Version 6.0.x node** field, enter the same ID that you entered in the **Primary administrative user name** field and enter the password for the user ID in the **Password** field.
    e. Select the default file repository that shows **File** in the Repository type column and click **Remove**.
    f. Click **OK** and then click the **Save** link in **Save to the master configuration**.
16. Return to Secure administration, applications, and infrastructure page and configure the following items:
    a. In the **Available realm definitions** list, select Federated repositories and then click **Set as current**.
    b. Click **Apply** and then click the **Save** link in **Save to the master configuration**
17. Stop and restart the Tivoli Integrated Portal, IBM Tivoli Storage Productivity Center for Replication servers and the IBM Tivoli Storage Productivity Center

Data and Device servers. For information regarding stopping and starting a server, go to the Tivoli Storage Productivity Center Information Center and search for *starting and stopping services*.

To verify that the federated repository is configured correctly, complete the following steps:

1. Log on to Tivoli Integrated Portal as the new super user ID (the administrative account is no longer the OS credentials, but is now configured to use the LDAP server super user credentials).
2. In the Tivoli Integrated Portal navigation tree, click **Users and Groups** → **Manage Users** .
3. In the **Search by** list, select **User ID**.
4. Click **Search** to search for users in the federated repository. The list of users includes users from both the LDAP and the local file registry.

Related tasks

"Enabling secure communication between Tivoli Storage Productivity Center and the LDAP repository"
You can use the Secure Socket Layer (SSL) protocol to secure the communication between IBM Tivoli Storage Productivity Center and the LDAP repository that you are using for user authentication. The SSL protocol provides security and data integrity for communications over Transmission Control Protocol/Internet Protocol (TCP/IP) networks.

"Disabling secure communication between Tivoli Storage Productivity Center and the LDAP repository" on page 164
You can disable the Secure Socket Layer (SSL) protocol between the LDAP repository and the IBM Tivoli Storage Productivity Center system at any time.

Enabling secure communication between Tivoli Storage Productivity Center and the LDAP repository:

You can use the Secure Socket Layer (SSL) protocol to secure the communication between IBM Tivoli Storage Productivity Center and the LDAP repository that you are using for user authentication. The SSL protocol provides security and data integrity for communications over Transmission Control Protocol/Internet Protocol (TCP/IP) networks.

To configure for SSL, you must complete the following steps in IBM Tivoli Integrated Portal:

1. 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.
2. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator permissions.
3. In the Tivoli Integrated Portal navigation tree, click **Security** → **SSL certificate and key management** → **Key stores and certificates** → **NodeDefaultTrustStore** → **Signer certificates**.

4. On Signer certificates page, click **Retrieve from port**.
5. On the Retrieve from port page, enter values for the following fields:
   - **Host**. The fully qualified host and domain name of your LDAP-compliant repository.
   - **Port**. The port where your LDAP repository is listening for secure communications; this is usually port 636.
   - **SSL configuration for outbound connection**. Accept the default value.
   - **Alias**. An alias name for the retrieved certificate.
6. Click **Retrieve signer information**.
7. When the signer information is displayed, click **OK**.
8. In the **Messages** box on the Signer certificates page, click the **Save** link in **Save to the master configuration**.
9. On the Signer certificates page, select the LDAP repository certificate that you just retrieved and click **Extract**.
10. On the Extract signer certificate page, enter values for the following fields:
    - **File name**. The file name for this certificate. For example, LDAPSSLCert. This file is saved in the C:*Tivoli Integrated Portal install directory*\profiles\TIPProfile\etc\ directory on the Windows platform. In the AIX, Linux, and UNIX platforms, this file is saved in the */Tivoli Integrated Portal install directory*/profiles/TIPProfile/etc/ directory. *Tivoli Integrated Portal install directory* is the directory where Tivoli Integrated Portal is installed.
    - **Data type**. Accept the default value.
11. Click **OK**. A message is displayed on the Signer certificates page stating that the file was successfully extracted.
12. In the Tivoli Integrated Portal navigation tree, click **Security → Secure administration, applications, and infrastructure**.
13. On the Secure administration, applications, and infrastructure page, make sure that **Federated Repositories** is selected in the **Available Realm Definitions** list.
14. Click **Configure**. The Federated repositories page is displayed.
15. Under **Related Items**, click **Manage repositories**.
16. On the Manage repositories page, click the identifier for the repository for which you want to enable the SSL protocol.
17. On the configuration page for the repository, configure the following items:
    a. In the **Port** field, enter 636
    b. Select the **Require SSL communications** check box.
18. Click **OK**.
19. In the **Messages** box on the Manage repositories page, click the **Save** link in **Save to the master configuration**.
20. Log off of Tivoli Integrated Portal.
21. From the command line, stop the Tivoli Integrated Portal server, the Tivoli Storage Productivity Center for Replication, and the Tivoli Storage Productivity Center Device and Data servers. For information regarding starting and stopping a server, go to the Tivoli Storage Productivity Center Information Center and search for *starting and stopping services*.
22. Use the following commands to import the LDAP repository certificate into the Tivoli Storage Productivity Center Device server keystore:

    Windows

```
cd C:\Tivoli Storage Productivity Center install directory\jre\bin\

keytool -import -keystore C:\Tivoli Storage Productivity Center install directory
\device\apps\was\profiles\deviceServer\config\cells\DefaultNode\nodes\DefaultNode\
trust.p12 -storetype pkcs12 -storepass WebAS -file
C:\Tivoli Integrated Portal install directory\profiles\TIPProfile\etc\
LDAP repository certificate file name
```

where *Tivoli Storage Productivity Center install directory* is the is the directory where Tivoli Storage Productivity Center is installed, *Tivoli Integrated Portal install directory* is the directory where Tivoli Integrated Portal is installed, and *LDAP repository certificate file name* is the file name that you assigned to the LDAP repository certificate in Step 11.

Type `yes` when you are asked if you want to trust this certificate.

AIX, Linux, and UNIX

```
cd Tivoli Storage Productivity Center install directory/jre/bin/

./keytool -import -keystore /Tivoli Storage Productivity Center
install directory/device/apps/was/profiles/deviceServer/config/
cells/DefaultNode/nodes/DefaultNode/trust.p12 -storetype pkcs12
-storepass WebAS -file /Tivoli Storage Productivity Center
install directory/profiles/TIPProfile/etc/
LDAP repository certificate file name
```

where *Tivoli Storage Productivity Center install directory* is the is the directory where Tivoli Storage Productivity Center is installed, *Tivoli Integrated Portal install directory* is the directory where Tivoli Integrated Portal is installed, and *LDAP repository certificate file name* is the file name that you assigned to the LDAP repository certificate in Step 11.

Type `yes` when you are asked if you want to trust this certificate.

23. Use the following commands to import the LDAP repository certificate into the Tivoli Storage Productivity Center for Replication server keystore:

    Windows

    ```
    keytool -import -keystore C:\Tivoli Storage Productivity Center install
    directory\eWAS\profiles\CSM\config\cells\DefaultNode\nodes\DefaultNode
    \trust.p12 -storetype pkcs12 -storepass WebAS -file <C:\Tivoli
    Storage Productivity Center install directory\profiles\TIPProfile\
    etc\LDAP repository certificate file name
    ```

    where *Tivoli Storage Productivity Center for Replication* install directory is the directory where Tivoli Storage Productivity Center for Replication is installed.

    Type `yes` when you are asked if you want to trust this certificate.

    AIX, Linux, and UNIX

    ```
    ./keytool -import -keystore /Tivoli Storage Productivity Center install
    directory/eWAS/profiles/CSM/config/cells/DefaultNode/nodes/DefaultNode
    /trust.p12 -storetype pkcs12 -storepass WebAS -file /Tivoli Storage
    Productivity Center install directory/profiles/TIPProfile/etc/
    LDAP repository certificate file name
    ```

    where *Tivoli Storage Productivity Center for Replication* install directory is the directory where Tivoli Storage Productivity Center for Replication is installed.

    Type `yes` when you are asked if you want to trust this certificate.

24. From the command line, start the Tivoli Integrated Portal server, the Tivoli Storage Productivity Center for Replication, and the Tivoli Storage Productivity Center Device and Data servers.

Important: All of the group-to-role mappings in Tivoli Storage Productivity Center for Replication are deleted as a result of securing communications between LDAP and the Tivoli Storage Productivity Center system. You must manually re-create the Tivoli Storage Productivity Center for Replication group-to-role mappings

following this procedure. The Tivoli Storage Productivity Center and Tivoli Integrated Portal group-to-role mappings are preserved during this procedure and require no adjustment.

Disabling secure communication between Tivoli Storage Productivity Center and the LDAP repository:

You can disable the Secure Socket Layer (SSL) protocol between the LDAP repository and the IBM Tivoli Storage Productivity Center system at any time.

To disable the use of SSL, you must complete the following steps in IBM Tivoli Integrated Portal:

1. 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.
2. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator permissions.
3. In the Tivoli Integrated Portal navigation tree, click **Security** → **Secure administration, applications, and infrastructure**.
4. On the Secure administration, applications, and infrastructure page, make sure that **Federated Repositories** is selected in the **Available Realm Definitions** list.
5. Click **Configure**. The Federated repositories page is displayed.
6. Under **Related Items**, click **Manage repositories**.
7. On the Manage repositories page, click the identifier for the repository for which you want to disable the SSL protocol.
8. On the configuration page for the repository, configure the following items:
   a. In the **Port** field, enter 389.
   b. Clear the **Require SSL communications** check box.
9. Click **OK**.
10. In the **Messages** box on the Manage repositories page, click the **Save** link in **Save to the master configuration**.
11. Log off of Tivoli Integrated Portal.
12. From the command line, stop and then start the Tivoli Integrated Portal server, the Tivoli Storage Productivity Center for Replication, and the Tivoli Storage Productivity Center Device and Data servers. For information regarding starting and stopping a server, go to the Tivoli Storage Productivity Center Information Center and search for *starting and stopping services*.

Important: All of the group-to-role mappings in Tivoli Storage Productivity Center for Replication are deleted as a result of disabling the secure communications between LDAP and the Tivoli Storage Productivity Center system. You must manually re-create the Tivoli Storage Productivity Center for Replication group-to-role mappings following this procedure. The Tivoli Storage Productivity

Center and Tivoli Integrated Portal group-to-role mappings are preserved during this procedure and require no adjustment.

Changing the authentication method from LDAP to local operating system

You can configure IBM Tivoli Storage Productivity Center to authenticate against the local operating system. To change the authentication method from LDAP to local operating system, you must use Tivoli Integrated Portal.

If you use operating system authentication, the use of the single sign-on feature is limited. Operating system authentication does not support single sign-on if Tivoli Storage Productivity Center and Tivoli Integrated Portal are on separate computers. Operating system authentication also does not support single sign-on for element managers regardless of the location of the element manager application.

To change the user authentication method from LDAP to local operating system, complete the following steps:

1. 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.
2. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator permissions.
3. In the Tivoli Integrated Portal navigation tree, click **Security → Secure administration, applications, and infrastructure**.
4. On the Secure administration, applications, and infrastructure page, select **Local operating system** from the **Available Realm Definitions** list.
5. On the Local operating system page, configure the following items:
   a. In the **Primary administrative user name** field, enter the name of a user that is defined in your local operating system. This user will be granted administrative privileges in the Tivoli Integrated Portal server, the Tivoli Storage Productivity Center Device server, and the IBM Tivoli Storage Productivity Center for Replication server.
   b. Click **Server identity that is stored in the repository**.
   c. In the **Server user ID or administrative user on a Version 6.0.x node** field, enter the same ID that you entered in the **Primary administrative user name** field and enter the password for the user ID in the **Password** field.
   d. Click **OK**.
   e. In the **Messages** box on the Secure administration, applications, and infrastructure page, click the **Save** link in **Save to the master configuration**.
6. Configure the following items on the Secure administration, applications, and infrastructure page:
   a. In the **Available realm definitions** list, select **Local operating system** and then click **Set as current**.
   b. Click **Apply** and then click the **Save** link in **Save to the master configuration**

7. Stop and restart the Tivoli Integrated Portal, IBM Tivoli Storage Productivity Center for Replication servers and the IBM Tivoli Storage Productivity Center Data and Device servers. For information regarding stopping and starting a server, go to the Tivoli Storage Productivity Center Information Center and search for *starting and stopping services*.

Changing the LTPA token expiration for single-sign on

A single sign-on environment requires a centralized authentication repository that is accessed by all applications within the environment. The user's authentication information is passed between applications using Lightweight Third-Party Authentication (LTPA) tokens. You can change the expiration time for the LTPA tokens using IBM Tivoli Integrated Portal.

To change the expiration time for the LTPA tokens, complete the following steps.

1. 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 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.
2. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator permissions.
   - If you are using Tivoli Integrated Portal that was installed with Tivoli Storage Productivity Center, type your Tivoli Storage Productivity Center ID and password in the **User ID** and **Password** fields and click **Log in**. Your Tivoli Storage Productivity Center user ID must be a member of the Tivoli Storage Productivity Center Superuser role.
   - If you are using a stand-alone Tivoli Integrated Portal system, type your Tivoli Integrated Portal user ID and password. Your Tivoli Integrated Portal user ID must be a member of the Administrator role.
3. In the Tivoli Integrated Portal navigation tree, click **Security → Secure administration, applications, and infrastructure**.
4. Under **Authentication**, click **Authentication mechanisms and expiration**.
5. On the **Configuration** tab under **Authentication expiration**, type the expiration time in minutes in the **Timeout value for forwarded credentials between servers** field. The expiration time must be greater than the authentication cache timeout value that is shown above the **Timeout value for forwarded credentials between servers** field. The default expiration time is 1440 minutes (24 hours).
6. Click **Apply**.
7. In the Messages box, click **Save directly to the master configuration**.

Chapter 4. Managing storage resources

The information in this chapter describes how to use IBM Tivoli Storage Productivity Center to manage your environment.

Managing DS8000 storage systems using element managers

Use the Configuration Utility in IBM Tivoli Storage Productivity Center to add and manage element managers for IBM System Storage DS8000. *Element managers* are applications that are used to configure and manage a storage device. You can add and manage connections to the element manager that is provided with DS8000, DS8000 Storage Manager, or you can specify other element managers for a device.

The **Element Manager** tab in the Configuration Utility is used to add and manage DS8000 element managers. The **Element Manager** tab contains multiple sections. The **DS8000 Element Manager** section shows the DS8000 element managers that have been added to Tivoli Storage Productivity Center.

To expand or collapse the **DS8000 Element Manager** section, click the + or - icon to the left of the section heading.

The DS8000 Element Manager section

The **DS8000 Element Manager** section contains a table listing the element managers that are visible to Tivoli Storage Productivity Center and a **Select Action** list. This list enables you to add an element manager to the section and to manage existing element managers. The menu items that are shown in the **Select Action** list depend on whether you have selected an element manager in the table. This table contains the following fields:

Element Manager
Shows the IP address of the element manager or the display name if a display name was defined.

Connection Status
Shows the status of the element manager. The values that appear in this field are:

- **Online** (The element manager is accessible.)
- **Offline** (The element manager is not accessible)

Device
The name for the device as defined in Tivoli Storage Productivity Center.

UID/PWD
Shows whether a user ID and password have been defined for authentication to the element manager. The values that appear in this field are:

- **Yes** (A user ID and password have been defined.)
- **No** (A user ID and password have not been defined.)

CIMOM
Shows whether a CIMOM is configured for the element manager. The values that appear in this field are:

- **Yes** (A CIMOM is configured.)

- **No** (A CIMOM is not configured.)

LIC Enabled

Shows whether a launch descriptor file (LDF) has been registered for the element manager. An LDF is an XML file that describes the menus and the detailed uniform resource locators (URLs) for applications. An LDF is required to enable the launch in context feature for the device. Tivoli Storage Productivity Center provides an LDF for the element manager that is provided with DS8000.

The values that appear in this field are:

- **Yes** (Launch in context is enabled.)
- **No** (Launch in context is not enabled.)

Configuring the DS8000 element manager

DS8000 Storage Manager is the element manager that is provided with IBM System Storage DS8000. This section describes how to add and manage DS8000 Storage Manager in IBM Tivoli Storage Productivity Center.

You can perform the following tasks for DS8000 element managers:

- Add DS8000 element managers that you want to manage.
- Start a DS8000 element manager (for example, DS8000 Storage Manager) to configure and administer the storage subsystems associated with the element manager. On Windows, the DS8000 Storage Manager is started in a browser window that is embedded in Tivoli Storage Productivity Center. On AIX, Linux, and UNIX, DS8000 Storage Manager is started in an external Web browser window.
- Modify or remove element managers.
- Test the connection between Tivoli Storage Productivity Center and element managers.
- Add CIM object manager (DS CIM Agent) connections for element managers. For example, you can perform this task to associate a DS CIM Agent that is embedded on a DS8000 Hardware Management Console (HMC) with a DS8000 element manager. You are not required to connect to a DS CIM Agent if you want to configure and administer DS8000 storage systems. However, a DS CIM Agent connection is required if you want to use Tivoli Storage Productivity Center to monitor and collect storage metrics about the DS8000 storage system and to discover other storage entities managed by the DS CIM Agent .
- Modify or remove DS CIM Agent connections for element managers.

DS8000 credentials

The configuration of DS8000 credentials depends on whether you are using the single sign-on feature in your environment. *Single sign-on* is an authentication process that enables you to start an application from another application using a single user ID and password. For more information about single sign-on, refer to “Single sign-on” on page 54.

To use single sign-on with DS8000, the DS8000 system must be release 4.2 or later and must be configured for Lightweight Directory Access Protocol (LDAP). Tivoli Storage Productivity Center passes the DS8000 user credentials as a Lightweight Third Party (LTPA) token.

If you are not using single sign-on, the DS8000 element manager user name and password are stored for each Tivoli Storage Productivity Center user. A set of user credentials must be stored for each Tivoli Storage Productivity Center user who wants to access an DS8000 element manager. The **Element Manager** tab lists all element managers known to Tivoli Storage Productivity Center, regardless of the presence of valid user credentials. Element manager credentials might be out of date or absent for that Tivoli Storage Productivity Center user.

If the element manager has been added to Tivoli Storage Productivity Center using different credentials from the ones you are using, or if you attempt to open the DS8000 element manager for the first time after the element manager was added to Tivoli Storage Productivity Center by DS CIM Agent discovery, the following message is shown:

```
The element manager GUI cannot be accessed because the user name and
password are missing. Update the required fields in the next dialog. The
GUI will be displayed after the fields have been correctly completed and
the dialog closed.
```

If you created an element manager with one set of credentials and are currently logged on to Tivoli Storage Productivity Center as another user with different credentials, you can open the DS8000 element manager in one of two ways:

- **Add credentials**. Close the warning message. The **Modify Element Manager** window opens. Add the current credentials and click **Save**. You can then open the DS8000 element manager for that element manager.

  Important: Use this method with caution because anyone who knows the Tivoli Storage Productivity Center credentials can manage the DS8000 device, including performing destructive operations.
- **Log on as the user who added the element manager**. Close the warning message and the **Modify Element Manager** window. Log off of Tivoli Storage Productivity Center. Log on with the user ID that was used to create the element manager. From the Configuration Utility **Element Manager** tab, start the DS8000 element manager.

If the element manager has been added by DS CIM Agent discovery, user credential information has not been automatically added. The first time you attempt to open the DS8000 element manager, you must provide user credentials. Close the error message and enter the credential information in the **Modify Element Manager** dialog box. Click **Save**. You can then open the DS8000 element manager.

Adding DS8000 element managers

This section discusses how to add IBM System Storage DS8000 element managers.

There are two ways to add DS8000 element managers:

Manual method
: Use the **Element Manager** tab in the Configuration Utility to add element managers individually.

Automatic method
: Use the CIM object manager (CIMOM) discovery process in IBM Tivoli Storage Productivity Center to add any element managers that are managing the same storage subsystems as DS CIM Agents in your environment.

Using Add Element Manager:

Use the **Add Element Manager** action in the Element Management perspective to add element managers to IBM Tivoli Storage Productivity Center.

Use this method when you want to add one element manager at a time or do not require Tivoli Storage Productivity Center to monitor the storage subsystem associated with the element manager.

To add a DS8000 element manager, follow these steps:

1. Start the Tivoli Storage Productivity Center user interface.
2. Log in to Tivoli Storage Productivity Center.
3. Use either of the following methods to add an element manager:
   - Click the **Element Management** button in the toolbar to access the Element Management perspective. Select **Add Element Manager** from the **Select action** menu.
   - Expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**. On the Element Manager tab in the right pane, click the **Add Element Manager** button.

   This displays the Configure Element Manager page.

Figure 11. Configure Element Manager page

Enter the following information:

Note: Required fields are marked with an asterisk (*).

***Username**
Enter the user name that IBM Tivoli Storage Productivity Center will use to log on to an element manager and access the interface for administering the associated storage subsystems.

***Password**
Enter the password associated with the ***Username** field.

Tip: If the DS8000 element manager that you want to add is a new R3 or R4 installation and the administrator user password has not been changed from the initial value of **admin**, you are prompted to change the password. Be prepared to enter a password and record the new password in a safe place.

***Password Confirm**
Enter the same password you entered in the ***Password** field.

Display Name
Enter the name of the element manager as you want it to appear in the **Name** column and **Select a view** menu on the Element Management page. The IP address of the system on which the element manager is located is used if you do not enter a value in this field. If this field is left blank, the display name is set to a URL with an IP address of the element manager, for example ″https://192.0.2.0″. The display name can be modified at a later time. The display name is used in the IBM Tivoli Storage Productivity Center, and does not affect the element manager itself.

Connection
Use this section to define connection information for the element manager.

***IP Address**
Enter the DNS name or IP address of the host system where the element manager is located. You can enter an IPv4 or IPv6 name or address, depending on what address types are configured on the element manager and on the system where Tivoli Storage Productivity Center is installed. Enter the following information:

- Enter the element manager's host name that the name service can use to resolve to an IPv4 or IPv6 address.
- Enter an IPv4 address in the following format: ###.###.###.###. For example: 192.0.2.0
- Enter an IPv6 address. These addresses are typically written as eight groups of four hexadecimal digits, but might vary from site to site. For example: 2001:db8:0:0:1812:5b9e:3391:3156. You can use IPv6 addresses if the system where the Tivoli Storage Productivity Center is installed is IPv6 or dual stack (IPv4 and IPv6) enabled.

Notes:

- The host system also represents the system where the corresponding DS8000 Storage Manager is installed.
- If the system on which the element manager is located is IPv6-only, the Tivoli Storage Productivity Center servers

must be installed on an IPv6-only or dual stack (IPv4 and IPv6 addresses are configured) machine.

***Protocol**
Select the protocol (**http** or **https**) for the element manager.

***Port** Enter the port number on which the element manager listens for requests. When a selection is made in the **Protocol** field, the default HTTP port number 8451 or default HTTPS port number 8452 is automatically entered in the **Port** field. Modify the **Port** field if your element manager is configured to use a different port number for the specified protocol.

OK Click this button to add the element manager. IBM Tivoli Storage Productivity Center tests the connection to the element manager. If successful, the element manager appears on the Element Management page. If unsuccessful, an error message appears indicating that the element manager at the specified IP address could not be added.

4. Click **Add** to add the element manager. Tivoli Storage Productivity Center tests the connection to the element manager. If successful, the element manager appears in the table on the Element Management page and in the **Select a view** menu. If unsuccessful, an error message appears indicating that the element manager at the specified IP address could not be added.

   Click **Cancel** to close the window without attempting to add the element manager.

Adding DS8000 element managers using CIMOM discovery:

You can add more than one DS8000 element manager automatically by running CIMOM discovery. The element managers discovered through the CIMOM discovery process are automatically added to the **Element Manager** tab in the Configuration Utility.

Use these steps to:

- Run a discovery against multiple CIMOMs and each of those CIMOMs has a registered storage system.
- Assign more than one registered storage system per CIMOM.

This procedure does not physically add the CIMOM to your environment, but provides Tivoli Storage Productivity Center with the information it needs to communicate with that CIMOM. Element managers that are added through CIMOM discovery might exist on a Hardware Management Console (HMC) other than the HMC where the CIMOM is located.

To add an element manager using CIMOM discovery in Tivoli Storage Productivity Center, complete the following steps:

1. In the navigation tree expand **Administrative Services → Data Sources → CIMOM Agents**
2. In the right pane, click **Add CIMOM**. The Add CIMOM dialog box is displayed.
3. Define the information that is required to add the CIMOM using the following fields. Required fields are specified in bold in the dialog box. In the following field descriptions, required fields are marked by an asterisk.

   ***Host** Enter the IP address of the computer on which the CIMOM is installed.

You can enter an IPv4 or IPv6 address depending on what is supported on the computer where the CIMOM is located:

- Enter an IPv4 address in the following format: ###.###.###.###. For example: 192.168.1.65
- Enter an IPv6 address. These addresses are typically written as eight groups of four hexadecimal digits, but might vary from site to site. For example: 2001:DB8::/32:1812:5b9e:3391:3156. You can use IPv6 addresses if the computer where Tivoli Storage Productivity Center is installed is IPv6 or dual stack (IPv4 and IPv6) enabled.

Note: If the computer on which the CIMOM is located is IPv6-only, the Tivoli Storage Productivity Center servers must be installed on an IPv6-only or dual stack (IPv4 and IPv6 enabled) computer.

Port
Enter the port through which Tivoli Storage Productivity Center contacts the CIMOM. You can modify the default port number if required for your environment.

Username
If required by the CIMOM, enter the user name that Tivoli Storage Productivity Center uses to log on to the CIMOM.

Password
Enter the password for the user name that is defined in the **Username** field.

Password Confirm
Enter the same password that you entered in the **Password** field.

***Interoperability Namespace**
Enter the interoperability namespace of the CIMOM. This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information.

Protocol
Select the protocol (**http** or **https**) for the CIMOM.

Truststore Location
Enter the location (path on this computer) of the truststore for the CIMOM. The truststore is the certificate file for certificate-based authentication in the https protocol.

Truststore Passphrase
Enter the passphrase of the truststore for the CIMOM.

Display Name
Enter the name for the CIMOM as you want it to be shown in Tivoli Storage Productivity Center.

Description
Enter a short description for the CIMOM as you want it to be shown in Tivoli Storage Productivity Center.

Test CIMOM connectivity before adding
Select this check box to ensure that communication to the CIMOM is working properly before the CIMOM is added. If this box is not checked, the CIMOM is added without verifying connectivity with the given parameters.

Important: Failure to check the CIMOM connection might cause a failure later if the CIMOM cannot be scanned or probed.

Save Click this button add the CIMOM.

- If you selected **Test CIMOM connectivity before adding**, Tivoli Storage Productivity Center attempts to connect to the CIMOM using the information provided in this dialog box. If the logon information for a CIMOM is incorrect, an error occurs and an associated error message dialog box is displayed.
- If you did not select **Test CIMOM connectivity before adding**, the CIMOM is associated with the element manager but connectivity with Tivoli Storage Productivity Center is not verified.

Cancel
Click this button to close the dialog box without adding the CIMOM.

4. Expand **Administrative Services → Discovery** in the navigation tree.
5. Right-click **CIMOM** and select **Run Now**. The CIMOM discovery process adds the element managers that are managing the same storage system as the CIMOM.

Starting element manager from the Configuration Utility

You can manage a storage subsystem by starting the element manager that you have created for it.

Start the Tivoli Storage Productivity Center user interface and log in. Complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** in the navigation tree.
2. Click **Configuration Utility**.
3. Click the **Element Manager** tab in the right pane.
4. Select the existing element manager.**Launch Element Manager**
5. Click **Select Action → Launch Default Element Manger** to open the Configure Element Manager page.
6. Click **OK** on the message page.
7. Enter login information for the element manager and click **OK**.

   Tip: The first time you start the element manager you will see a security warning. Accept the warning to create a new security exception for the DS8000.

The DS8000 storage manager interface appears and you can administer the storage subsystem.

Modifying an element manager

You can modify the authentication settings of an existing element manager.

Start the Tivoli Storage Productivity Center user interface and log in. Complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** in the navigation tree.
2. Click **Configuration Utility**.
3. Click the **Element Manager** tab in the right pane.
4. Select an element manager.
5. Click **Select Action → Modify Element Manager** to open the Configure Element Manager page.

Tip: If you only want to view the information for this element manager, click **Cancel** to close the Configure Agent panel without making any changes.

Username
The user ID used for authentication, if required by the CIMOM.

Password and Password Confirm
The password, if required by the CIMOM.

Display Name
The name that appears in the DS8000 Welcome page list.

Note: Do not use special characters such as & and # in the display name. This causes the display name to be truncated.

6. Click **OK** to save the changes. When complete, a confirmation popup will open. Click **OK**. The changes are updated in the display on the Element Manager page.

Remove Element Manager

You can remove a DS8000 element manager from the list of element managers. Removing an element manager from Tivoli Storage Productivity Center does not affect the DS8000 Storage Manager on the HMC.

Start the Tivoli Storage Productivity Center user interface and log in. Complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** in the navigation tree.
2. Click **Configuration Utility**.
3. Click the **Element Manager** tab in the right pane.
4. Select the element manager you are removing.

   Note: You can remove only one element manager from the list at a time.
5. Click **Select Action → Remove Element Manager**.
6. The confirmation window is displayed. Click **OK** to remove the element manager.

Test Element Manager Connection

This topic explains how to test an element manager connection. ###WHAT TO TEST FOR### - ###WHY WOULD YOU TEST###

Start the Tivoli Storage Productivity Center user interface and log in. Complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** in the navigation tree.
2. Click **Configuration Utility**.
3. Click the **Element Manager** tab in the right pane.
4. Select an element manager.
5. Click **Select action → Test Element Manager Connection**.

Tivoli Storage Productivity Center tests the status of the element manager connection and updates the status column in the Welcome page list.####WHERE IS THE WELCOME PAGE###

Add CIMOM connection

The Add CIMOM Connection option allows you to manually associate a CIMOM with an Element Manager. The CIMOM that you associate with an Element

Manager manages the same storage subsystem as that Element Manager. Use the Add CIMOM Connection window to define information about the CIMOM.

Start the Tivoli Storage Productivity Center user interface and log in. Complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** in the navigation tree.
2. Click **Configuration Utility**.
3. Click the **Element Manager** tab in the right pane.
4. Select the element manager from the list of known element managers.
5. Click **Select action → Add CIMOM Connection**.
6. Complete the form in the Configure Agent page with the following information:

   Username
   The user ID used for authentication, if required by the CIMOM.

   Password and Password Confirm
   The password, if required by the CIMOM.

   Display Name
   The name of the CIMOM, as specified by the CIMOM provider, that will appear in the DS8000 list.

   Description
   A description of the CIMOM.

   Connection

   Interoperability Namespace
   This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how IBM Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information.

   IP Address
   The IP address of the CIMOM. For example: 7.97.128.551

   Protocol
   The version of the cim-xml protocol. Can be http or https

   Port The port on which the CIMOM is listening. The default CIMOM ports are:

   http 5988

   https 5989

   Note: For DS8000 R3 or later, port 6989 is used for the embedded CIM agent instead of port 5989.

   Truststore Location
   The location (path on this computer) of a certificate file for certificate based authentication in the https protocol.

   Truststore Password
   The password for the truststore location.

   Test Connection
   Click the **Test Connection** box. This will have Tivoli Storage

Productivity Center ensure that communication to the CIMOM is working properly before adding information about that CIMOM.

7. Click **Add** to apply the settings. A confirmation pop-up opens. Click **OK** to associate the CIMOM with the element manager.

   Tip: Click **Test Connection** to attempt to connect to the CIMOM using the settings you just entered. If a connection to the new CIMOM could not be made, an error message dialog box is displayed and you can correct the information.

After you complete these steps the following occurs:

- A check mark appears in the CIMOM column next to the associated element manager on the DS8000 element manager tab.
- The CIMOM is added to the ###NAME PAGE###. You can find it by clicking **Administrative Services** → **Data Sources** → **CIMOM agents**. ###VERIFY
- System Storage Productivity Center runs a discovery job to discover the storage entities that are managed by that CIMOM. The discovery job might take some time to complete.

If Tivoli Storage Productivity Center can communicate with the CIMOM and the discovery job completes successfully, you can then run data collection jobs (probes, scans, subsystem performance monitors, and so forth) against the DS8000 storage subsystems and other storage entities managed by that CIMOM.

Modify CIMOM Connection

This topic shows you how to view or modify an existing CIMOM connection.

Start the Tivoli Storage Productivity Center user interface and log in. Complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** in the navigation tree.
2. Click **Configuration Utility**.
3. Click the **Element Manager** tab in the right pane.
4. Select the CIMOM from the list of known CIMOMs.
5. Click **Select action** → **Modify CIMOM Connection**.
6. To modify the CIMOM, complete the form in the Configure Agent page with following information:
7. Select **Modify CIMOM Connection** from the **Select action** menu.
8. You can modify the following fields:

   Username
   The user ID used for authentication.

   Password and Password Confirm
   The password, if required by the CIMOM.

   Display Name
   The name that will appear in the DS8000 Welcome page list.

   Description
   A description of the CIMOM.

   Interoperability Namespace
   This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and

determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information.

Test Connection
If not selected, click the **Test Connection** box. This will have Tivoli Storage Productivity Center ensure that communication to the CIMOM is working properly before adding information about that CIMOM.

9. Click **OK** to save the changes. When complete, a confirmation pop-up opens. Click **OK**. The changes are displayed in the Welcome page list

Remove CIMOM Connection

You can remove a CIMOM connection from an element manager. Removing the CIMOM connection is removed, the element manager remains.

Start the Tivoli Storage Productivity Center user interface and log in. Complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** in the navigation tree.
2. Click **Configuration Utility**.
3. Click the **Element Manager** tab in the right pane.
4. Select an element manager.
5. Click **Select action → Remove CIMOM Connection**.
6. A confirmation pop-up prompts you to confirm the removal of the CIMOM connection. To remove the CIMOM connection click **OK**.

   To exit without removing the CIMOM connection, click **Cancel**.

Configuring other element managers for DS8000

DS8000 Storage Manager is the element manager that is provided with IBM System Storage DS8000. However, you can add other element managers for DS8000 as described in this section.

Adding user-defined element managers for DS8000

A *user-defined element manager* for IBM System Storage DS8000 is an element manager other than DS8000 Storage Manager. Use the Configuration Utility to add user-defined element managers for IBM System Storage DS8000 in IBM Tivoli Storage Productivity Center.

User-defined element managers are not displayed in the **DS8000 Element Manager** table. These element managers are displayed in the various launch in context menus that are provided in IBM Tivoli Storage Productivity Center.

To add user-defined element managers in Tivoli Storage Productivity Center, follow these steps:

1. In the navigation tree, expand **IBM Tivoli Storage Productivity Center**. Left-click **Configuration Utility**. In the right pane, click the **Element Manager** tab.
2. In the **DS8000 Element Manager** section, select the row in the DS8000 element manager table that contains the device to which you want to add a user-defined element manager.
3. Select **Add User Defined Element Manager** from the **Select Action** list. The Add User Defined Element Manager dialog box is displayed.
4. Enter a Uniform Resource Locator (URL) or command to start the element manager, and click **Save**.

   To exit without adding a element manager, click **Cancel**.

Removing user-defined element managers for DS8000

Use the Configuration Utility to remove user-defined element managers for IBM System Storage DS8000 in IBM Tivoli Storage Productivity Center.

To remove a user-defined element manager from Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **DS8000 Element Manager** section, select the row in the DS8000 element manager table that contains the device from which you want to remove a user-defined element manager.
3. Select **Remove Element Manager** from the **Select Action** list. The Remove Element Manager dialog box opens.
4. Select the element manager that you want to remove from the list and click **OK**. The element manager is removed. When you remove an element manager, it is removed only from the list of element managers that are visible to Tivoli Storage Productivity Center.

   To exit the window without removing the element manager, click **Cancel**.

Managing other storage systems using element managers

Use the Configuration Utility in IBM Tivoli Storage Productivity Center to add and manage element managers for storage devices other than IBM System Storage DS8000. *Element managers* are applications that are used to configure and manage a storage device. You manage connections to the element managers that are provided with the storage device or you can specify other element managers for a device.

The **Element Manager** tab in the Configuration Utility is used to add and manage element managers. The **Element Manager** tab contains multiple sections. The **Other Device Element Manager** section shows element managers that have been added using a discovery job or existing element managers that were added to Tivoli Storage Productivity Center prior to release 4.1.

If you are gathering data from IBM XIV Storage System, the XIV Storage System element manager is displayed in the **Other Device Element Manager** section. However, the tasks available for the other element managers in this section are not available for the XIV Storage System element manager. Use the Storage Subsystem page to work with the XIV Storage System element manager (see “Managing storage subsystems” on page 411).

To expand or collapse the **Other Device Element Manager** section, click the + or - icon to the left of the section heading.

Other Device Element Manager section

The **Other Device Element Manager** section contains a table listing the element managers or external tools that are visible to Tivoli Storage Productivity Center and a **Select Action** list. This list enables you to manage existing element managers. The menu items that are shown in the **Select Action** list depend on whether you have selected an element manager in the table. This table contains the following fields:

Element Manager
Shows the IP address of the element manager or the display name if a display name was defined.

Connection Status
Shows the status of the element manager. The values that appear in this field are:

- **Online** (The element manager is accessible.)
- **Offline** (The element manager is not accessible)

Device
The name for the device as defined in Tivoli Storage Productivity Center.

UID/PWD
Shows whether a user ID and password have been defined for authentication to the element manager. The values that appear in this field are:

- **Yes** (A user ID and password have been defined.)
- **No** (A user ID and password have not been defined.)

CIMOM
Shows whether a CIMOM is configured for the element manager. The values that appear in this field are:

- **Yes** (A CIMOM is configured.)
- **No** (A CIMOM is not configured.)

LIC Enabled
Shows whether a launch descriptor file (LDF) has been registered for the element manager. An LDF is an XML file that describes the menus and the detailed uniform resource locators (URLs) for applications. An LDF is required to enable the launch in context feature for the device.

The values that appear in this field are:

- **Yes** (Launch in context is enabled.)
- **No** (Launch in context is not enabled.)

Adding DS8000 element managers using CIMOM discovery

You can add more than one DS8000 element manager automatically by running CIMOM discovery. The element managers discovered through the CIMOM discovery process are automatically added to the **Element Manager** tab in the Configuration Utility.

Use these steps to:

- Run a discovery against multiple CIMOMs and each of those CIMOMs has a registered storage system.
- Assign more than one registered storage system per CIMOM.

This procedure does not physically add the CIMOM to your environment, but provides Tivoli Storage Productivity Center with the information it needs to communicate with that CIMOM. Element managers that are added through CIMOM discovery might exist on a Hardware Management Console (HMC) other than the HMC where the CIMOM is located.

To add an element manager using CIMOM discovery in Tivoli Storage Productivity Center, complete the following steps:

1. In the navigation tree expand **Administrative Services → Data Sources → CIMOM Agents**
2. In the right pane, click **Add CIMOM**. The Add CIMOM dialog box is displayed.
3. Define the information that is required to add the CIMOM using the following fields. Required fields are specified in bold in the dialog box. In the following field descriptions, required fields are marked by an asterisk.

   ***Host** Enter the IP address of the computer on which the CIMOM is installed. You can enter an IPv4 or IPv6 address depending on what is supported on the computer where the CIMOM is located:

   - Enter an IPv4 address in the following format: ###.###.###.###. For example: 192.168.1.65
   - Enter an IPv6 address. These addresses are typically written as eight groups of four hexadecimal digits, but might vary from site to site. For example: 2001:DB8::/32:1812:5b9e:3391:3156. You can use IPv6 addresses if the computer where Tivoli Storage Productivity Center is installed is IPv6 or dual stack (IPv4 and IPv6) enabled.

   Note: If the computer on which the CIMOM is located is IPv6-only, the Tivoli Storage Productivity Center servers must be installed on an IPv6-only or dual stack (IPv4 and IPv6 enabled) computer.

   Port Enter the port through which Tivoli Storage Productivity Center contacts the CIMOM. You can modify the default port number if required for your environment.

   Username
   If required by the CIMOM, enter the user name that Tivoli Storage Productivity Center uses to log on to the CIMOM.

   Password
   Enter the password for the user name that is defined in the **Username** field.

   Password Confirm
   Enter the same password that you entered in the **Password** field.

   ***Interoperability Namespace**
   Enter the interoperability namespace of the CIMOM. This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information.

   Protocol
   Select the protocol (**http** or **https**) for the CIMOM.

   Truststore Location
   Enter the location (path on this computer) of the truststore for the CIMOM. The truststore is the certificate file for certificate-based authentication in the https protocol.

   Truststore Passphrase
   Enter the passphrase of the truststore for the CIMOM.

   Display Name
   Enter the name for the CIMOM as you want it to be shown in Tivoli Storage Productivity Center.

Description
Enter a short description for the CIMOM as you want it to be shown in Tivoli Storage Productivity Center.

Test CIMOM connectivity before adding
Select this check box to ensure that communication to the CIMOM is working properly before the CIMOM is added. If this box is not checked, the CIMOM is added without verifying connectivity with the given parameters.

Important: Failure to check the CIMOM connection might cause a failure later if the CIMOM cannot be scanned or probed.

Save Click this button add the CIMOM.

- If you selected **Test CIMOM connectivity before adding**, Tivoli Storage Productivity Center attempts to connect to the CIMOM using the information provided in this dialog box. If the logon information for a CIMOM is incorrect, an error occurs and an associated error message dialog box is displayed.
- If you did not select **Test CIMOM connectivity before adding**, the CIMOM is associated with the element manager but connectivity with Tivoli Storage Productivity Center is not verified.

Cancel
Click this button to close the dialog box without adding the CIMOM.

4. Expand **Administrative Services → Discovery** in the navigation tree.
5. Right-click **CIMOM** and select **Run Now**. The CIMOM discovery process adds the element managers that are managing the same storage system as the CIMOM.

Modifying other element managers

Use the Configuration Utility to modify information for element mangers in IBM Tivoli Storage Productivity Center.

To modify an element manager other than an IBM System Storage DS8000 element manager in Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **Other Device Element Manager** section, select the element manager that you want to modify from the element manager table.
3. Select **Modify Element Manager** from the **Select Action** list. The Modify Element Manager dialog box is displayed.
4. Modify any of the following fields:

Username
Enter the user name that Tivoli Storage Productivity Center will use to log on to the element manager.

Password
Enter the password associated with the user name that is defined in the **Username** field.

Password Confirm
Enter the same password that you entered in the **Password** field.

Display Name
Enter a display name for the element manager. The display name is shown in the **Element Manager** field in the element manager table. If this field is blank, the IP address of the computer on which the element manager is located is shown in the **Element Manager** field of the table. The display name is used in Tivoli Storage Productivity Center and does not affect the element manager itself.

Save Click this button to modify the element manager information. Tivoli Storage Productivity Center tests the connection to the element manager. If the connection is successful, the element manager information is saved. If the connection is unsuccessful, an error message appears indicating that the element manager at the specified IP address could not be added.

Cancel
Click this button to close the dialog box without modifying the element manager information

Starting other element managers

Use the Configuration Utility to start element managers from IBM Tivoli Storage Productivity Center. You can also start element managers from other points in Tivoli Storage Productivity Center.

To start element managers other than IBM System Storage DS8000 element managers from the Configuration Utility follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **Other Device Element Manager** section, click the element manager that you want to start from the element manager table.
3. Use one of the following methods to start the element manager:
   - To start the default element manager for a storage system, select **Launch Default Element Manager** from the **Select Action** list.
   - To start another element manager for a storage system, right click the element manager that you want to start and click **Launch Element Manager**. Select the element manager that you want to start from the menu.

If the connection information for an element manager is valid, the element manager is started. If unsuccessful, an error message is displayed. The error message that you receive describes one of the following problems:

- that the element manager at the specified IP address could not be started
- that the required connection information for the element manager has not been provided
- that the element manager is not installed and configured correctly (for local element manager installations such as XIV Storage Manager)

If you are prompted to enter additional information such as an ID and password for the element manager, complete the following steps:

1. Click **OK** in the message dialog box. The Modify Element Manager dialog box is displayed.
2. Enter the required information for the element manager and click **Save**.

Testing the connection to other element managers

Use the Configuration Utility to test the connection to element managers from IBM Tivoli Storage Productivity Center.

To test the connection to an element manager other than an IBM System Storage DS8000 element manager from Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **Other Device Element Manager** section, select the element manager for which you want to test the connection from the element manager table.
3. Select **Test Element Manager Connection** from the **Select Action** list. IBM Tivoli Storage Productivity Center tests the status of the element manager connection. A message dialog box that states whether the test passed or failed is displayed.

Adding CIMOM connections for other element managers

Use the Configuration Utility to add CIM object manager (CIMOM) connections for other element managers in IBM Tivoli Storage Productivity Center.

To add a CIMOM connection for an element manager other than an IBM System Storage DS8000 element manager in Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **Other Device Element Manager** section, select the element manager to which you want to add a CIMOM connection.
3. Select **Add CIMOM Connection** from the **Select Action** list. The Add CIMOM Connection dialog box is displayed.
4. Define the information that is required to connect to the element manager using the following fields. Required fields are specified in bold in the dialog box. In the following descriptions, required fields are marked by an asterisk.

   ***Host** Enter the IP address of the computer on which the CIMOM is installed. You can enter an IPv4 or IPv6 address depending on what is supported on the computer where the CIMOM is located:

   - Enter an IPv4 address in the following format: ###.###.###.###. For example: 192.168.1.65
   - Enter an IPv6 address. These addresses are typically written as eight groups of four hexadecimal digits, but might vary from site to site. For example: 2001:DB8::/32:1812:5b9e:3391:3156. You can use IPv6 addresses if the computer where Tivoli Storage Productivity Center is installed is IPv6 or dual stack (IPv4 and IPv6) enabled.

   Restriction: If the computer on which the CIMOM is located is IPv6-only, the Tivoli Storage Productivity Center servers must be installed on an IPv6-only or dual stack (IPv4 and IPv6 enabled) computer.

   Port Enter the port through which Tivoli Storage Productivity Center contacts the CIMOM. You can modify the default port number if required for your environment.

Username
If required by the CIMOM, enter the user name that Tivoli Storage Productivity Center uses to log on to the CIMOM.

Password
Enter the password for the user name that is defined in the **Username** field.

Password Confirm
Enter the same password that you entered in the **Password** field.

***Interoperability Namespace**
Enter the interoperability namespace of the CIMOM. This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information.

Protocol
Select the protocol (**http** or **https**) for the CIMOM.

Truststore Location
Enter the location (path on this computer) of the truststore for the CIMOM. The truststore is the certificate file for certificate-based authentication in the https protocol.

Truststore Passphrase
Enter the passphrase of the truststore for the CIMOM.

Display Name
Enter the name for the CIMOM as you want it to be shown in Tivoli Storage Productivity Center.

Description
Enter a short description for the CIMOM as you want it to be shown in Tivoli Storage Productivity Center.

Test CIMOM connectivity before adding
Select this check box to ensure that communication to the CIMOM is working properly before the CIMOM is added. If this box is not checked, the CIMOM is added without verifying connectivity with the given parameters.

> **Important:** Failure to check the CIMOM connection might cause a failure later if the CIMOM cannot be scanned or probed.

Save Click this button to associate the CIMOM with the element manager.

- If you selected **Test CIMOM connectivity before adding**, Tivoli Storage Productivity Center attempts to connect to the CIMOM using the information provided in this dialog box. If the logon information for a CIMOM is incorrect, an error occurs and an associated error message dialog box is displayed.
- If you did not select **Test CIMOM connectivity before adding**, the CIMOM is associated with the element manager but connectivity with Tivoli Storage Productivity Center is not verified.

Cancel
Click this button to close the window without associating a CIMOM with the element manager.

If the CIMOM connection is successful, the following actions occur after you click **Save**:

- **Yes** is displayed in the CIMOM column of the element manager table.
- The CIMOM is added to the **Administrative Services → Data Sources → CIMOM Agents** window.
- A discovery job runs automatically to discover the storage entities managed by that CIMOM. The discovery job might take some time to complete. After the discovery job completes, you can run data collection jobs (probes, scans, system performance monitors, and so forth) against the storage systems managed by that CIMOM.

Modifying CIMOM connections for other element managers

Use the Configuration Utility to modify CIM object manager (CIMOM) connections for element mangers in IBM Tivoli Storage Productivity Center.

To modify a CIMOM connection for an element manager other than an IBM System Storage DS8000 element manager in Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **Other Device Element Manager** section, select the element manager for which you want to modify the CIMOM connection from the element manager table.
3. Select **Modify CIMOM Connection** from the **Select Action** list. The Modify CIMOM Connection dialog box is displayed.
4. You can modify the following fields:

   Username
   If required by the CIMOM, enter the user name that Tivoli Storage Productivity Center uses to log on to the CIMOM.

   Password
   Enter the password for the user name that is defined in the **Username** field.

   Password Confirm
   Enter the same password that you entered in the **Password** field.

   Interoperability Namespace
   Enter the interoperability namespace of the CIMOM. This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information.

   Display Name
   Enter the name for the CIMOM as you want it to be shown in Tivoli Storage Productivity Center.

   Description
   Enter a short description for the CIMOM as you want it to be shown in Tivoli Storage Productivity Center.

   Test CIMOM connectivity before adding
   Select this check box to ensure that communication to the CIMOM is

working properly before the CIMOM is added. If this box is not checked, the CIMOM is added without verifying connectivity with the given parameters.

Important: Failure to check the CIMOM connection might cause a failure later if the CIMOM cannot be scanned or probed.

Save Click this button to save the modified information for the CIMOM.

- If you selected **Test CIMOM connectivity before adding**, Tivoli Storage Productivity Center attempts to connect to the CIMOM using the information provided in this dialog box. If the logon information for a CIMOM is incorrect, an error occurs and an associated error message dialog box is displayed.
- If you did not select **Test CIMOM connectivity before adding**, the CIMOM is modified but its connectivity with Tivoli Storage Productivity Center is not verified.

Cancel

Click this button to close the dialog box without modifying the CIMOM connection.

If the CIMOM connection is successful, the following actions occur after you click **Save**:

- **Yes** is displayed in the CIMOM column of the element manager table.
- The CIMOM is added to the **Administrative Services** → **Data Sources** → **CIMOM Agents** window.
- A discovery job runs automatically to discover the storage entities managed by that CIMOM. The discovery job might take some time to complete. After the discovery job completes, you can run data collection jobs (probes, scans, system performance monitors, and so forth) against the storage systems managed by that CIMOM.

Removing CIMOM connections for other element managers

Use the Configuration Utility to remove CIM object manager (CIMOM) connections for element managers in IBM Tivoli Storage Productivity Center.

To remove a CIMOM connection for an element manager other than an IBM System Storage DS8000 element manager, complete the following steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **Other Device Element Manager** section, select the element manager that you want to remove from the element manager table.
3. Select **Remove CIMOM Connection** from the **Select Action** list. A confirmation dialog box is displayed.
4. Click **OK**. The CIMOM connection is removed and **No** is displayed in the CIMOM column of the element manager table.

   To exit the window without removing the CIMOM connection, click **Cancel**.

Adding user-defined element managers for other storage systems

A *user-defined element manager* is an element manager other than the manager that is provided with the storage system. Use the Configuration Utility to add user-defined element managers in IBM Tivoli Storage Productivity Center.

User-defined element managers are not displayed in the **Other Device Element Manager** table. These element managers are displayed in the various launch in context menus that are provided in IBM Tivoli Storage Productivity Center.

To add user-defined element managers for a storage device other than DS8000 in Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **Other Device Element Manager** section, select the select the row in the element manager table that contains the storage system to which you want to add a user-defined element manager.
3. Select **Add User Defined Element Manager** from the **Select Action** list. The Add User Defined Element Manager dialog box is displayed.
4. Enter a Uniform Resource Locator (URL) or command to start the element manager, and click **Save**.

   To exit without adding a element manager, click **Cancel**.

Removing user-defined element managers for other storage systems

Use the Configuration Utility to remove user-defined element managers for storage systems in IBM Tivoli Storage Productivity Center.

To remove a user-defined element managers for a storage device other than DS8000 in Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **Other Device Element Manager** section, select the row in the element manager table that contains the storage system from which you want to remove a user-defined element manager.
3. Select **Remove Element Manager** from the **Select Action** list. The Remove Element Manager dialog box opens.
4. Select the element manager that you want to remove from the list and click **OK**. The element manager is removed. When you remove an element manager, it is removed only from the list of element managers that are visible to Tivoli Storage Productivity Center.

   To exit the window without removing the element manager, click **Cancel**.

Discovering storage resources

This topic describes how you can collect details about storage resources in your environment.

The process of finding resources within an enterprise, including detection of network topology, is called discovery. Use discovery jobs to collect basic

information about the storage resources in your environment, including computers, CIMOMs, storage subsystems managed by CIMOMs, fabrics and switches managed by CIMOMs, NetWare trees, filers, clusters, and tape libraries. You must run discovery jobs before you can collect more detailed information about storage resources using probes, scans, and pings.

Discovery overview

IBM Tivoli Storage Productivity Center supports a variety of discovery types. Use discovery jobs to find new CIMOMs, Out of Band Fabrics, NetWare Filers, Windows Domains, NAS filers, and SAN File System that have been introduced into your environment.

The discovery types in Tivoli Storage Productivity Center are summarized in the following table. Additionally the table summarizes the types of information that each discovery provides.

Table 12. Discovery types, their data sources, and top-level entities that are discovered

Type of discovery	Data source	Top-level entities
CIMOM	CIM Agent	• Fabric (switches) • Storage subsystems • Tape subsystems
Out-of-band	SNMP agents on the switches	Fabric (switches). A fabric discovery gets all the fabric information switches that have been added as out-of-band agents. All information is collected by using the SNMP protocol to send queries across the IP network to the management information bases (MIBs) supported on the switch.
In-band Fabric agents	In-band Fabric agents	Fabric (switches). If you have installed Fabric agents on the computers attached to the fabric, you do not have to perform a specific discovery step for these in-band agents. The fabric is automatically discovered.
NetWare Filer	NetWare agent	NetWare trees
Windows domain, NAS, and SAN File System	Data agent	• Computers • Filers • Clusters • Names or agents
VMWare VI	VMWare VI data source	The data source can be a hypervisor (ESX Server or VirtualCenter).

As a general rule, discovery only provides information about the existence of a resource. Use probes and scans to obtain detailed information about a specific resource after it has been discovered.

CIMOM discovery

This topic describes how the IBM Tivoli Storage Productivity Center manages and monitors disk, switch, and tape devices through Common Information Model Object Managers (CIMOMs).

Note: CIMOM can be referred to by a variety of names, such as CIM agents or SMIS agents. A CIM agent consists of a CIMOM and an SMIS provider for the managed device.

CIM agents on the same subnet as the Tivoli Storage Productivity Center Device server are automatically discovered with CIMOM discovery. If CIM agents are not in the same subnet, they can be discovered using a Service Location Protocol (SLP) Directory Agent (DA) or they can be added manually. If you do not want Tivoli Storage Productivity Center to discover CIM agents in the local subnet, clear the Scan local subnet box in the CIMOM discovery Options panel. To include SLP Directory Agents in the discovery, add the IP address of the DA in the CIMOM discovery Options panel and click **Add**. CIMOM discovery has several purposes. It can scan the local subnet to find CIM agents and contact CIM agents that have been added through the dialog. After it has found a CIM agent, it attempts to log in and discover the storage entities being managed by the CIM agent. If a CIM agent is discovered but requires login information, you need to enter it in the View and Edit CIMOM panel and then rerun the discovery to collect the information for the storage entities. The amount of time CIMOM discovery takes depends on the number of CIM agents, the number of storage entities, and whether you are scanning the local subnet.

CIMOM discovery can be run on a schedule. How often you run it depends on how dynamic your environment is. It must be run to detect a new subsystem. CIMOM discovery also performs basic health checks of the CIM agent and subsystem.

For Tivoli Storage Productivity Center to successfully communicate with the CIM agents, the CIMOM service must be accessible through the IP network. The TCP/IP network configuration on the host where Tivoli Storage Productivity Center is installed must include in its list of domain names, all the domains that contain storage devices that are discoverable by Tivoli Storage Productivity Center.

The CIM agents are installed outside of Tivoli Storage Productivity Center. Contact your device vendor for information about installing the CIM agent.

Limitation when running a CIMOM Discovery against a McDATA CIM agent: When you run a CIMOM discovery against a McDATA CIM agent, the following switch information is collected:

- Manageable McDATA switches that exist in the fabric.
- Unmanageable switches that connect to the manageable McDATA switches. In a homogenous fabric, a McDATA switch is considered unmanageable when it is managed by an Enterprise Fabric Connectivity Manager (Proxy Mode) or a McDATA CIM agent (Direct Mode) other than the McDATA CIM agent against which you ran a CIMOM discovery. In a heterogenous fabric, unmanageable switch is either a non-McDATA switch or a McDATA switch which is managed by different Enterprise Fabric Connectivity Manager in proxy mode or different McDATA CIM agent in direct mode

You can use Tivoli Storage Productivity Center to view detailed information about the manageable McDATA switches and collect performance data about them. For unmanageable switches, you can use Tivoli Storage Productivity Center to view

their WWN and port numbers which connect to manageable McDATA switches only, and you cannot collect performance data about them.

SLP registration with slptool

This topic describes how IBM Tivoli Storage Productivity Center uses Service Location Protocol (SLP) discovery.

SLP discovery requires that all of the CIMOMs that Disk Manager discovers are registered using the SLP. However, for CIMOMs outside of the IP subnet, you need to use an SLP DA and register the CIMOM using **slptool**. In a non-multicast network environment, you must use the **slptool** as SLP can only discover CIMOMs that are registered in its IP subnet. When you use **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 machine 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.

SMI-S discovery

This topic describes how the IBM Tivoli Storage Productivity Center can discover storage devices that comply with the Storage Management Initiative Specification (SMI-S).

SMI-S-compliant CIM Agents are discovered using the Service Location Protocol (SLP). The Tivoli Storage Productivity Center server software performs SLP discovery on the network. The SLP User Agent (UA) looks for all registered services with a service type of **service:wbem**.

The Tivoli Storage Productivity Center performs the following discovery tasks:

- Locates individual storage devices
- Retrieves vital characteristics for those storage devices
- Populates the internal Tivoli Storage Productivity Center databases with the discovered information

Device server and discovery on a SAN

This topic describes how the Device server uses SMI-S CIMOMs, in-band events, and out-of-band SNMP traps to start a discovery operation on the SAN.

The Storage Management Initiative Specification (SMI-S) standard is now used for discovery and comprehensive probing of Brocade and McDATA switches and fabrics. This means that information previously gathered through in-band and out-of-band agents and SNMP traps can now be gathered through SMI-S CIMOMs for the Brocade and McDATA devices, including topology information, zoning and zone control, zone alias support (Brocade only), and CIM indications of fabric events. This also allows support for heterogeneous fabrics composed of switches from Brocade and McDATA. Note that the data gathering capabilities of in-band and out-of-band agents and SNMP traps are still supported.

The following SMI-S profiles and their respective subprofiles are supported:

Table 13. Supported SMI-S profiles and subprofiles

Profile	Subprofile
Fabric	• Zone Control • Enhanced Zoning and Enhanced Zoning Control
Switch	• Blades • Access Points • Multiple Computer System • Physical Package • Software

The level of support depends on the subprofiles registered with the CIMOM. For example, zone control support requires that the CIMOM has the Zone Control subprofile registered with it.

If a fabric is managed by SMI-S, in-band, and out-of-band agents at the same time, then each time a probe is run on that fabric, all three mechanisms will run concurrently. The use of multiple agent types has benefits, such as redundant collection of information in the event that one type of agent fails. Also, some types of agents provide certain features which others do not. However, the use of multiple agent types can increase network traffic as well as management load on the switches. In general, it is recommended that you use SMI-S Agents as the primary agent type for Brocade fabrics and McDATA fabrics. You can also use in-band agents and out-of-band agents, depending on the storage management needs. See Table 14 on page 193 for a list of the different types of jobs that collect data for the fabrics. You can use this table to determine which agent types to use for your fabric.

In-band events and SNMP traps provide the same information, such as a change that has occurred in the SAN by indicating that a discovery operation should be performed. The in-band events and SNMP traps let the Device server know when something has happened in the SAN and then a discovery is done to identify the changes. The Fabric agents that are running on the managed hosts detect in-band events.

the Device server out-of-band change agent detects the SNMP traps. For the Device server to receive the SNMP traps, configure the device, such as a switch, to send traps to the Device server.

The SMI-S Agents also notify the Device server of changes on the fabric, such as changes to the topology or to the zoning. This is done through CIM Indications. The CIM Indications are subscribed to for a CIMOM SMI-S Agent upon the first successful discovery of the fabric through the CIMOM Discovery job. Some extra configuration may be needed to set this up for environments with systems that have multiple IP addresses. The CIM Indications send specific details of a change on the fabric to the Device server. Rather than re-collect the complete topology and zoning information for every CIM Indication received, the Device server analyzes the CIM Indication and in most cases either sends a small set of post-Indication discovery requests to the SMI-S Agent to collect information pertinent to that event, or sets certain devices as missing without any post-Indication requests made to the SMI-S Agent. For CIM Indications that represent a zoning change, the complete zoning information is re-collected from the SMI-S Agent.

the Device server discovers SAN information by performing the following operations:

- Communicates with Fabric agents. The agents run on the managed hosts (in-band discovery).
- Sends Management Information Base (MIB) queries directly to switches and other devices (out-of-band discovery).
- Sends SMI-S queries to the CIM agents (CIMOM Discovery and Probe).

Host and device information is gathered by the in-band operation. Topology information is gathered by using either or both the in-band, out-of-band, and SMI-S agents. The topology discovery operations of all three agent types provide a similar level of information. However, zone information is available only through certain agent types for certain vendors. See Table 14 for details.

Using both SMI-S agents, in-band agents, and out-of-band agents extend the range of devices that Fabric Manager supports. Some switches are supported only through the in-band mechanism and some are supported only through the out-of-band mechanism. For Brocade and McDATA switches, it is recommended that you use SMI-S agents in the environment, and use in-band and out-of-band agents in addition to these if necessary for redundancy or for the features that the SMI-S agents do not provide. See Table 14 for details.

If there are no agents running on the host systems, out-of-band discovery operations, SMI-S discovery operations, CIM Indication monitoring, and SNMP trap monitoring can be used to monitor the SAN. In this configuration, the Device server can indicate only switch and connection level information, and for Brocade and McDATA fabrics, the zoning information can also be collected. The devices and hosts will appear as unknown entities. You can change the icon and labels of the unknown entities to something more appropriate for your enterprise. If you want to use the Device server in this way, configure the Device server to receive SNMP traps and CIM Indications. CIM Indications are usually automatically subscribed to after CIMOM Discovery, however in some environments, manual configuration of certain settings may be required. The SNMP traps and CIM Indications determine when to start a rediscovery operation. Consider this configuration if you have the following conditions:

- You want to include unsupported systems on your SAN
- You want to monitor machines that you do not directly control

The SNMP traps that are sent by the devices and by the Device server can be sent to the IBM Tivoli Storage Productivity Center graphical user interface (GUI). The events can be displayed in the event browser, or they can be sent to another trap console for monitoring. You then can provide error recovery or other advanced processing, such as paging.

The following table shows the type of information that is gathered for CIMOM/SMI-S, in-band, and out-of-band discovery operations.

Table 14. Information gathered by the different types of discovery and probe jobs

Type of action or job	In–band agent	Out–of–band agent		CIMOM	
		SNMP	BROCADE API	Brocade and McDATA	Cisco and QLogic

Table 14. Information gathered by the different types of discovery and probe jobs (continued)

Type of action or job	In–band agent	Out–of–band agent		CIMOM	
Out-of-band agent discovery job		• Discovers switches and connections within an IP range • Retrieves switch and topology information • Retrieves Cisco VSAN information	• Retrieves Brocade zoning information, including zone aliases		
CIMOM agent discovery job				• Discovers CIMOMs through SLP • Retrieves switch information • Retrieves fabric information	• Discovers CIMOMs through SLP • Retrieves switch information • Retrieves fabric information
Probe job	• 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 information	• Retrieves Brocade zoning information, including zone aliases	• 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
In-band agent is detected or its connectivity changes	• 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 information	• Retrieves Brocade zoning information, including zone aliases		

Table 14. Information gathered by the different types of discovery and probe jobs (continued)

Type of action or job	In–band agent	Out–of–band agent		CIMOM	
SNMP trap	• 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 Cics VSAN information	• Retrieves Brocade zoning information, excluding zone aliases		
Fabric CIM indication					• Retrieves zoning information, including zone alias information • Retrieves switch and topology information • Sets device as missing without further requests to the SMI-S agent

Fabric discovery

This topic describes how the Fabric Manager uses SNMP queries to discover information about selected fabric switches.

Management Information Base (MIB) information is collected from those switches. Out-of-band agents are configured for the Device server and contacted from the server through SNMP.

Discovering computers in the Windows domain, NAS, and SAN File System

This topic describes how a NAS discovery finds SAN File System clusters.

When you run a discovery to find Windows computers on a domain, it discovers only new Windows computers. For IBM Tivoli Storage Productivity Center to recognize a non-Windows computer, you must install an agent on that computer.

A NAS discovery finds NAS Filers, and a SAN File System discovery finds the SAN File System clusters.

For information about configuring your system for NAS and SAN File System discovery, see the section about the Manual NAS/NetWare Server Entry node in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide*.

Discovering NetWare Filer systems that do not have Data agents

This topic describes how you can discover NetWare Filer systems that do not have installed Data agents.

To discover NetWare Filer systems, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** → **Discovery**.
2. Right-click **NetWare Filers** and click **Run Now**.

Discovering Windows systems that do not have Data agents

This topic describes how you can discover Windows systems that do not have installed Data agents.

To discover Windows computers that do not have Data agents in your Windows domain, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** → **Discovery**.
2. Right-click **Windows Domain, NAS, and SAN FS** and click **Run Now**.
3. After you have identified the computers, you can install Data agents on the computers to get detailed information about the computers and file systems.

Discovering VMware VI data sources

This topic describes how you can discover VMware VI data sources.

To display information about VMware VI data sources, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** → **Discovery** → **VMware VI Data Source**.
2. Right-click **VMware VI Data Source**. Click **Run Now**.

Discovering computers

This topic describes how you must install Data agents on the computers that you want to be discovered.

You must register the Data agents with the Agent Manager. The IBM Tivoli Storage Productivity Center uses information from Agent Manager to provide information about the computers. An agent runs directly on the computer it manages and provides information about that computer. To gather information from the agents, you create pings, probes, and scans. A scan discovers the largest amount of information but is the most time-consuming and intensive job.

When an agent shuts down and Tivoli Storage Productivity Center detects that the agent is no longer responding, the agent it is put into the missing state (a black circle) to indicate that it is not detectable. IBM Tivoli Storage Productivity Center for Data also does a heartbeat ping to the agents every 30 minutes. If this ping detects the host while the agent is not running, the host is put into the warning state (a yellow triangle) to indicate that the host is responding but the agent is not.

Note: This ping is only sent to computers that were running Data agents.

If Tivoli Storage Productivity Center can no longer communicate with an agent, the agent is marked as missing (a black circle) . Information about the host continues to be displayed until the host is removed from the database. To remove

the host from the database, you must uninstall the Data agent and Fabric agent. The host computer with the Data agent is automatically removed.

Discover storage subsystem, tape library, and fabric information

This topic describes how you can discover storage subsystems, tape libraries and fabrics in your environment using CIMOM agents.

CIMOM discovery performs basic health checks of the CIMOMs. Even if you do not expect the number of CIMOMs to vary, be sure to schedule a CIMOM discovery often enough to be kept informed about your environment.

To automatically discover CIMOMs that are visible to the Device server with SLP, follow the instructions in "Discover CIM agents automatically" on page 126.

To add and discover CIMOMs manually, follow the instructions in "Adding a CIM agent manually" on page 124.

Once your CIMOMs have been manually added you must perform a CIMOM discovery to gather information about the CIMOMs so that IBM Tivoli Storage Productivity Center can communicate with them.

To discover CIMOM information for any CIMOMs that you manually added, complete the following steps:

1. Expand **Administrative Services** → **Data Sources** → **CIMOM Agents**. CIMOMs will be listed in the CIMOM Agents table along with their connection status. For each automatically-discovered CIMOM you want to use, click on the magnifying glass to the left of the CIMOM name in the table. The View and Edit window for that CIMOM will open. Enter missing information, such as user name and password, for that CIMOM. Also make sure you check the **Test CIMOM connectivity before updating** box. Click **File** → **Save** on the menu bar. This will save your information and perform a connectivity check based on the information you entered. You will need to repeat this process for each SLP-discovered CIMOM you choose to use.
2. Expand **Administrative Services** → **Discovery** → **CIMOM**. The **Edit CIMOM** page opens in the right pane.
3. Optional: Click the **Options** tab. Enter the IP addresses or host names for the directory agents to be discovered.

   Note: If you do not want to discover every CIMOM in your system's local subnet, clear the Scan Local Subnet check box. If you clear this box the CIMOM discovery will only run on those CIMOMs already listed under **Administrative Services** → **Data Sources** → **CIMOM Agents**. If you chose not to clear this box CIMOM discovery will search for and discover all the CIMOMs in your subnet which may take a long time.

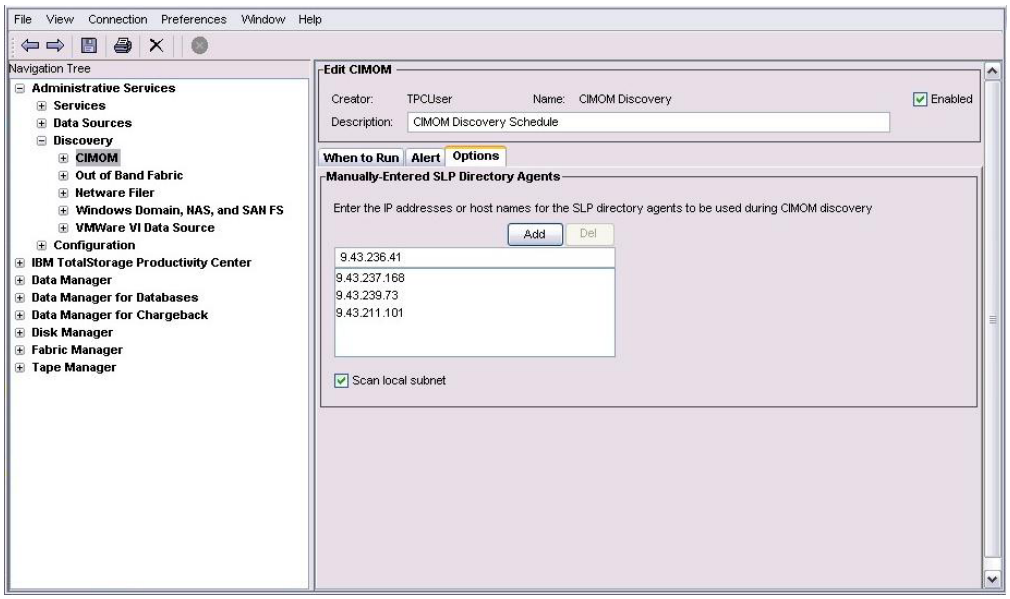

Figure 12. Editing the CIMOM discovery job (Options tab)

4. Click the **When to Run** tab and indicate how often and when to run this discovery job.
5. Click **File** → **Save** to save the job.

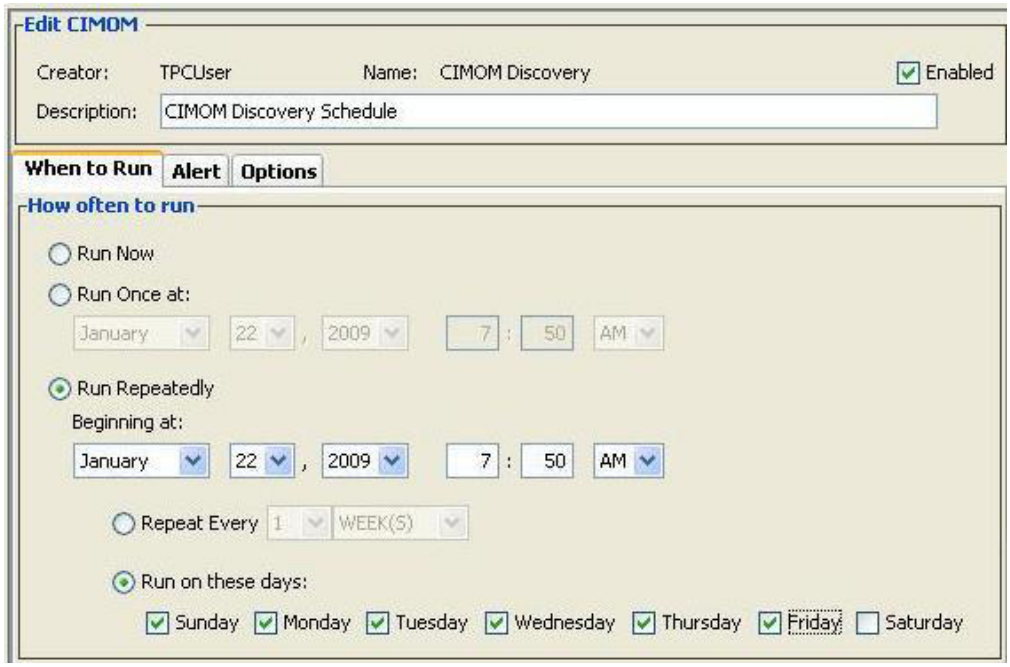

Figure 13. Editing the CIMOM discovery job (When to Run tab)

6. After you have set the parameters for the CIMOM discovery job, perform the CIMOM discovery to gather information on all the CIMOMs you have entered in the Options tab.

   Expand **Administrative Services** → **Discovery**. Right-click on **CIMOM** and click **Run Now**.

7. In the **CIMOM Discovery job submitted** window, click **OK**. To determine if the job has completed, right-click **CIMOM** and click **Update Job Status**. The job color will reflect the status.
8. When the job has finished, view the list of discovered storage subsystems or tape libraries by expanding either **Disk Manager** and click **Storage Subsystems**, or **Tape Manager** and click **Tape Libraries**.

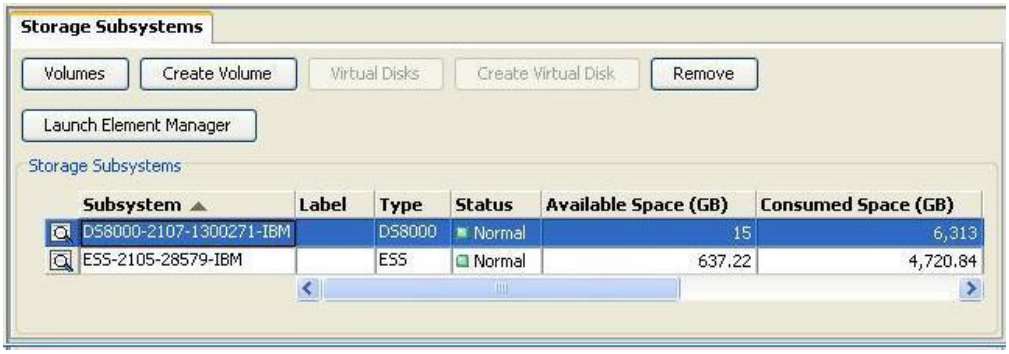

Figure 14. Checking storage subsystems

Confirming discovered storage subsystems and tape libraries

Once the discovery has completed, confirm that it was successful for the CIMOMs you are concerned about by checking the CIMOM discovery job status table.

Storage subsystems

To confirm that the storage subsystems were discovered, expand **Disk Manager** → **Storage Subsystems** in the Navigation Tree. The Storage Subsystems window opens listing all the discovered storage subsystems and giving each one's status, as well as other information.

You can also view information about the storage subsystems in the L0 Topology view.

Verifying storage subsystem discovery in the topology viewer

IBM Tivoli Storage Productivity Center can detect a storage subsystem that is connected to a fabric through queries to the switches in the fabric and through CIMOM queries. However, only basic information is available at this point, and the information from the two discoveries is not correlated. As a consequence, the storage subsystem is displayed in the L0:Storage view as two separate icons:

- After the in-band discovery, as a storage subsystem of unknown status (a blue diamond) identified only by its world-wide name (WWN)
- After the CIMOM discovery, a storage subsystem identified as, for example, an Tivoli Storage Enterprise Storage Server or a DS8000

In order to correlate the information collected through the in-band and CIMOM discoveries, you must run a subsystem probe. Then the storage subsystem is displayed in the topology by only one icon with its type identified.

Tape libraries

To confirm that the tape libraries were discovered, expand **Tape Manager** → **Tape Libraries** in the Navigation Tree.

Verifying tape library discovery in the topology viewer

Tivoli Storage Productivity Center can detect a tape library that is connected to a fabric through queries to the switches in the fabric and through CIMOM queries. However, only basic information is available at this point, and the information from the two discoveries is not correlated. As a consequence, the tape library is displayed in the L0:Storage view as two separate icons:

- After the in-band discovery, as a tape library of unknown status (a blue diamond) identified only by its world-wide name (WWN)
- After the CIMOM discovery, a tape library is identified as, for example, a 3584 or a TS3310

In order to correlate the information collected through the in-band and CIMOM discoveries, you must run a tape probe. Then the tape library is displayed in the topology by only one icon with its type identified.

Discovering in-band fabric agents

This topic describes how to discover Fabric agents (in-band agents) installed on the computers in your environment.

Fabrics are automatically discovered and the Fabric agents can provide information about the fabrics or switches to which the computers are attached. When Fabric agents (in-band agents) are installed on the computers attached to the fabric and these agents are registered with IBM Tivoli Storage Productivity Center, you do not have to perform a specific discovery step for these agents.

For example, if you have McDATA and QLogic switches and have Fabric agents installed on the computers that are attached to these switches, the switches are automatically discovered when you install the Fabric agents.

If you have a McDATA switch, ensure that the switch has the Open Systems Management Server (OSMS) feature enabled. OSMS is required for in-band discovery. The host control option needs to be enabled if you want to do zone control and if no CIM Agent for the fabric is present. Refer to your McDATA documentation for information about these features and how to enable them.

In-band discovery provides this following information:

- Topology information for switches, connections, nodes, and ports of the devices connected to the fabric, zones, zone sets
- Host and device information (collected from in-band agents only)
- Basic information to identify the host and devices in the fabric

Discovering out-of-band Fabric agents

If no CIM Agents for your Brocade switches are configured, you must discover out-of-band Fabric agents to get zone information for the Brocade switches.

To perform Fabric discovery for a Brocade switch, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** → **Agents**.
2. Right-click **Out of Band Fabric** and click **Add**.
3. In the dialog box that is displayed, enter the host name of the switch. If this is a Brocade switch, and no CIM Agent is configured for this fabric, and you want zoning information, check the box for Brocade agents and enter a user ID and

password for the switch. It is not recommended that you add the Brocade user ID and password to all switches; for more information see "Zone configuration" on page 69.

Note: Not all Brocade switches in a fabric should have the admin user ID and password added to gather zoning information. One switch is capable of gathering the information for the entire fabric, but enabling two might be desirable for redundancy. Enabling all of the Brocade switches in a fabric will cause unnecessary activity on the switches. When choosing which switches to enable, thought should be given to the models and firmware levels that make up the fabric. It is best to choose the highest-powered switches, particularly those running the highest levels of firmware, as the ones to enable.

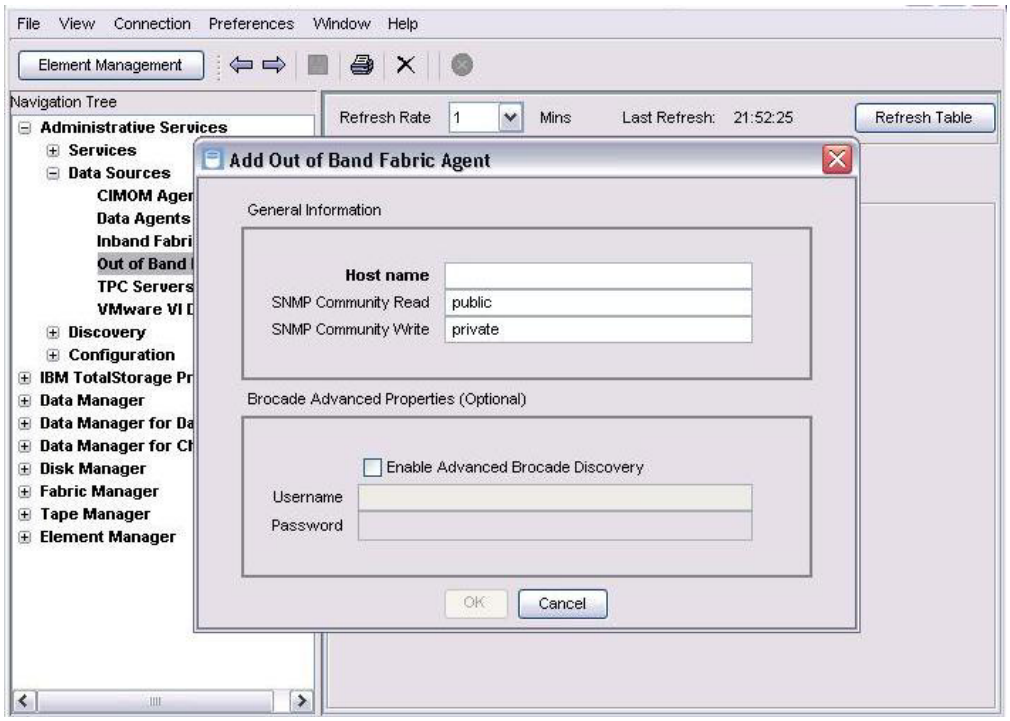

Figure 15. Adding information for a Brocade switch

4. You can also enter an SNMP community if you have it; otherwise the default is "Public." Click **OK**.
5. If you have many switches, and do not want to enter the information manually, you can edit a discovery job to search for switches in specific subnets. To do so, complete the following steps:
   a. In the Navigation Tree pane, expand **Administrative Services** → **Discovery**.
   b. Click **Out of Band Fabric**. The Edit Out of Band panel appears.
   c. Click **Options** to display the Options page.
   d. Specify the IP address ranges that you want to scan in the **From** and **To** fields. You can enter an IPv4 or IPv6 address depending on what is supported within your environment:
      - Enter IPv4 addresses in the following format: ###.###.###.###. For example, if you enter **192.168.1.68** in the **From** field and **192.168.1.254** in the **To** field., IBM Tivoli Storage Productivity Center scans IP addresses that fall in the range between 192.168.x.68 and 192.168.x.254.

- Enter IPv6 addresses according to the standards in your environment. 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:DB8:0000:1234:0000:0000:5678: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.

  For example, if you enter **2001:DB8::0** in the **From** field and **2001:DB8::FF** in the **To** field, Tivoli Storage Productivity Center scans the IP addresses that fall in the range between 2001:DB8:x:0 to 2001:DB8:x:FF.

  The address range that you specify determines how long an SNMP scan takes to run. For example, if the range includes over 256 addresses, the scan might take some time to complete. A warning message will appear if you indicate a range that includes between 257 and 4096 (0xFFF) addresses. This warning message notifies you that the scan will require significant time to run as more addresses are included. Keep in mind that the scan range you enter cannot exceed 4096 (0xFFF) addresses.

  Note: It takes about 5 minutes to search a subnet. For out of band fabric discovery to work, the FC Management MIB and FE MIB must be enabled on the switch. The Brocade API does not support IPv6.

6. Expand **Administrative Services** → **Discovery**
7. Right-click **Out of Band Fabric** and click **Run Now**.
8. After the discovery completes, expand **Fabric Manager** → **Fabrics** to display a list of fabrics.

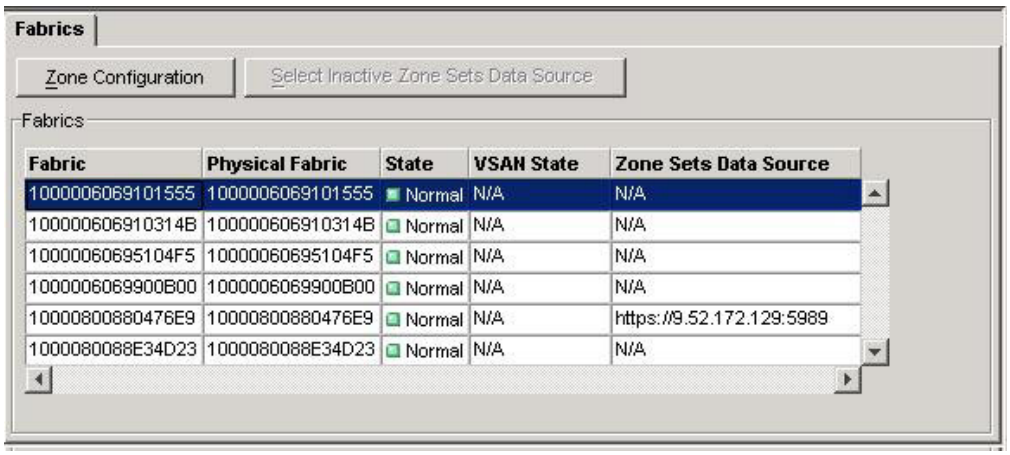

Figure 16. Displaying the fabrics

Gathering switch performance data

This topic describes how to gather performance data from a switch to evaluate the switches in your environment.

To complete the discovery process for switch performance functions, a CIMOM must be enabled for your switch.

- For some switches, if you have the prerequisite firmware, you only need to enable the CIMOM (also referred to as an embedded CIMOM). Switches from Cisco and QLogic have embedded CIMOMs.

- For other switch vendors, you need to install software (the CIMOM proxy) on any host computer and configure it to run providers for one or more switches. Switches from Brocade and McDATA need a CIMOM proxy.

Contact your switch vendor for information on how to enable the CIMOM or install the CIMOM proxy.

To gather switch performance data, complete the following steps:

1. Run a CIMOM discovery for a switch CIMOM.
2. Run a Fabric Probe for the Fabric that the switch is a part of.

For information on creating a switch performance monitor, see "Creating a switch performance monitor" on page 241.

Collecting data

Create data collection jobs to gather information about the storage resources in your environment. There are different data collection jobs within IBM Tivoli Storage Productivity Center depending on the type of information is collected and from which types of resources the information is retrieved: discovery, probes, Tivoli Storage Productivity Center Server probes, scans, and pings.

Collecting storage statistics (probes)

Use probes to collect statistics about the monitored storage resources in your environment, such as computers, storage subsystems, fabrics, volumes, disk controllers, hard disks, and filesystems. Probe jobs can also discover information about new or removed disks and filesystems. You can view the information collected by probes in Asset, Capacity, Storage Subsystems, and System reports.

You can define any number of probes to collect statistics for different storage resources at different times.

- From the **IBM Tivoli Storage Productivity Center** node you can collect statistics on computers, disk controllers, hard disks, clusters, fabrics, storage subsystems, volumes, tape libraries, and filesystems.
- From the **Data Manager for Databases** node you can itemize and create an inventory of the files, instance, logs, and objects that make up your enterprise's monitored RDBMSs.

  Note: The login ID that Data Manager for Databases uses to log into Microsoft SQL Server instances that you want to probe must have "Permit" access. For information about how to set the logins for Microsoft SQL Server instances, see the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide*.

Note: To ensure the accuracy and consistency of the data that appears in the reports, run regularly scheduled probes against the hosts that use/import volumes and the storage subsystems upon which you want to report. You should run probes against storage subsystems after running them against any hosts.

Creating probes

Create and schedule a probe to job to gather information about your storage resources. You can define any number of probes to gather information about different storage resources at different times.

To create a probe, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center → Monitoring → Probes**. For Data Manager for Databases, expand **Data Manager for Databases → Monitoring → Probes**.
2. Right-click **Probes** A pop-up menu appears.
3. Select **Create Probe**.
4. In the **What to Probe** or **Instance** page, perform the following steps:
   a. In the **Available** list box highlight the resources you want to probe.
   b. Click >>. The resources are moved to **Current Selections** list box. Use << to move resources to the **Available** list box.
5. In the **When to Run** page, schedule when you want the probe to run. You have the following options:
   - Run the job immediately
   - Run the job once at a specified time and date. Use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM.
   - Run the job repeatedly according to a schedule you specify. You can:
     – use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM when the job should begin running.
     – indicate how often the job should run. You can use the list boxes to specify the number of minutes, hours, days, weeks, or months.
     – indicate which days on which run the job. A check mark will appear next to the days you select.

   Use the fields in the **How to handle timezones** section to indicate the time zone that Data Manager for Databases should use when scheduling and running the action. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
   - 9:00 AM of the time zone where the server resides
   - 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on a Data agent located in the PST time zone.

   Use the timezone that the server runs in
   Select this option to use the time zone of the location where the Data server resides.

   Use this timezone
   Select this option to indicate a specific time zone for when to run the job. When you select this option, the list box to the right of the field will activate.
6. In the **Alert** page, set the following options:
   - The condition that triggers an alert. The default setting is **Probe Failed**. No other conditions are currently available.
   - The method by which you should be notified of an alert: SNMP Trap, IBM Tivoli Enterprise Console Event, Login Notification, Windows Event Log, UNIX Syslog, running a script, and e-mail. If you select SNMP Trap, Tivoli Enterprise Console Event, or e-mail for the method by which you are notified of an alert, you must define a destination for the method on the **Administrative Services → Configuration → Alert Dispositions** window.
7. Click  to save the probe.
8. Enter a name for the probe and click **OK**. An entry for the probe appears under the **Probes** node.

Editing probes

Change the settings in an existing probe job.

To edit an existing probe, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** → **Monitoring** → **Probes**. For Data Manager for Databases, expand **Data Manager for Databases** → **Monitoring** → **Probes**.
2. Click the name of the probe you want to edit. The content pane displays information about that job.
3. Edit the probe definition.
4. Click  to save the updated probe definition.

Creating fabric probes

Use probes to collect statistics about fabrics and fabric groups.

After you have defined your fabric groups, Fabric Manager enables you to run probes against fabrics and fabric groups to collect information.

To run probes against fabrics, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** → **Monitoring**, right-click **Probes**, and select **Create Probe**.
2. In the **What to Probe** tab enter a description for the probe, and from the **Available** column, select the fabrics and fabric groups to be probed. Move them to the **Current Selections** column by clicking  .
3. In the **When to Run** tab specify when the probe is to begin. You can begin the probe immediately, schedule it to begin at a later time, and specify if it should be run repeatedly.
4. In the **Alert** tab specify the action to be taken if the probe fails. See "Running scripts with alerts" on page 363 for a description of these triggered actions.
5. To save the probe, click **Save**.
6. In the **Save As** window enter the probe name and click **OK**.

Creating tape library probes

Probes collect data from the tape libraries that have been discovered.

You can create tape library probes that collect information from all your tape libraries, or a subset. Probes can be scheduled to run immediately, or at a scheduled point in time. The data that probes collect is put into the database and displayed in various places throughout the product, such as the topology viewer and the **Tape Manager** → **Tape Libraries** node.

To create a tape library probe:

1. Expand **IBM Tivoli Storage Productivity Center** → **Monitoring** → **Probes** and select **Create Probe**
2. Enter the information in the three tabbed pages:
   - **What to Probe** - select the tape libraries and tape library groups to include in this probe
   - **When to Probe** - specify the scheduling information for this probe

- **Alert** - specify the triggering condition that will cause an alert to be generated for this probe, and the actions that should be taken when an alert is generated

3. Click  . Enter a descriptive name for the probe in the **Save As** dialog and click OK to submit the probe.
4. Your new probe will immediately be listed in the **Probes** node. You can right-click the new probe and select **Update Job Status** to show the latest status of the probe.

There are some special situations to note when creating and running probes of tape libraries:

- Do not include more than three tape libraries in the same probe if they are all registered with the same IBM SMI-S Agent for Tape. Doing so can increase the load on the SMI-S agent and can increase the possibility of time outs. To avoid this, create multiple probes with different start times, and split the libraries up among the probes.
- After the list of tape libraries that are registered with the IBM SMI-S Agent for Tape has been altered, a condition can occur where not all of the tape cartridges are returned to the CIM client, in this case, Tivoli Storage Productivity Center. To resolve this situation, restart the IBM SMI-S Agent for Tape; refer to the documentation for the SMI-S Agent for instructions on how to do this.
- When the probe of a tape library has successfully completed, some of the buttons in the Tape Libraries GUI panel might not be enabled. To resolve this situation, close and then redisplay the panel. This will enable the appropriate buttons.

Locating a tape cartridge:

View a list of all the cartridges for each discovered tape library.

Tape libraries can contain many cartridges (sometimes called tape volumes), and it can be difficult to find a particular tape cartridge among all the cartridges for a particular tape drive. IBM Tivoli Storage Productivity Center simplifies this task by listing all the cartridges for each discovered tape library. To find a particular tape cartridge:

1. Expand the **Tape Manager** node and click **Tape Libraries**.
2. Select a tape library from the list of discovered tape libraries.
3. Click **Cartridges**.
4. Click the **Label** column header to sort the list of cartridges by label. The triangle next to the label header indicates ascending or descending order.
5. Scroll through the list of cartridges to find the one you are looking for.
6. The **Location** column will indicate exactly where the tape is located in the library.

Note: You can also sort by other columns by clicking on the column headers. For example, if you know that the cartridge you are looking for is dual sided, you can sort the list by that criterion. This can make finding a particular cartridge much easier, especially for libraries that have many cartridges.

Creating a VMware VI probe

Probes collect data from the VMware VI data sources that have been discovered.

To create a VMware probe, complete the following steps:

1. In the Navigation Tree pane, expand **IBM Tivoli Storage Productivity Center** → **Monitoring** → **Probes**.
2. Right-click on **Probes** and click **Create Probes**.
3. Enter the information in the three tabbed pages:

   What to Probe
   Enter a description for the probe and from the **Available** column, select **Hypervisors** to include in the group to be probed. Click **>>** to move the group to the **Current Selections** column.

   When to Run
   Specify when the probe is to begin. You can begin the probe immediately, schedule it to begin at a later time, and specify if it should be run repeatedly.

   Alert Specify the triggering condition that will cause an alert to be generated for this probe, and the actions that should be taken when an alert is generated.
4. To save the probe, click **Save**.
5. In the **Save As** window, enter the probe name and click **OK**.

Collecting storage statistics from IBM Tivoli Storage Productivity Center servers (TPC server probes)

Define and run TPC server probes from the master server to collect storage information that has been gathered by subordinate servers. The information gathered by probes of subordinate servers is displayed in the **Tivoli Storage Productivity Center** > **Rollup Reports** node of the master server's navigation tree.

A subordinate server is a server that monitors and reports on storage resources like a standard server, but also communicates with the master server during TPC server probes. During these probes, the master server collects the storage information gathered by a subordinate server's agents and stores that information in its own database repository.

Note: If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

Before you run TPC server probes, make sure to perform the following:

1. Associate subordinate servers with the master server on the **Administrative Services** → **Data Sources** → **TPC Servers** window. See "Adding an IBM Tivoli Storage Productivity Center Server" on page 150 for information on how to associate subordinate servers with a master server.
2. Run discoveries, probes, and scans on the subordinate servers to gather information about their locally-managed entities.

The information collected by TPC server probes from subordinate servers includes the following:

- Asset information: detailed statistics about agents (Tivoli Storage Productivity Center), computers, storage subsystems, disk and volume groups, disks, filesystems, logical volumes, volumes, and fabrics.
- Database asset information: detailed statistics about the RDBMSs in your environment, including Oracle, SQL Server, Sybase, and UDB.
- Capacity information: storage metrics related to the disk capacity, filesystem capacity, filesystem used space, and filesystem freespace of the storage entities in your environment.
- Database capacity information: storage metrics related to the storage capacity of the RDBMSs in your environment, including Oracle, SQL Server, Sybase, and UDB.

Additional storage information collected by a subordinate server and used within such Tivoli Storage Productivity Center functions as the topology viewer, data path explorer, volume provisioning, volume performance, SAN Planner , etc. is available for that subordinate server only. That information is not collected by the master server and thus not rolled up into the reports available through the master server.

Also, the Tivoli Storage Productivity Center functions within the master server will use the storage information collected by the master server only.

Creating TPC server probes

Use TPC server probes from the master server to collect storage information that has been gathered by subordinate servers.

To create a TPC server probe, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** → **Monitoring** → **TPC Server Probes**.
2. Right-click **TPC Server Probes** and select **Create TPC Server Probes** from the pop-up menu.
3. Enter a brief description of the probe in the **Description** field.
4. In the **What to Probe** page perform the following steps:
   a. In the **Available** list box highlight the TPC servers you want to probe. You can do any of the following: 1. Select the **All TPC Servers** item to automatically include all subordinate servers in the probe. 2. Select the name of a specific subordinate server and all its related storage entities that you want to probe. 3. Select the specific entities on a subordinate server that you want to probe. For example, you can expand the node for a subordinate server and select the specific entities for which you want to collect data: All Resources, Clusters, Computer Groups, Computers (including Hypervisors), Database Instances, Fabrics, RDBMS Computer Groups, Storage Subsystem Groups, and Storage Subsystems.
   b. Click  to include highlighted TPC servers or storage entities in the probe. The entities that you selected are moved to **Current Selections** list box.
5. In the **When to Run** page, schedule when you want the probe to run. You have the following options:
   - Run the job immediately
   - Run the job once at a specified time and date. Use the list boxes to specify the month, day year, and the hour, minute, and AM/PM.
   - Run the job repeatedly according to a schedule you specify. You can:

- use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM when the job should begin running.
- indicate how often the job should run. Use the list boxes to specify the number of minutes, hours, days, weeks, or months.
- indicate which days on which run the job. A check mark will appear next to the days you select.

Use the fields in the **How to handle timezones** section to indicate the time zone that Data Manager for Databases should use when scheduling and running the action. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:

- 9:00 AM of the time zone where the server resides
- 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on a Data agent located in the PST time zone.

Use the timezone that the server runs in
Select this option to use the time zone of the location where the Data server resides.

Use this timezone
Select this option to indicate a specific time zone for when to run the job. When you select this option, the list box to the right of the field will activate.

6. In the **Alert** page, set the following options:
   - The condition that triggers an alert. The default setting is **Probe Failed**. No other conditions are currently available.
   - The method by which you should be notified of an alert: SNMP Trap, IBM Tivoli Enterprise Console Event, Login Notification, Windows Event Log, UNIX Syslog, running a script, and e-mail. If you select SNMP Trap, Tivoli Enterprise Console Event, or e-mail for the method by which you are notified of an alert, you must define a destination for the method on the **Administrative Services → Configuration → Alert Dispositions** window.
7. Click  to save the probe.
8. Enter a name for the TPC Server probe and click **OK**. An entry for the probe appears under the **TPC Server Probes** node.

Editing TPC Server probes

Learn how to edit the definition of a TPC server probe.

To edit an existing TPC server probe, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center → Monitoring → TPC Server Probes**.
2. Click the name of the TPC server probe you want to edit. The content pane displays information about that probe.
3. Edit the values on the What to PROBE, When to Run, and Alert pages to change the TPC server probe definition.
4. Click  to save the updated TPC server probe definition.

Managing TPC server probe jobs and runs

Learn how to manage the TPC server probe jobs and runs that appear in the **Tivoli Storage Productivity Center** > **Monitoring** > **TPC Server Probes** node of the master server's navigation tree.

When you expand **TPC Server Probes** node, the existing probe job definitions are displayed as subnodes. Note that there are no default jobs available for TPC server probes.

To create a TPC server probe job:
Right-click the **TPC Server Probes** node and select **Create TPC Server Probes**. Define information about that job on the What to PROBE, When to Run, and Alert pages that appear in the content pane. Select **File** > **Save** to name and save the probe.

To view or edit detailed information about an existing TPC server probe job:
Expand the **TPC Server Probes** node and left-click on a job. Information about that job appears in the content pane on the What to PROBE, When to Run, and Alert pages. If you edit the information in these windows, you must select **File** > **Save** to save the changes.

To view information about the runs for a TPC server probe job:
Expand a probe job. An entry appears in the navigation for every run of the probe. Click on a job run to view information about that run in the content pane.

To refresh the list of runs for a job:
Right-click the job node and select **Refresh**.

To delete a probe job and all its job runs:
Right click a job and select **Delete** from the pop-up menu.

To rename a probe job:
Right click a job and select **Rename**, and enter the new name for the job.

To refresh the list of runs for a job:
Right click a job and select **Refresh Job List**.

To update the job's status:
Right click the job and select **Update Job Status**. When you select this option, Tivoli Storage Productivity Center will update the visual indicator next to a job that indicates its status. For example, select this option to turn the status of any non-running jobs to red.

To view a job run history:
Right click a job and select **History**.

To run the job immediately:
(temporarily bypassing the schedule defined in the When to Run page)
Right click a job and select **Run Now**.

Collecting storage usage and trending information (scans)

Use scans to collect statistics about the usage and trending of storage consumption within your environment. You can view the information collected by scans in Capacity, Usage, Usage Violations, Backup, and System reports, as well as use that information as part of quota analysis.

Scans are always directed against a Data agent and deliver very detailed information about the filesystems, files, and RDBMS objects (instances, databases, devices, tablespaces, tables, indexes, datafiles, containers) of computers. Use the statistics gathered by scans to:

- view information about when storage resources are created, accessed, and modified and by what group or user
- map the actual storage resource to the consumers of that resource
- generate a historical view of storage consumption to help determine usage trends over time

You can define any number of scans to gather information about different storage resources at different times.

Creating scans

Create and schedule a scan to gather usage information about your storage resources. You can specify which file systems, computers, clusters, directories, databases, and so on. that you want to be scanned. You can define any number of scans to gather information about different storage resources at different times.

To create a scan, complete the following steps:

1. Expand **Data Manager** → **Monitoring** → **Scans**. For Data Manager for Databases, expand **Data Manager for Databases** → **Monitoring** → **Scans**.
2. Right-click the **Scans** node and select **Create Scan**.
3. For Data Manager, define the following about the scan:
   - Select the computers and file systems whose storage usage and consumption you want to monitor. See "Selecting file systems to scan" on page 212 for more information.
   - Select the directory groups to scan. See "Selecting directory groups to scan" on page 212 for more information.
   - Select the profiles to use during a scan. See "Using profiles in a scan" on page 214 for more information.
   - Schedule when you want a scan to run. See "When to run the scan" on page 215 for more information.
   - Define an alert to notify you if the scan fails. See "Defining alert values for a scan" on page 216 for more information.
4. For Data Manager for Databases, define the following about the scan:
   - Select the databases, tablespaces, instances, and computers whose storage usage and consumption you want to scan. See "Selecting databases, tablespaces to scan" on page 213 for more information.
   - Select the table groups you want to scan. See "Select the table groups to scan" on page 214 for more information.
   - Select the profiles to use during a scan. See "Using profiles in a scan" on page 214 for more information.
   - Schedule when you want a scan to run. See "When to run the scan" on page 215 for more information.
   - Specify options for Oracle and UDB. See "Specify options for Oracle and UDB" on page 216 for more information.
   - Define an alert to notify you if the scan fails. See "Defining alert values for a scan" on page 216 for more information.
5. Click the save icon on the menu bar to save the scan.

6. Enter a name for the scan and click **OK**. An entry for the scan job appears under the **Scans** node in the navigation tree.

Selecting file systems to scan:

This topic describes how to select the file systems that you want to scan.

Use the **File systems** page to select the computers and file systems whose storage usage and consumption you want to monitor. You can include clusters, individual file systems and computers, or file system and computer groups (that have been predefined using the "Working with groups of resources" on page 236 facility) against which to run a scan.

The **Current Selections** list box shows the file systems, file system groups, computer groups, computers, and clusters that have been selected for the scan. The **Available** list box shows which storage resources are still available for selection in the scan.

Note: A cluster name is displayed next to individual computers that are members of an HACMP or Microsoft Cluster Server cluster. Cluster resource groups are identified with a special icon. Clusters displayed in the **Available** list box cannot be expanded to show member nodes. You can move a cluster to the **Current Selections** list box to view and, if necessary, exclude member nodes.

To include storage resources in a scan:

1. Highlight the file systems, file system groups, computer groups, computers, and clusters you want to include in the scan in the **Available** list box.
2. Click >>. The storage resources you selected appear in the **Current Selections** list box.

To remove storage resources from a scan:

- Highlight the file systems, file system groups, computer groups, computers, and clusters that you want to remove and click <<. This will remove the resources from the **Current Selections** list box and move them to the **Available** list box.

To exclude one or more file systems or computers from a selected cluster or group:

Perform the following steps if you want to exclude specific filesystems or computers from a cluster or group that you included in a scan:

1. Double click the cluster or group's name in the **Current Selections** list box.
2. Right click on the file systems or computers that you want to exclude. A pop-up menu appears.
3. Select **Exclude**. This will create an **Excludes** section in the **Current Selections** list box that displays the names of the file systems or computers that you have chosen to exclude from a scan. This is useful when you want to select a group or cluster for scanning, but you want to exclude a few file systems or computers from that monitoring.

Selecting directory groups to scan:

This topic describes how to select the directory groups that you want to scan.

Use the Directory Groups tab to select the groups of directories that you want to scan. Including directory groups within a scan enables you to monitor and gather information about specific directories across a number of machines.

The **Selected Directory Groups** list box shows what directory groups have been included in the scan. The **Available Directory Groups** list box shows the groups that are still available for selection. You can select directory groups that have been predefined using the groups facility.

To select a directory group:

1. Highlight the desired directory group in the **Available Directory Groups** list box.
2. Click >>. The group you selected appears in the **Selected Directory Groups** list box.

To remove a directory group from a scan:

- Highlight the groups you want to remove and click <<. This will remove the groups from the **Selected Directory Groups** list box and move them to the **Available Directory Groups** list box.

To create a new directory group:

- Click **New Directory Group**. The **New Directory Group** window opens, enabling you to create a new group. See "Creating directory groups" on page 238 for more information on how to create a directory group.

Selecting databases, tablespaces to scan:

This topic describes how to select the databases and tablespaces that you want to scan.

Use the Database-Tablespaces page to select what databases, tablespaces, instances, and computers whose storage usage and consumption you want to monitor. You can include single or groups of databases, tablespaces, and computers against which to run the scan.

The **Current Selections** list box shows the databases, tablespaces, tablespace groups, computers, and computer groups that have been selected for the job. The **Available** list box shows what items are still available for selection.

To include a database, tablespace, computer, instance, or group in a scan:

1. Highlight the desired databases, tablespaces, computers, instances, or groups in the **Available** list box.
2. Click >>. The databases, tablespaces, computers, instances, or groups you selected appear in the **Current Selections** list box.

   For example, if you select a computer to include in a scan, that job will scan all the databases and/or tablespaces on all the instances that reside on that computer.

To remove a database, tablespace, computer, instance, or group from a scan:

- Highlight the items you want to remove and click  . This will remove the databases, tablespaces, computers, instances, or groups from the **Current Selections** list box and move them to the **Available** list box.

Select the table groups to scan:

Select the table groups that you want to scan.

Use this page to select the table groups you want to scan. A table group represents a group of defined tables. By choosing a table group on this window, you can perform more detailed scans of tables that might not appear within other Data Manager for Databases reports (such as the Largest Table or Monitored Tables reports).

Note: Before you can view detailed information about tables in an RDBMS, you must have 1) defined a table group that contains the tables about which you want to view storage information and 2) scheduled and run a scan against that table group.

The **Selected Table Groups** list box shows the table groups that have been selected for the job. The **Available Table Groups** list box shows the table groups that are still available for selection.

To include a table group in a scan:

1. Highlight the desired group in the **Available Table Groups** list box.
2. Click  . The group you selected appears in the **Selected Table Groups** list box.

To remove a table group from a scan:

- Highlight the groups you want to remove and click  . This will remove the groups from the **Selected Table Groups** list box and move them to the **Available Table Groups** list box.

To create a new table group:

- Click **New Table Group**. The **New Table Group** window opens, enabling you to create a new group.

Using profiles in a scan:

Select the profile to use as part of a scan. Profiles enable you to control what files are scanned and what statistics are gathered.

Use the **Profiles** page to select which profiles to use during a scan. Using profiles is a powerful method for determining the statistics you want to gather during scans. Keep in mind that the more profiles you include within a scan, the longer the scan will take to gather the statistics specified by those profiles.

To include profiles in a scan:

1. Highlight the desired profiles in the **Available Profiles** list box.

2. Click the top >> to apply the selected profiles to file systems; click the bottom >> to apply the selected profile(s) to directories. The profiles you selected appear in the either the **Profiles to apply to Filesystems** or **Profiles to apply to Directories** list box

   Note: To collect statistics according to Data Manager default settings, select all of the default profiles.

To remove profiles from a scan:

- Highlight the profiles you want to remove and click the corresponding <<. This will remove the profile from the list boxes on the right and move them to the **Available Profiles** list box.

To create a new profile:

- Click **New Profile**. The **Statistics** tab on the **Create Profile** window is displayed, enabling you to create a new profile. See "Using profiles" on page 218 for more information.

Related concepts

"Using profiles" on page 218
Use profiles to specify what statistical information is gathered and to fine tune and control what files are scanned during a scan. You can also use profiles with some reports to determine what statistics are displayed.

When to run the scan:

Schedule when and how often to run a scan.

You have the following options for determining when to run a scan:

- Run the job immediately
- Run the job once at a specified time and date. Use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM.
- Run the job repeatedly according to a schedule you specify. You can:
  - use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM when the job should begin running.
  - indicate how often the job should run. You can use the list boxes to specify the number of minute(s), hour(s), day(s), week(s), or month(s).
  - select which days you wish to run the job.

Use the fields in the **How to handle timezones** section to indicate the time zone that Data Manager should use when scheduling and running the scan. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:

- 9:00 AM of the time zone where the server resides
- 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

Use the timezone that the server runs in. Select this option to use the time zone of the location where the Data server resides.

Use this timezone. Select this option to indicate a specific time zone for when to run the job. When you select this option, the list box to the right of the field will activate.

Specify options for Oracle and UDB:

Specify options for Oracle and UDB when defining a scan.

Use the Options page to do the following:

- Determine if the job should scan the freelist chains associated with the segments in your Oracle environment. Free list chains represent the pages allocated to a segment that have enough free space for additional rows. Note that checking this option will increase the time required by the scan to complete.
- Set the number of pages Data Manager for Databases should scan on a chain before stopping.
- Determine if the job should gather segment extent data for locally-managed tablespaces.
- Run the DB2 utility runstats against the tables being scanned. Checking this box will enable Data Manager for Databases to collect the most up-to-date storage statistics about tables from the DB2 catalog. Note that checking this option will increase the time required by the scan to complete, but will ensure that table and user statistics are up-to-date in UDB/DB2 databases. Alerts based on table and user stats will not function until runstats is run.

Defining alert values for a scan:

This topic describes how to set an alert that will be triggered if the scan job fails.

Use the Alert page to can define an alert to notify you if a scan fails and how to notify you. You can choose to be notified by SNMP trap, IBM Tivoli Enterprise Console event, login notification, Windows event log, UNIX Syslog, running a script, and e-mail. If you select SNMP trap, Tivoli Enterprise Console event, or e-mail, you must also define a destination (for example, an e-mail address). Define this destination by expanding **Administrative Services** → **Configuration** → **Alert Dispositions**.

To define an alert for a scan, complete the following steps:

1. Click the **Alert** tab; the Alert page opens.
2. In the **Triggering Condition** area, select Scan Failed. This indicates that an alert will be triggered if the scan fails to run.
3. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

   SNMP trap
   An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

   TEC event
   A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

Login notification
The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who will receive alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who will receive alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that will be recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog will be sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

4. Click **File** > **Save** to save the scan.
5. Enter a name for the scan and click **OK**. An entry for the scan appears under the **Scans** node.

Editing scans

Change the settings in an existing scan job.

To edit an existing scan, complete the following steps:

1. Expand **Data Manager** → **Monitoring** → **Scans**. For Data Manager for Databases, expand **Data Manager for Databases** → **Monitoring** → **Scans**.
2. Click the name of the scan you want to edit. The content pane displays information about that job.
3. Edit the scan definition by changing the values that appear on the following pages:
   - Filesystems. Use this page to edit the computers and file systems whose storage usage and consumption you want to monitor as part of a scan.
   - Directory Groups. Use this page to edit the groups of directories that you want to scan.
   - Databases-Tablespaces. Use this page to edit the databases, tablespaces, instances, and computers whose storage usage and consumption you want to monitor as part of a scan.
   - Table Groups. Use this page to edit the table groups that you want to scan. A table group represents a group of defined tables.

- Profiles. Use this page to edit which profiles are used during a scan. Using profiles is a powerful method for determining the statistics you want to gather during scans.
- When to Run. Use this page to edit when a scan is run.
- Options (Oracle and UDB only). Use this page to edit the scan options for Oracle or UDB.
- Alert. Use this page to edit how you are notified if a scan fails.

4. Click  to save the updated scan definition.

Scanning NetWare volumes

Use scans to gather detailed information about the NetWare volumes within your environment.

However, before you can gather information about NetWare volumes, you must:

1. Configure Data Manager to discover all the NetWare servers in your NDS trees and then run a discovery job. See the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for information about configuring and running discovery against NDS trees.
2. Access the **Administrative Services → Configuration → Scan/Probe Agent Administration** window to assign agents to each of the volumes discovered on your licensed NetWare servers. You can assign any agent to a volume that has the ability to access that volume.

Field Name	Description
Server	The servers discovered in the NDS trees.
Domain	The domain where a server resides.
Tree Name	The names of NDS trees discovered by agents.
Filesystem	The file systems/volumes discovered in the NetWare servers.
Scan/Probe Agent	Agents assigned to the file systems/volumes in the NetWare servers.

   Note: If you do not assign an agent to a volume, Data Manager will not scan and gather information about that volume.
3. Define a scan and include the NetWare volumes and filesystems for which you want to gather detailed storage information.

Using profiles

Use profiles to specify what statistical information is gathered and to fine tune and control what files are scanned during a scan. You can also use profiles with some reports to determine what statistics are displayed.

Until you begin creating your own profiles, you can use the default profiles that are supplied with IBM Tivoli Storage Productivity Center. Each of these default profiles allows you to select a specific statistic that you want to gather. If you want to scan a computer or filesystem, and want to gather all of the statistical information available, simply select all of the default profiles when you create your scan job.

See "Default profiles" on page 219 for more information about default profiles.

Default profiles:

Default Profiles are provided for your use with IBM Tivoli Storage Productivity Center. Each default profile represents a specific statistic that you can gather (which is included in the name of the profile for convenient identification).

If you want to scan a computer or filesystem and gather all of the statistical information available, select all of the default profiles when you create your scan. TPCUser is shown as the creator of the default profiles, and TPCUser is the prefix name for each of the profiles. The following table identifies the default profiles and describes their functions.

Default profile name	Description
	Data Manager
By Access	Gathers statistics by length of time since last access of files.
By Creation	Gather storage statistics by the length of time since the creation of files. This profile applies to scans run against Windows computers only. Scans with this profile do not gather any creation time data from non-Windows computers.
By Mod Not Backed Up	Gathers statistics by length of time since last modification (only for files not backed up since modification). (Windows only)
By Modification	Gathers statistics by length of time since last modification of files.
Size Distribution	Gathers information on the size distribution of files.
Largest Directories	Gathers statistics on the *n* largest directories. (20 is the default amount.)
Largest Files	Gathers statistics on the *n* largest files. (20 is the default amount.)
Largest Orphans	Gathers statistics on the *n* largest orphan files. (20 is the default amount.)
Most At Risk	Gathers statistics on the *n* files that have been modified the longest time ago and have not yet been backed up since they were modified. (20 is the default amount.) (Windows only)
Oldest Orphans	Gathers statistics on the *n* oldest orphan files. (20 is the default amount.)
Most Obsolete Files	Gathers statistics on the *n* "most obsolete" files (i.e., files that have not been accessed or modified for the longest period of time). (20 is the default amount.)
Summary By File Type	Summarizes space usage by file types (for example, .exe., .dll, .doc, .mp3, etc.).
Summary By Filesystem/Directory	Summarizes space usage by file system or directory.
Summary By Group	Summarizes space usage by OS Group.
Summary By Owner	Summarizes space usage by Owner.
Temporary Files	Gathers statistics on the non-OS files not accessed in the last year and orphan files.
Wasted Space	Gathers statistics on the non-OS files not accessed in the last year and orphaned files.

Default profile name	Description
	Data Manager for Databases
DB User Space	Gathers statistics about the user space usage within a database or tablespace.
Largest Tables	Gathers statistics on the *n* largest tables. (20 is the default amount)
Segment Most Extents	Gathers statistics on the *n* segments with the most extents. (20 is the default amount)
Segment Most Unused Space	Gathers statistics on the *n* segments with the most empty used space. (20 is the default amount). Empty used space represents the space allocated to a segment that is currently empty/not being used. Gathering this information can help you discover space that can be reclaimed and allocated to other objects. Specifically, use this Profile to: • find objects that are over-allocated (not using all their associated free space) • find the free space that exists above a table's ″high-water″ mark (the ″high water mark″ is the highest block that was written for a table) • find segments that have the largest amount of blocks that are not filled
Summary	Summarizes space usage by database or tablespace, such as how many segments are in a tablespace (Oracle only), how many of those segments are tables (Oracle only), indexes, etc.

Creating a profile:

Create a profile for use within a scan.

To create a profile perform the following steps:

1. Expand **Data Manager** → **Monitoring**. For Data Manager for Databases, expand **Data Manager for Databases** → **Monitoring**.
2. Perform one of the following actions:
   - Select **Create Profile** from the pop-up menu.
   - Click **New Profile** from the **Profiles** tab of the Create scan window.

   Use the Create Profile window to specify exactly what level of detail to collect when running a scan. The more information that you specify, the more reports that you will be able to run against the information that you collect during the scan. Statistics are tied to a specific profile.
3. The **Statistics** tab on the **Create Profile** window is displayed.
4. Select statistics you want the scan to collect. Statistics are tied to a specific profile.

   See "Default profiles" on page 219 for more information about statistics.
5. If you want to further limit specific files to be scanned, select the **File Filters** tab. You can define exclusive conditions that filter what files are analyzed during the scan.

   See the "Creating file filters" on page 223 section for more information on using this tab.

6. The first time you view a new profile, this window is blank except for the text **All files selected**. That is the default and indicates that information will be collected for all files. To create your first group or condition, simply right-click **All files selected** and choose to create a new group or condition.
7. In the **File Filters** page, the default condition **All files selected** is displayed. Right-click **All files selected** and select **New Condition** from the pop-up menu. You can also specify groups for these conditions, by selecting **New Group**.

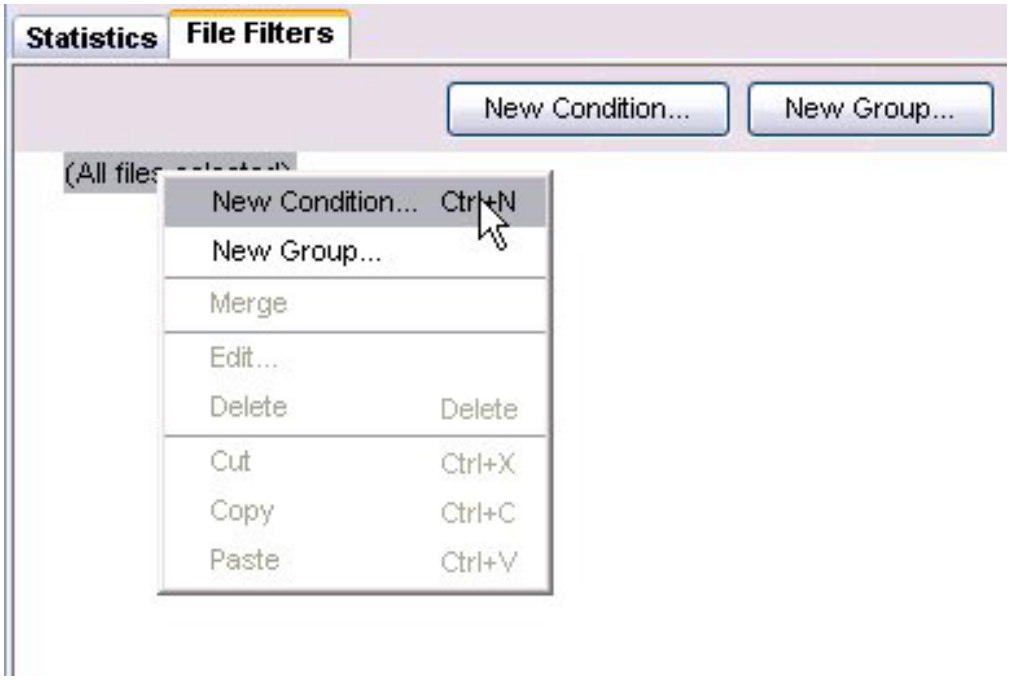

Figure 17. File Filters tab, pop-up menu

8. In the figure above, **New Condition** is selected. The **Create Condition** window opens, enabling you to define your file conditions. You can create conditions based on access time, creation time, file name or group, modification time, file attributes, size, type, owner, or length.

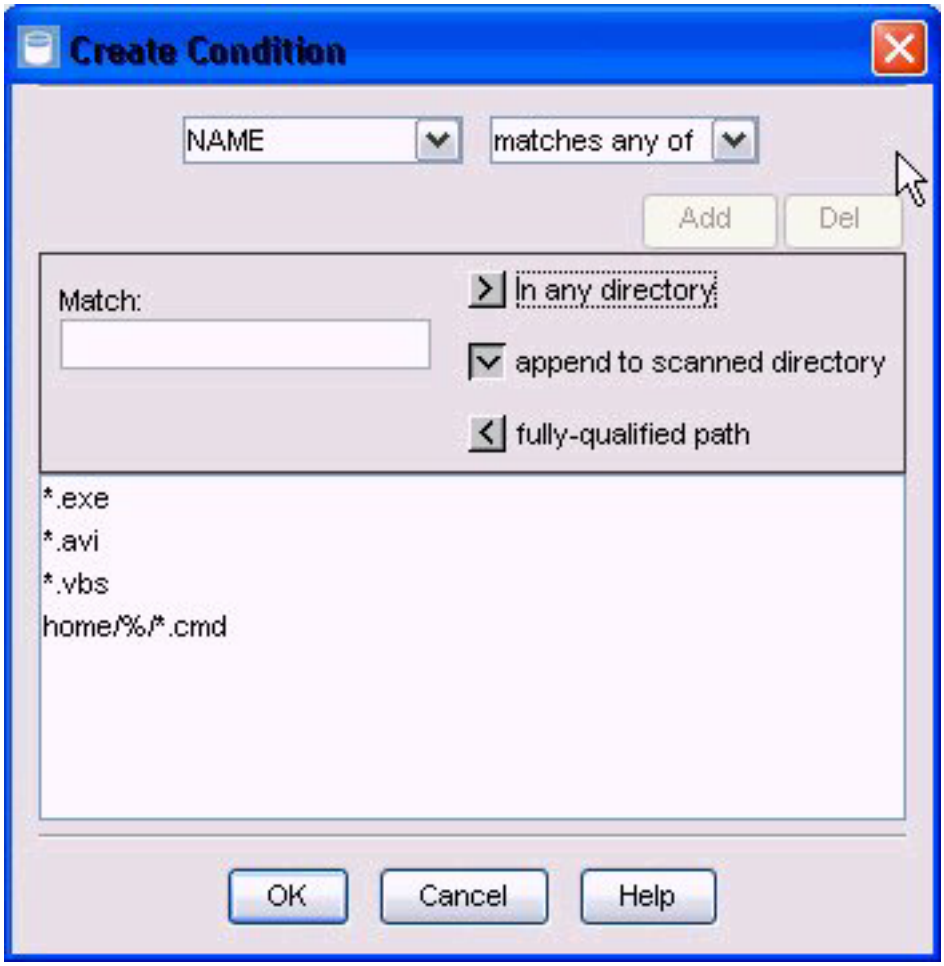

Figure 18. Create Condition window

9. Press **OK** after defining each condition.
10. If you are defining more than one condition, the **Create Group** dialog box displays to help you associate the conditions you created:

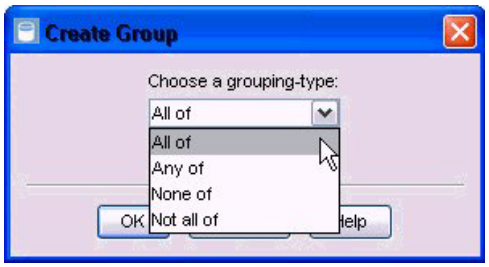

Figure 19. Create Group Window

11. Select how you want to associate the conditions you have defined. See "Associating conditions with groups" on page 228 for more information on creating groups.
12. The condition is displayed in the **File Filter** window.

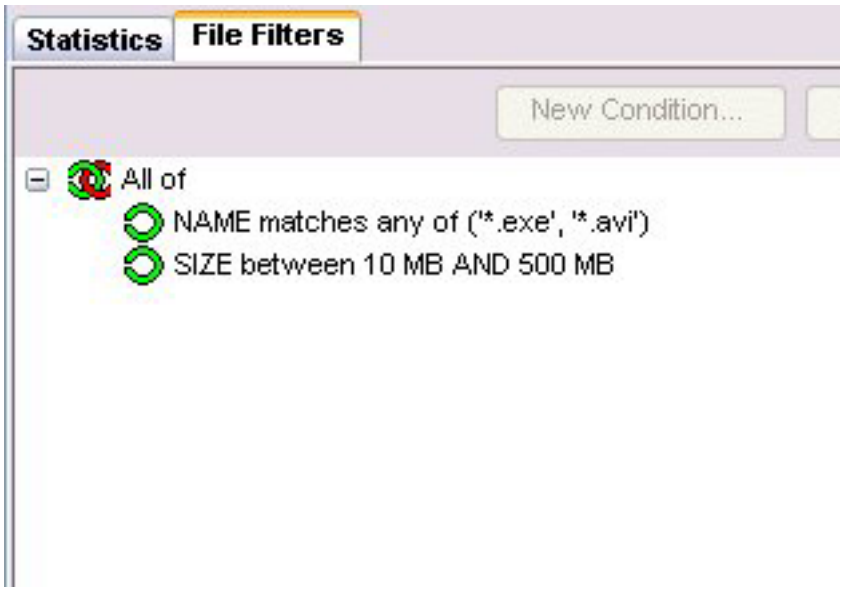

Figure 20. File Filter window

13. Click  to save the profile.
14. Enter a name for the profile and click **OK**. An entry for the profile appears as a node under the **Profiles** node.

Using file filters:

Use file filters as part of a profile to further define the files that you want to include or exclude in a scan. Filters limit the files on which you are gathering data during the scan based on criteria that you define.

Creating file filters:

Use file filters within profiles to build a complex where clause that will be used to filter what files are analyzed in the collection of data during a scan.

In the following example, a group (`All of`) has been created with three conditions:

- The first controls the type of files to collect information on (MP3, GIF, AVI, and JPG files).
- The second specifies that the files must have been created within the last year.
- The third limits the profile to analyzing only files greater than 500 KB.

To create file filters within profile, complete the following steps:

1. To edit the filer filters in an existing profile, expand **Data Manager** → **Monitoring** → **Profiles** and select an existing profile in the navigation tree. To create a new profile and add file filters, expand **Data Manager** → **Monitoring**, right-click **Profiles**, and select **Create Profile**.
2. Click the **File Filters** tab.
3. The first time you view a new profile, this window is blank except for **All files selected**. Selecting all files is the default and indicates that information is collected for all files. To create your first group or condition, right-click **All files selected** and choose to create a new group or condition.

4. In the File Filters page, the default condition **All files selected** is displayed. Right-click **All files selected** and select **New Condition**. You can also specify groups for these conditions by selecting **New Group**.

   The **Create Condition** window opens, enabling you to define your file conditions. You can create conditions based on access time, creation time, file name or group, modification time, file attributes, size, type, owner, or length.
5. Press **OK** after defining each condition. See "Defining a filter condition" for more information on creating groups.

   Each condition is displayed in the **File Filter** window.

   If you are defining more than one condition, the **Create Group** dialog box helps you associate the conditions you create.
6. Associate the conditions you have defined.

   See "Associating conditions with groups" on page 228 for more information on creating groups.

Defining a filter condition:

Define a condition for a filter.

On the File Filters page, right-click **All files selected** and select **New Condition**. The Create Condition window opens.

The first drop down list displays file criteria that you can select to define your condition. You can define file criteria based on access time, creation time, file name or group, modification time, file attributes, size, type, owner, or length. Depending on the criteria you choose to define, the second drop down list and window fields change dynamically to help you set parameters for the selected criteria. The following table describes the file criteria you can select for your file filter.

Condition	Description	Method for filtering the files	Condition Modifiers
Last Accessed	Use this condition to limit files based on the last time they were accessed. For example, you can define a filter that includes files that were last accessed between 12:01 on January 1, 2005 and 11:59 on January 30, 2005.	• later than • earlier than • between • not between	• Length of time before scan: enter the years, weeks, days, hours, or minutes before a scan • Timestamp - select files based on the timestamp
Created	Use this condition to limit files based on when they were created. For example, you can define a filter that includes files that were created between 12:01 on January 1, 2005 and 11:59 on January 30, 2005.	• later than • earlier than • between • not between	• Length of time before scan: enter the years, weeks, days, hours, or minutes before a scan • Timestamp: select files based on the timestamp

Condition	Description	Method for filtering the files	Condition Modifiers
Name	Use this condition to limit files based on their name. For example, you can define a filters that include files that end in *.html, or include files that end in .cmd and appear within a specified subdirectory.	• matches any of • matches none of • matches • does not match	Enter the file name upon which you want the condition to be based. For example, you can specify, *.exe, *.mp3, *.avi. in the **Match** field. See "Using wildcards in file filters" on page 229 for information about how to use wildcards (*, %, etc.) in this field. Select the location where the file name will be searched during the scan: • **In any directory**. During a scan, Data Manager will look for the file name in any directory. In this condition, the search string will occur at the end of the full directory path. • **Append to scanned directory**. This option enables you to restrict what directory to scan. Data Manager will look for the file type in the specified file system root directory or in a monitored directory in a directory group tied to a scan. In this option, you must first specify what directory groups you want scanned in the **Directory Groups** tab of the **Scan** window. See "Selecting directory groups to scan" on page 212 for more information. • **Fully-qualified path**. Enter the fully-qualified path when you enter the file type in the **Match** field. During a scan, Data Manager will look for the file type in the exact path you specify here. This is useful when you want to exclude system directories and files in a Most Obsolete Files report.
Group	Use this condition to limit files based on the OS group to which those files belong.	• matches any of • matches none of • between • not between • matches • does not match • >= • < • > • <=	Name of the OS group
Last Modified	Use this condition to limit files based on when they were last modified.	• later than • earlier than • between • not between	Specify the time as: • Length of time before scan: enter the years, weeks, days, hours, or minutes before a scan • Timestamp: select files based on the timestamp

Condition	Description	Method for filtering the files	Condition Modifiers
Attributes	Use this condition to limit files based on file attributes such as read only, system, temporary, etc.	• include any of • include none of • include all of • don't include all of	Select the file attribute upon which you want the condition to be based. You can select: • Read only: Users cannot write to the file. • Hidden: A file that cannot be seen by users during a normal directory listing. • System: A file used by the operating system that cannot normally be viewed by users. • Archive: A file that has been backed-up and the archive bit is turned off. An Archive bit signifies that the file has been changed since it was last archived. • Normal: A file with no special attributes set to it. • Temporary: (Windows only) A temporary or backup copy of a file. • Sparse: A file with a special kind of compression that does not take up a lot of CPU time. The logical size of the file looks bigger than is actually allocated. It contains some blank pages which are not on disk. • Reparse Point: New NTFS file system objects that have a definable attribute containing user-controlled data and are used to extend functionality in the input/output (I/O) Subsystem. • Compressed: A file to which a compression algorithm was applied. This file takes up less space on disk, but it will be slower to read. For Novell NetWare, the file itself is not usually compressed, rather the whole volume is compressed. • Offline: A file that was migrated to tape. • Not Content Indexed: A file that is not indexed--the Windows indexing service was told that this is not a file that should be searched through for text strings. • Encrypted: A file that was saved with a password (is password-protected). • Orphaned: (for Windows, UNIX, Linux, and NetWare) The user who owns the file is no longer present in the system database.

Condition	Description	Method for filtering the files	Condition Modifiers
Attributes (continued)			• Block Device: (UNIX/Linux only) A disk device file (not a real file). Executable: A file that you can run or execute. Sticky: A special type of executable file that is read into memory when it is first run and never released. Use Scans to find these types of files and set them to ″unsticky″. SETGID: (UNIX/Linux only, for executable files) When a file is run, its group changes to the group of the file. For example, if a user belongs to User group and runs a file from a different group (e.g., Root), than the user's group will be changed to Root while running the file (for security purposes). SETUID: (UNIX/Linux only, for executable files) When a file is run, the user's effective user will be changed when the file is run. For example, if a file is SETUID =Root, a users effective user ID becomes Root when running the file. This is considered a security hole by UNIX Administrators, so you can define a Profile to search for files that have a SETUID = Root.
Size	Use this condition to limit files based on their physical file size.	• between • not between • >= • < • > • <=	• KB (kilobytes) • MB (megabytes) • GB (gigabytes) • TB (terabytes)
Type	Use this condition to limit files based on pre-defined file types.	• in • not in • = • < • >	• Link • Directory • Fifo • Socket • Device
Owner	Use this condition to limit files based on their owner.	• matches any of • matches none of • between • not between • matches • does not match • >= • < • > • <=	Enter the name of the owner upon which you want the condition to be based in the **Entry** field.

Condition	Description	Method for filtering the files	Condition Modifiers
Length	Use this condition to limit files based on the logical size of the files contents.	• between • not between • >= • < • > • <=	Enter the file length range. You can qualify both values by selecting from one of the following: • KB (kilobytes) • MB (megabytes) • GB (gigabytes) • TB (terabytes)

Associating conditions with groups:

Associate multiple conditions in a group or groups.

Groups can further define your file filter by including or excluding certain files based on specific criteria.

1. Create a new group using one of the following options.
   - In the **File Filters** tab, right-click a condition and select **New Group** from the pop-up menu.

     The **Create Group** dialog box is displayed.
   - Create each file condition for your file filter.

     As you create a new file condition, the **Create Group** dialog box is displayed.
2. Select the grouping for the file conditions you created. The grouped condition is displayed in the **File Filters** tab.

Select one of the following groups.

Group	Description
All of	All inclusive. All the conditions must be met. For example: • a *and* b *and* c
Any of	Any one or more of the following conditions will be met. For example: • a *or* b *or* c
None of	None of the conditions are met. For example: • *not* (a *or* b *or* c)
Not all of	All of the following conditions are not met. For example: • *not* (a *and* b *and* c)

For example, say you want to create a file filter with the following conditions:

- Name in (*.gif, *.jpg)
- Size > 1 MB
- Attribute is Read Only

How you group the conditions produces different results. The following table shows the different results depending on the group you select:

File	Group	Result
A.gif 500KB read only	• All of • Any of • None of • Not all of	• False • True • False • True
B.jpg 2 MB read only	• All of • Any of • None of • Not all of	• True • True • False • False
C.gif 2 MB read/write	• All of • Any of • None of • Not all of	• False • True • False • True
D.txt 500KB read/write	• All of • Any of • None of • Not all of	• False • False • True • True

Using wildcards in file filters:

Use wildcards within the conditions of file filters to perform pattern matching of file names. A wildcard is a special symbol that represents one or more characters.

You can use wildcards in File Filter conditions to identify multiple files and directories with a single specification. The following table describes some of the wildcards supported by Data Manager:

Wildcard Character	Description	Examples
*	Use the asterisk (*) to stand-in for unknown characters in file names. It can act as a wildcard for zero or more characters of any value within a directory or file name.	• **s*** refers to all files that begin with **s** • ***.exe** refers to all files that have **.exe** as an extension • **/user/*** refers to all files under the **/user** directory
%	Use the percent (%) character to specify a wildcard for a directories within the path of a file. It can act as a wildcard for zero or more directories within a path.	• **/home/%/*.cmd** will match all files with the extension **.cmd** and located in any of the subdirectories under the **/home** directory • **/usr/%/*** will match all files and subdirectories under the **/usr** directory (for example, /usr/files/c, /usr/c, /usr/files/documents/c)

Wildcard Character	Description	Examples
?	Use the question mark (?) to stand in for a character within a file name or directory name. It can as a wildcard for a character within a directory or file name.	• **/home/usr?/** will match directories like: /home/usr1/, /home/usr2/, etc. • **/usr/temp/temp??.tmp** will match files like /usr/temp/temp01.tmp, /usr/temp/temp02.tmp, etc.

Data Manager supports standard UNIX shell path name expansion wildcarding within file filter conditions. Please refer to the appropriate UNIX man pages for more information about the wildcards you can use.

Determining the availability of storage resources (pings)

Use pings to collect information about the availability of the storage resources in your environment. You can view the information collected by pings in Availability and System reports.

Pings enable you to monitor and report on the availability of your storage from a network point of view or from a computer uptime perspective. See the percentage of off-network time due to network problems or system downtime. Define any number of pings that each check the availability of a unique storage resource.

Creating pings

Create pings to collect information about the availability of the storage resources in your environment. You can define any number of pings to gather information about different storage resources at different times.

To create a ping, complete the following steps:

1. Expand **Data Manager** → **Monitoring** → **Pings**.
2. Right-click **Pings** A pop-up menu appears.
3. Select **Create Ping**.
4. Select the computers whose availability you want to check. “Select computers to ping.”
5. Schedule when you want a ping to run. “When to run the ping” on page 231.
6. Specify how often you want the server to write out the ping results to the database repository for viewing in reports. “Specify ping options” on page 232.
7. Click  to save ping.
8. Enter a name for the ping and click **OK**. An entry for the ping appears under the **Pings** node.

Select computers to ping:

Select the computers and clusters that you want to ping.

Use the **Computers** page to select the computers and clusters whose availability statistics you want to collect. The **Available** list box shows the computers, computer groups, or clusters that are currently being monitored and available for inclusion in the ping. After you include an object in the ping it is removed from this list and moved to the **Current Selections** list box. The **Current Selections** list box shows the computers, computer groups, or clusters that have been selected for the ping.

Note: A cluster name is displayed next to individual computers that are members of a HACMP or Microsoft Cluster Server cluster. Cluster resource groups are identified with a special icon. Clusters displayed in the **Available** list box cannot be expanded to show member nodes. You can move a cluster to the **Current Selections** list box to view and, if necessary, exclude member nodes.

You can include individual computers, computer groups, or clusters against which to run a ping. See "Working with groups of resources" on page 236 for information about groups.

- **To include storage resources in a ping:**
  1. Highlight the desired computers, computer groups, and clusters in the **Available** list box.
  2. Click  . The resource you selected appear in the **Current Selections** list box.
- **To remove storage resources from a ping:**
  - Highlight the computers, computer groups, and clusters that you want to remove and click  . This will remove the resource from the **Current Selections** list box and move it to the **Available** list box.
- **To exclude one or more computers from a selected computer group or cluster:**
  1. Double click the group or cluster's name in the **Current Selections** list box.
  2. Right click on the computers that you want to exclude. A pop-up menu appears.
  3. Select **Exclude**. This will create an **Excludes** section in the **Current Selections** list box that displays the names of the computers that you have chosen to exclude from the ping.

When to run the ping:

Schedule when and how often to run a ping.

You have the following options for determining when to run a ping:

- Run the job immediately
- Run the job once at a specified time and date. Use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM.
- Run the job repeatedly according to a schedule you specify. You can do the following:
  - Use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM when the job should begin running.
  - Indicate how often the job should run. Use the list boxes to specify the number of minutes, hours, days, weeks, or months.
  - Select which days you wish to run the job.

Use the fields in the **How to handle timezones** section to indicate the time zone that Data Manager should use when scheduling and running the ping. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:

- 9:00 AM of the time zone where the server resides
- 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

Use the timezone that the server runs in. Select this option to use the time zone of the location where the Data server resides.

Use this timezone. Select this option to indicate a specific time zone for when to run the job. When you select this option, the list box to the right of the field will activate.

Specify ping options:

Determine how often the results of pings are written from memory and stored in the database repository.

For greater efficiency, the server temporarily stores ping results in memory. The unit of time you specify in the **Options** tab determines how often these results are externalized to the repository.

You can specify a unit of time in a number of pings, minutes, hours or days.

Define alert values for a ping:

Set an alert that will be triggered if the ping job fails.

Use the Alert page to can define an alert to notify you if a computer that a ping is trying to check is unreachable for a certain amount of time. You can choose to be notified by SNMP trap, IBM Tivoli Enterprise Console event, login notification, Windows event log, UNIX Syslog, running a script, and e-mail. If you select SNMP trap, Tivoli Enterprise Console event, or e-mail, you must also define a destination (for example, an e-mail address). Define this destination by expanding **Administrative Services** → **Configuration** → **Alert Dispositions**.

To define an alert for a ping, complete the following steps:

1. Click the **Alerts** tab; the Alert page opens.
2. Define the conditions that will trigger the alert. The default setting is **Computer Is Unreachable More Than**. No other alerting conditions are currently available. Enter the **Value** and **Value Units**. These values specify how many pings occur before an error is generated. By default, **Value Units** is set to **Times**.
3. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

   SNMP trap
   An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

   TEC event
   A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

   Login notification
   The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who will receive alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who will receive alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that will be recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog will be sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

4. Click **File** > **Save** to save the ping.
5. Enter a name for the ping and click **OK**. An entry for the ping appears under the **Pings** node.

Editing pings

Change the settings in an existing ping job.

To edit an existing ping, complete the following steps:

1. Expand **Data Manager** → **Monitoring** → **Pings**.
2. Click the name of the ping you want to edit. The content pane displays information about that job.
3. Edit the ping definition by changing the values that appear on the following pages:
   - Computers. Use this page to edit the computers, computer groups, and clusters whose availability you want to monitor as part of a ping.
   - When to Run. Use this page to edit when a ping is run.
   - Options Use this page to edit how often the results of pings are written from memory and stored in the database repository.
   - Alert. Use this page to edit how long a storage entity is unreachable before you are notified that the ping has failed.
4. Click  to save the updated ping definition.

Data collection groups

Groups enable you to define a set of related resources such as a computer, fabric, or storage subsystem once and save these selections as a named group for future use in scans, pings, and probes, making data collection more efficient. For example,

you can create a group that contains a number of related computers. You can then reference that group within a monitoring job when you want to collect data about all those computers at the same time.

Each type of group is represented by a node under the **Monitoring → Groups** and **IBM Tivoli Storage Productivity Center → Monitoring → Probes** sections. You can expand the tree for each type of group to see what groups exist for that type. Resources can belong to only one group at a time. For example, if a computer already belongs to a computer group and you attempt to add it to a different computer group, the system removes it from the first computer group before adding it to the new computer group (this includes removing it from the default group if it had never been previously assigned to a user defined group). This is to prevent a resource from being counted twice statistically.

Additionally, you can use Data Manager for Chargeback to create invoices for these groups. When implementing your chargeback system, it is recommended that you group the IT resources within your organization into functional groups. Accurately identifying the types and locations of your stored data and resources is the first step in implementing an effective and equitable chargeback policy. Defining groups with Data Manager enables you to classify your data and resources into billable entities, which in turn provides information on who is using what storage. This enables you to bill specific groups according to their usage of storage space.

The following table lists and describes the types of groups you can create.

Group Type	Description
	Data Manager
Computer	A logical grouping of related computers. A computer can belong to one computer group only. A computer will automatically be removed from a group if you add it to another group. You can use computer groups to monitor computers that are related to each other or belong to a common group such as a department, location, territory, and so on. **Examples of computer groups:** Sales, Accounting, Development, IT, West, East, Midwest, NAS Filers, Remote Servers
Filesystem	A grouping of file systems across multiple computers. A file system can belong to one file system group only. A file system will automatically be removed from a group if you add it to another group. **Examples of file system groups:** `C:\`, `D:\`, `/boot`, `/home`, `/usr`

Group Type	Description
Directory	A grouping of directories across multiple computers. You can include a directory in multiple directory groups. **Examples of directory groups:** `WindowsTempDir` Included Computer Name:Directory - `a:\machine1\winnt\temp, b:\machine2\winnt\temp, c:\machine3\winnt\temp` **Note:** To monitor directories from the IBM Tivoli Storage SAN File Systems, you must create a directory group and add the desired directories as follows: */<SAN File System cluster name>/<directory name>* where *<SAN File System cluster name>* represents the name of the SAN File System cluster that contains the directory that you want to add and *<directory name>* represents the name of the directory you want to add (for example, `/IBM/myFilesetName/myDirectoryName`, `/IBM/myDirectoryName`.
User	A grouping of related users. A user can belong to one user group only. A user will automatically be removed from a group if you add it to another group. Users that were discovered by Tivoli Storage Productivity Center can be included in a user group. **Note:** The user groups that you define in Data Manager for Databases will also available in Data Manager, and vice versa. **Examples of user groups:** Administrators, Developers, Marketing staff, HR staff
OS User Group	A grouping of OS user groups that were defined outside of Tivoli Storage Productivity Center at the operating system level. OS user groups are the operating system groups that Tivoli Storage Productivity Center finds that own files and directories in the environment. **Examples of OS user groups:** Windows, Linux, UNIX
	Data Manager for Databases
Computer	A logical grouping of related computers. A computer can belong to one computer group only. A computer will automatically be removed from a group if you add it to another group. You can use computer groups to monitor computers that are related to each other or belong to a common group such as a department, location, territory, and so on. **Examples of computer groups:** Sales, Accounting, Development, IT, West, East, Midwest, NAS Filers, Remote Servers
Databases-Tablespaces	A grouping of RDBMS databases or tablespaces. You can use database-tablespace groups to monitor entire databases or tablespaces (depending on the RDBMS) across multiple instances. A database or tablespace will automatically be removed from a group if you add it to another group. A database or tablespace can belong to one database-tablespace group only. **Examples of database-tablespace groups:** DB2 databases, system tablespaces
Table	A grouping of related RDBMS tables.

Group Type	Description
User	A grouping of related users. A user can belong to one user group only. A user will automatically be removed from a group if you add it to another group. Users that were discovered by Tivoli Storage Productivity Center can be included in a user group. **Note:** The user groups that you define in Data Manager for Databases will also available in Data Manager, and vice versa. **Examples of user groups:** Administrators, Developers, Marketing staff, HR staff
	Disk Manager
Storage Subsystem	A storage subsystem group can contain subsystems and SAN Volume Controllers.
	Fabric Manager
Fabric	A grouping of fabrics.
	Tape Manager
Tape Libraries	A grouping of tape libraries.

Default groups

Default groups for each storage resource are also provided within each group type node. As IBM Tivoli Storage Productivity Center discovers these resources within your enterprise through discovery, probes, and scans, the resources are automatically added to the corresponding default groups.

The following table lists the names of the default groups for each group type.

Manager	Group Type	Default Group Name
Data Manager	Computer	TPCUser.Default Computer Group
	Filesystem	TPCUser.Default FS Group
	Directory	n/a
	User	TPCUser.Default User Group
	OS User Group	TPCUser.Default OSGroup Group
Data Manager - Databases	Computer	TPCUser.Default Computer Group
	Databases-Tablespaces	TPCUser.Default Tablespace Group
	Table	n/a
	User	TPCUser.Default User Group
Disk Manager	Storage Subsystems	TPCUser.Default Storage Subsystem Group
Fabric Manager	Fabrics	TPCUser.Default Fabric Group
Tape Manager	Tape Library	TPCUser.Default Tape Group

Working with groups of resources

Create resource groups to logically arrange resources and to permit actions against the entire group.

Creating a group

You can create named groups that represent a set of resources, such as computers, databases, table spaces, tables, and users, to later use in pings, scans, and probes.

To create a new group, perform the following steps.

1. Expand **Monitoring** → **Groups**. For example, **Data Manager** → **Monitoring** → **Groups**.
2. Right-click the type of group you want to create and select **Create** ***group type***, where ***group type*** represents the group you want to create, such as a **Create Computer Group**, **Create Filesystem Group**, or **Create Table Group**.
3. In the Create *group type* Groups page, select the resources in the **Available** column that you want to include in the group. Hold down the Ctrl key as you click to select multiple resources.

   Note: For computer and file system groups in Data Manager, a **Group By Cluster** check box is displayed. When this option is not selected, the **Available** pane displays a cluster name next to any clustered computers and cluster resource groups. When this option is selected, any computer or cluster resource group that belongs to an HACMP or MSCS cluster is grouped under its cluster name. Cluster resource groups are identified with a special icon along with the name of the current host node. All other computers will be grouped under **Non-Clustered Computers**.
4. Click >> to move the selected resources to the **Current Selections** column on the right. Click << to move resources from the **Current Selections** column and move them to the **Available** column.
5. Optional: Type a description of the group in the **Description** field.
6. Click the save icon in the menu bar to save the group.
7. Type a descriptive name for the group and click **OK**.

Creating tape library groups

Create groups of tape libraries that have similar attributes to make it easier to perform storage management tasks on the grouped libraries.

1. Expand the **Tape Manager** → **Monitoring** → **Groups** node and right-click **Tape Library**.
2. Select **Create Tape Library Groups**.
3. In the Create Tape Library Groups page, select the libraries in the **Available** column to include the libraries in the group.

   Hold down **Ctrl** to click multiple libraries.
4. Click >> to move selected libraries to the **Current Selections** column on the right.
5. Optional: Enter a description of the tape library group in the **Description** field.
6. Click  to type a descriptive name for the group in the **Save As** dialog. Then click **OK** to save the group.

The new tape library group is listed under the **Tape Library** node in the navigation tree.

Grouping fabrics

You can group fabrics to run reporting jobs against all the fabrics in the group. A fabric can belong to multiple groups at a time.

For more information on groups, see "Data collection groups" on page 233.

To create a fabric group, complete the following steps:

1. Expand **Fabric Manager** → **Monitoring** → **Groups**, right-click on **Fabric**, and select **Create Fabric Group**.
2. Enter a description for the group, and from the **Available** column, select the fabrics to be added to the group. Move them to the **Current Selections** column by clicking  .
3. To save the group, click **File** → **Save**.
4. In the **Save As** window, enter the group name and click **OK**.

Editing groups

When you update the resources that are included within a group, those updates are automatically picked up by the jobs that reference that group.

To edit an existing group, perform the following steps.

1. Expand **Monitoring** → **Groups**. For example, to edit groups within Disk Manager, expand **Disk Manager** → **Monitoring** → **Groups** → **Storage Subsystem**.
2. Expand a group type to view the list of groups defined for that type. For example, to view storage subsystem groups, expand the **Storage Subsystem** node under Disk Manager.
3. Click the name of the group you want to edit. Details for the selected group are displayed in the Create *group type* Groups page, where group type is the type of group (for example, Directory or Computer).
4. Edit the values defined for the group as desired. Click  to move the selected resources to the **Current Selections** column on the right. Click  to move resources from the **Current Selections** column and move them to the **Available** column.
5. Click the save icon in the menu bar to save the group.

Creating directory groups

Create and use directory groups within scans to monitor and gather information about specific directories within filesystems across a number of machines.

Note: IBM Tivoli Storage SAN File Systems information: To monitor directories from the SAN File System, you must create a directory group and add the desired directories as follows:

/<IBM Tivoli Storage SAN File Systems cluster name>/<directory name>

where:

- *<IBM Tivoli Storage SAN File Systems cluster name>* represents the name of the IBM Tivoli Storage SAN File Systems cluster that contains the directory you want to add
- *<directory name>* represents the name of the directory you want to add

Examples:

- `/IBM/myFilesetName/myDirectoryName`

- `/IBM/myFilesetName`
- `/IBM/myDirectoryName`

To create a directory group, perform the following steps.

1. Expand **Data Manager → Monitoring → Groups → Directory** or click **Create Directory Group** from the **Directory Groups** tab of the Scan window. The Create Directory Group window opens.
2. Use the New Directory window to edit or add directories to a directory group.

Field	Description
Creator	Displays the creator of the group.
Name	Displays the name of the group you are adding or editing. This name can be helpful in determining the content of the group. For example, a directory group named `UnixUserHomes` might consist of the `/home` directory on all the UNIX and Linux computers within an organization.
Description	A description of the group you are adding or editing.
New Computer, New Directory	Click these buttons to add a directory or computer to the group.

Viewing directory groups:

When adding directories to a group, you can determine how you want to view groups of directories. You can choose either of the following options:

- **Directories by computer.** Use this option to display the directories in the directory group according to the computer on which they exist.
  1. Expand Data Manager.
  2. Expand Monitoring.
  3. Expand Groups.
  4. Expand Directory.
  5. Double-click a directory object.
  6. In the right pane, under View, select Directories by Computer which is the default setting.

     Note: You might need to try more than one directory object to find one with data.
- **Computers by directory.** Use this option to display the computers in the group according to the directories included in the directory group.
  1. Expand Data Manager.
  2. Expand Monitoring.
  3. Expand Groups.
  4. Expand Directory.
  5. Double-click a directory object.
  6. In the right pane, under View, select Computers by Directory.

     Note: You might need to try more than one directory object to find one with data.

Tip: Choosing these different buttons enables you to add directories to a directory group in different ways. For example, if you want to monitor or receive alerts on

different directories on multiple computers, choose the Directories by Computer option. If you want to monitor or receive alerts on the same directory across multiple computers, choose the Computers by Directory option.

Adding new directories by computer:

If you want to add new directories to a directory group by first selecting the computers on which they exist, perform the following steps:

1. Select **directories by computer** on the Create Directory Group window.
2. Click **New computer**. The Edit Directories for Computer window opens. This window enables you to select the computer for which you want to monitor specific directories, and enter the directories for the selected computer.

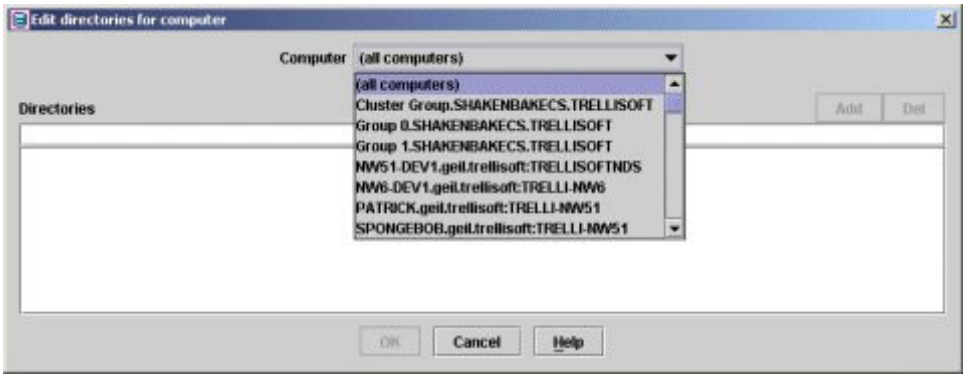

3. Click the **Computer** list box.
4. Select the computer for which you want to monitor specific directories.
5. Enter the directory you want to add to the group in the **Directories** field. For example: `c:\`or `/home/*`.
6. Click **Add**. The directory is added to the list box.
7. Click **OK** when you are done adding directories. The New Directory Group window opens. The computer you added will be displayed as a node in the display area, with the specified directories under the node.

Adding new computers by directory:

If you want to add new directories to a directory group by first specifying the directory and then selecting the computers on which they exist, perform the following steps:

1. Select **computers by directory** on the Create Directory Group window.
2. Click **New computer**. The Edit Computers for Directory window opens. This window enables you to enter the directories you want to add to the directory group and specify which computers contain those directories.
3. Enter the directory you want to add to the group in the **Directory** field. For example: `c:\`, `/home/*`, `\user\*`, `\windows\temp\`, `/usr/*`.
4. Select the computers you want to include in the group and click >>. The computer will be displayed in the **Current Selections** column. Select the **All computers** option to include all computers that contain the directory in the Directory Group.
5. Click **OK**. The New Directory Group window opens. The directory you added will be displayed as a node in the display area, with the selected computers under the node.

Collecting performance data

Use the tasks in this section to collect performance data about your storage subsystems, fabric switches, and tape libraries.

Monitoring switch performance

You can monitor the performance of those switches in your fabrics from which you can collect performance data.

Creating a switch performance monitor

You can monitor the performance of selected fabric switches. Fabric Manager can collect performance data for the switches and generate reports from the results of those monitors.

Report data is collected on data rates, operations rates, switch port errors, and other performance statistics. For information about generating those reports, see "Creating switch performance reports" on page 654.

Note: A fabric probe is now required after a CIMOM discovery before a switch performance monitor job can be run.

To create a switch performance monitor, complete the following steps:

1. Expand **Fabric Manager** → **Monitoring**, right-click **Switch Performance Monitors**, and select **Create Switch Performance Monitors**.
2. In the **Switches** tab enter a description for the performance monitor, and from the **Available** column, select the switches to be added to the monitor definition.

   Move them to the **Current Selections** column by clicking 

.

   Note: Only those switches from which Fabric Manager can collect performance data are listed in the **Available** column. If a switch is added to a performance monitor, it is no longer listed as an available switch.
3. In the **Sampling and Scheduling** tab specify an interval for the performance averages, the duration of the data collections, and a schedule for the monitor. You can begin the monitor immediately, schedule it to begin at a later time, and specify if it should be run repeatedly.
4. In the **Alert** tab specify the action to be taken if the monitor fails. See "Running scripts with alerts" on page 363 for a description of these triggered actions.
5. To save the performance monitor, click **Save**.
6. In the **Save As** window enter the performance monitor name and click **OK**.
7. To display the name of the new monitor, right-click **Switch Performance Monitors** and select **Refresh**. The new name is listed under the node.

Editing a switch performance monitor

Edit a saved switch performance monitor to change settings, such as which switches to include in the collection.

1. Expand **Fabric Manager** → **Monitoring** → **Switch Performance Monitors**. There will be a node for each saved monitor definition.
2. Click the name of the monitor you want to edit.
3. Edit the monitor definition.
4. To save the edited monitor, click **Save**.

Monitoring storage-subsystem performance

Discover, probe, define performance thresholds, and view performance data and alerts to monitor a storage subsystem.

1. Discover the CIMOM of the storage subsystem.

   CIMOM discovery enables communications between IBM Tivoli Storage Productivity Center and the storage subsystem.

   The CIMOM can be discovered through Service Location Protocol (SLP) if the CIMOM is located in the same subnet as the Device server. You can also manually add the CIMOM to Tivoli Storage Productivity Center.
2. Discover the storage subsystem.

   Discovery gets general information about the devices affiliated with the CIMOM: switches, SAN Volume Controller models, and other storage subsystems (DS6000, DS8000, and Tivoli Storage Enterprise Storage Server).
3. Probe the storage subsystem.

   Probes collect information about logical devices, such as volumes and virtual disks, that are associated with the storage subsystem. Probes also collect information about hosts that are mapped to the logical devices.
4. Define storage-subsystem performance thresholds.

   Thresholds draw your attention to exceptions to normally expected values in the data that you collect.
5. Collect storage-subsystem performance data.

   Collect performance data for an individual storage subsystem.
6. View storage-subsystem performance data.

Creating storage subsystem performance monitor

Create a storage-subsystem performance monitor to gather performance data and view the results in the log file.

Complete the following steps to collect storage-subsystem performance data for an IBM SAN Volume Controller:

Note: If you schedule a performance monitor to collect data about a storage subsystem at the same time you plan to create a large number of volumes on that subsystem, the performance of the volume creation job and the general performance of the associated CIMOM might decrease. We recommend scheduling a performance monitoring job at a different time than when you plan to create a large number of volumes on a subsystem.

1. In the left pane, expand **Disk Manager → Monitoring**.
2. Right-click **Subsystem Performance Monitor** and click **Create Performance Monitor**.

   The Create Storage Subsystem Performance Monitor window opens in the right pane.
3. Optional: In the **Description** field, type a description of the threshold.
4. Click the **Storage Subsystems** tab. The Storage Subsystems page is displayed and all storage subsystems that have been discovered are listed.
5. In the **Available subsystems** field, click storage subsystems to select them for collection. Then click 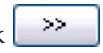 .

   The selected storage subsystems are displayed in the **Current selections** field.
6. Verify that you have created and run a probe for each selected storage subsystem, as described in “Creating probes” on page 203.

Storage-subsystem reports use data gathered by probes. This includes data on storage assets in the enterprise, such as computers, storage subsystems, disk controllers, hard disks, and file systems. Data collection fails if a selected system has not yet been probed.

7. Click the **Sampling and Scheduling** tab. The **Sampling and Scheduling** page is displayed.
8. In the **Sampling** area, specify the length of the sampling window and how frequently the performance data is gathered.
   a. In the **Interval length** field, select the length of the sampling window.

      The interval length is the number of minutes over which performance data is averaged. For example, if you select 15 minutes, all performance-data samples gathered from the storage subsystem represent an average of the performance of the subsystem over the previous 15 minute interval.
   b. In the **Duration** field, specify how long the monitoring task will run.

      You can specify that the task runs indefinitely or specify the length of time that the performance data is gathered.
   c. In the **Scheduling** area, click either **Begin immediately** or **Schedule to begin later**.

      If you click **Schedule to begin later**, specify the data and time that you want to the data collection to start. You also can specify that the data collection will recur, how frequently it recurs, and whether certain days of the week are excluded.
9. Click the **Alert** tab. The **Alert** page is displayed. In the **Condition** area **Monitor Failed** is the triggering condition.
10. In the **Triggered Actions** area, select actions to trigger when the data collection task fails.

    You can choose from the following check boxes:
    - **SNMP Trap**
    - **TEC Event**
    - **Login Notification**
    - **Windows Event Log**
    - **Run Script**
    - **Email**

    Depending on what action you select, you might have other choices to make. Suppose that you select the **Windows Event Log** check box, for example. The **Event Type field** becomes activate to specify a severity for the event in the Windows event log.
11. Click **File → Save**.
12. When prompted, type a name for the threshold, and click **OK**.

    The name of the storage-subsystem performance monitor is displayed under the **Subsystem Performance Monitors** node in the left pane.
13. After the monitoring job has completed, expand the name of the storage-subsystem performance monitor.

    The following information is displayed:
    - A status icon that indicates whether the task completed successfully
    - Number of times that the task ran
    - Date that the task ran
    - Time that the task began
14. To view additional information, click the task information.

In the right pane, the following additional information is displayed:

- Time that the task finished
- Number of jobs in the task
- Number of jobs that failed

A table contains information about the storage subsystems in the data collection.

15. In the table, click a storage subsystem, and then click  to display the contents of the log file.

Example log file

The following is an example of the log file for a SAN Volume Controller:

```
Job log file: <Location of log file>
<date><time>HWNPM21131 The performance monitor for device <name> (<IP address>:<?>) is starting in an
active state
<date><time>HWNPM21151 Monitor Policy:name="svcProbe",creator="administrator", description=""
<date><time>HWNPM21161 Monitor Policy:retention period:sample data=14 days,hourly data=30 days,daily
data=90 days
<date><time>HWNPM21171 Monitor Policy:interval length=900 secs,frequency=900 secs, duration=1 hours
<date><time>HWNPM21181 Threshold Policy:name="Default Threshold Policy for SVC",creator="System",
description="Current default performance threshold policy for SAN Volume Controller"
<date><time>HWNPM21191 Threshold Policy:retention period:exception data=14 days
<date><time>HWNPM21201 Threshold Policy:threshold name=Total I/O Rate Threshold,enabled=no,
boundaries=-1,-1,-1,-1 ops/sec
<date><time>HWNPM21201 Threshold Policy:threshold name=Total Data Rate Threshold,enabled=no,
boundaries=-1,-1,-1,-1 MB/sec
<date><time>HWNPM21201 Threshold Policy:threshold name=Total Backend I/O Rate Threshold,enabled=no,
boundaries=3,2,1,0 ops/sec
<date><time>HWNPM21201 Threshold Policy:threshold name=Total Backend Data Rate Threshold,enabled=yes,
boundaries=3,2,1,0 MB/sec
<date><time>HWNPM21201 Threshold Policy:threshold name=Overall Backend Response Time Threshold,
enabled=yes,boundaries=-1,-1,-1,-1 ms/sec
<date><time>HWNPM02001 This operation (startPerfCollection())on Performance Manager was successful.
<date><time>HWNPM22001 Successfully retrieved the configuration data for the storage system. Found 1
I/O Groups,2MDisk Groups,25MDisks, and 1 VDisk
<date><time>HWNPM21231 Performance data was collected and processed successfully. 27 performance data
records were inserted into the database.
<date><time>HWNPM21231 Performance data was collected and processed successfully. 54 performance data
records were inserted into the database.
<date><time>HWNPM21261 The performance monitor for device <name> (<IP address>:<?>) is stopping because
its intended duration has elapsed.
```

Descriptions of variables in the log file include:

- <date> is the date, for example, 1/23/06.
- <time> is the time, for example, 12:45:20 PM.
- <name> is the device ID of the storage subsystem, for example, SVC-2145-RedC3-IBM.
- <IP address> is the IP address of the storage subsystem, for example, 9.47.97.159.

This log also provides the following information:

- The monitoring task was targeted against a single SAN Volume Controller, SVC-2145-RedC3-IBM.
- The name of the monitoring task is **svcProbe**, and it was created by a user with an ID **administrator**.
- The retention period for exception data.

- The user had provided the following values for the task: interval length: sampling window of 900 seconds (15 minutes) and duration of one hour.
- The monitoring task successfully retrieved the configuration data from the SAN Volume Controller. The task associated the following logical entities with the SAN Volume Controller:
  - One I/O group
  - Two managed-disk groups (MDisk groups)
  - 25 managed disks (MDisks)
  - One virtual disk (VDisk)
- The monitoring task successfully collected and processed data for two performance metrics. A total of 81 performance records were written to the database.
- The monitoring task stopped when its duration (one hour) had elapsed.

Editing a storage-subsystem performance monitor

Edit a storage-subsystem performance monitor.

Complete the following steps to edit storage subsystem performance monitor:

1. In the left pane, expand Disk Manager > **Monitoring** > **Subsystem Performance Monitors**. The list of existing subsystem performance monitors is displayed.
2. Click the monitor that you want to edit. The Edit Subsystem Performance Monitor window opens.

   The Create Storage Subsystem Performance Monitor window opens in the right pane.
3. Make any necessary changes.
4. Click **File → Save**.

Monitoring tape libraries

Monitor the performance of tape libraries to be aware of events that occur, and to respond to alerts.

1. Create one or more tape library groups, as described in “Creating tape library groups” on page 237.
2. Create a probe job to collect data about your tape libraries and library groups, as described in “Creating tape library probes” on page 205.
3. View alerts that are generated by your tape libraries, as described in “Viewing alert logs” on page 262.

Viewing performance data

View collected performance data and recognized threshold exceptions in various user-friendly formats to stay aware of usage trends and exceptions.

Display data in a tabular format (actual values of the collected metrics), a graphical format (primarily as line-graphs), or export the data to a file. Select whether recent data, historical data, or information about threshold exceptions (constraint violations) is displayed.

Select a device or a device component, select a particular metric, and specify a time range to plot the data over the time range.

Use constraint violation reports to view a list of any threshold exceptions that occurred during a particular time range. Bar graphs of exceptions are commonly used to show the total number of threshold violations.

Viewing storage-subsystem performance data

View collected storage-subsystem performance data and recognized threshold exceptions to stay aware of storage-subsystem usage trends and exceptions.

1. In the left pane, expand **Disk Manager → Reporting → Storage Subsystem Performance**.

   A list of reports is displayed under **Storage Subsystem Performance node**.

   The following reports generate information concerning SAN Volume Controller models:

   - By Storage Subsystem
   - By Controller
   - By I/O Group
   - By Node
   - By Array
   - By Managed Disk Group
   - By Volume
   - By Managed Disk
   - By Port
   - Constraint Violations

2. Click **By I/O Group** to view information about I/O groups.

   In the right pane, the Selection page is displayed. By default, all applicable performance metrics are selected for inclusion in the report.

3. Select the data or performance metrics that you want to exclude from the report in the **Included Columns** field. See "Performance metrics" on page 815 for a list of performance metrics.

   a. Click the items that you want to exclude from the report.

   b. Click  .

      Excluded items are displayed in the **Available Columns** field.

   c. Select a metric and click either  or  to change the order of the selected metric in the report.

4. Type the number of I/O groups to include in the report, in the **Return maximum of** field.

   Limit the number of I/O groups to reduce the amount of performance information in the report. By default, this value is set to 2500.

5. Click **Selection** to open the Select Resources window.
6. Select check boxes to select I/O groups for the report, then click **OK**.
7. Click **Filter** to open the Edit Filter window.
8. Specify filter conditions for the report.
9. Click **Generate Report**.

   An I/O Group page opens and the report is displayed.

10. Optional: View the report in chart format.

    a. Click  , to specify charting options, then click **OK**.

b. Click **Customize this chart** to further customize the chart.

11. Optional: Click  and type a name to save the report settings for future use.

    The report is saved under the **IBM Tivoli Storage Productivity Center** node.

Viewing tape library information

View collected tape library data to stay aware of discovered tape libraries and usage information.

- Click **Tape Manager**. Left-click **Tape Libraries**. In the right pane, the Tape Libraries page is displayed, which lists the discovered tape libraries.

  Select a library from the table and then click one of the buttons to see further details about that library. For example, select a library and click **Drives** to open a page that lists all the drives for that library.

  You can also right-click a library in the **Tape Libraries** table and select **View Details**.
- Use the **View Details** feature in the topology viewer, as described in "Working with tape libraries in the topology viewer" on page 404.

Working with alerts

Create an alert, configure how the alert is publicized, or edit, delete, disable, rename, or view an alert.

Related reference

"Performance thresholds" on page 821
Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

Creating alerts

You can define alerts so that IBM Tivoli Storage Productivity Center notifies you when a specified event occurs. The specific triggering conditions that you can use depend on the type of storage resource that you are monitoring.

Note: Alerts are not generated in a Tivoli Storage Productivity Center instance for actions that you perform from that instance. For example, if you start Tivoli Storage Productivity Center and use Disk Manager to assign or unassign volumes for a subsystem, you will not receive alerts for those volume changes in that instance. However, if you assign and unassign volumes outside of that Tivoli Storage Productivity Center instance, an alert is generated.

Related reference

"Performance thresholds" on page 821
Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

Creating computer alerts

You can use these alerts to monitor computers for a changes in configuration or status. When a condition is detected during a probe, an alert is generated. Computer alerts can be applied to computers, computer groups, and clusters.

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.

- If you want to use SNMP traps, IBM Tivoli Enterprise Console events, or e-mail as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** → **Configuration** → **Alert Dispositions**

To create a computer alert, complete the following steps:

1. In the Navigation Tree pane, expand **Data Manager** → **Alerting**.
2. Right-click **Computer Alert** and click **Create Alert**. The Create Alert window opens and the Alert page is displayed.
3. At the far right hand side of the window, click the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert:
   a. In the **Condition** field, select the triggering condition.
   b. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be greyed out.
   c. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be greyed out.
6. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition. For an explanation of triggered actions, see Appendix F, "Triggered actions for alerts," on page 801.
7. Click the **Computer** tab; the Computers page opens.
8. Select the storage entities to which you want the alert to apply:
   a. In the **Available** field, click the storage entities.
   b. Click  . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** → **Save**. The Save As window opens.
   b. In the **Specify Alert name** field, type a name for the alert and click **OK**. The alert is displayed in the Navigation Tree pane

Creating directory alerts

You can use these alerts to be notified when either a directory is not found or a directory consumes more than a specified amount of storage space. When a condition is detected during a scan, an alert is generated. These alerts can be applied to directories only.

Before you can create a directory alert, you must have completed the following prerequisites:

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use SNMP traps, IBM Tivoli Enterprise Console events, or e-mail as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** → **Configuration** → **Alert Dispositions**
- Directories must be grouped into a directory group before they can be included in an alert.

To create a directory alert, complete the following steps:

1. In the Navigation Tree pane, expand **Data Manager** → **Alerting**

2. Right-click **Directory Alert** and click **Create Alert**. The Create Alert window opens and the Alert page is displayed.
3. At the far right hand side of the window, click the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert:
   a. In the **Condition** field, select the triggering condition.
   b. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be greyed out.
   c. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be greyed out.
6. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition. For an explanation of triggered actions, see Appendix F, "Triggered actions for alerts," on page 801.
7. Click the **Directories** tab; the Directories windows opens.
8. Select the storage entities to which you want the alert to apply:
   a. In the **Available** field, click the storage entities.
   b. Click  . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** → **Save**. The Save As window opens.
   b. In the **Specify Alert name** field, type a name for the alert and click **OK**. The alert is displayed in the Navigation Tree pane

Creating file system alerts

You can use these alerts to monitor file systems for available free space, changes in configuration, or a file system being removed or unmounted. When the condition is detected during a scan, an alert is generated. File system alerts can be applied to file systems, file system groups, computers, computer groups, and clusters.

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use SNMP traps, IBM Tivoli Enterprise Console events, or e-mail as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** → **Configuration** → **Alert Dispositions**

To define a file system alert, complete the following steps:

1. In the Navigation Tree pane, expand **Data Manager** → **Alerting**.
2. Right-click **Filesystem Alert** and click **Create Alert**. The Create Alert window opens and the Alert page is displayed.
3. At the far right hand side of the window, click the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert:
   a. In the **Condition** field, select the triggering condition.
   b. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be greyed out.
   c. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be greyed out.

6. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition. For an explanation of triggered actions, see Appendix F, "Triggered actions for alerts," on page 801.
7. Click the **Filesystems** tab; the Filesystems window opens.
8. Select the storage entities to which you want the alert to apply:
   a. In the **Available** field, click the storage entities.
   b. Click  . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** → **Save**. The Save As window opens.
   b. In the **Specify Alert name** field, type a name for the alert and click **OK**. The alert is displayed in the Navigation Tree pane

Creating instance alerts

You can use these alerts to monitor instances for changes. When the condition is detected during a probe, an alert is generated. These alerts can be applied to instances, computers, and computer groups.

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use SNMP traps, IBM Tivoli Enterprise Console events, or e-mail as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** → **Configuration** → **Alert Dispositions**

To create an instance alert, complete the following steps:

1. From the Navigation Tree pane, expand **Data Manager for Databases** → **Alerting**.
2. Right-click **Instance Alerts** and click **Create Alert.**. The Create Alert window opens and the Alert page is displayed.
3. At the far right hand side of the window, click the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert:
   a. In the **Rdbms Type** field, select the type of database application. You can choose **All Rdbms**, **Oracle**, **SQL\Server**, **Sybase**, or **UDB**.
   b. In the **Condition** field, select the triggering condition.
   c. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be greyed out.
   d. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be greyed out.
6. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition. For an explanation of triggered actions, see Appendix F, "Triggered actions for alerts," on page 801.
7. Click the **Instances** tab; the **Instances** window opens.
8. Select the storage entities to which you want the alert to apply:
   a. In the **Available** field, click the storage entities.
   b. Click  . The storage entities are moved to the **Current Selections** field.
9. Save the alert:

a. Click **File** → **Save**. The Save As window opens.
b. In the **Specify Alert name** field, type a name for the alert and click **OK**. The alert is displayed in the Navigation Tree pane

Creating database and tablespace alerts

You can use these alerts to monitor database and tablespaces for changes. When a scans or probes are run, events are recognized and alerts triggered. These alerts can be applied to databases, database groups, table spaces, table space groups, instances, computers, and computer groups.

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use SNMP traps, IBM Tivoli Enterprise Console events, or e-mail as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** → **Configuration** → **Alert Dispositions**

Note: (DB2 only) Alerts triggered by table and user statistics will not function until the **runstats** utility is run. This utility ensures that table and user statistics are up-to-date. When you define a scan job for DB2 databases, select the **Perform runstats on all tables being scanned** check box on the **Options** page.

Note: (DB2 only) When the following conditions apply, an alert is not generated:

- The triggering condition selected is Database-Tablespace Free Space Less Than.
- The storage entities selected are DB2 databases.
- The table spaces affected by the triggering condition are System Managed Space (SMS) tablespaces.

The alert is not triggered because IBM Tivoli Storage Productivity Center considers SMS containers to be directories on a file system

To create a database-table space alert, complete the following steps:

1. From the Navigation Tree pane, expand **Data Manager for Databases** → **Alerting**.
2. Right-click **Database-Tablespace Alerts** and click **Create Alert.**. The Create Alert window opens and the Alert page is displayed.
3. At the far right hand side of the window, click the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert:
   a. In the **Rdbms Type** field, select the type of database application. You can choose **All Rdbms**, **Oracle**, **SQL\Server**, **Sybase**, or **UDB**.
   b. In the **Condition** field, select the triggering condition.
   c. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be greyed out.
   d. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be greyed out.
6. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition. For an explanation of triggered actions, see Appendix F, "Triggered actions for alerts," on page 801.
7. Click the **Databases-Tablespaces** tab; the **Databases-Tablespaces** window opens.
8. Select the storage entities to which you want the alert to apply:
   a. In the **Available** field, click the storage entities.

b. Click >>. The storage entities are moved to the **Current Selections** field.

9. Save the alert:
   a. Click **File** → **Save**. The Save As window opens.
   b. In the **Specify Alert name** field, type a name for the alert and click **OK**. The alert is displayed in the Navigation Tree pane

Creating table alerts

You can use these alerts to monitor tables for certain conditions. When the condition is detected during a scan, an alert is generated. These alerts can be applied to table groups and tables only.

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use SNMP traps, IBM Tivoli Enterprise Console events, or e-mail as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** → **Configuration** → **Alert Dispositions**
- Tables must be grouped into a table group before they can be included in an alert.

To create a table alert, complete the following steps:

1. From the Navigation Tree pane, expand **Data Manager for Databases** → **Alerting**.
2. Right-click **Table Alerts** and click **Create Alert.**. The Create Alert window opens and the Alert page is displayed.
3. At the far right hand side of the window, click the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert:
   a. In the **Rdbms Type** field, select the type of database application. You can choose **All Rdbms**, **Oracle**, **SQL\Server**, **Sybase**, or **UDB**.
   b. In the **Condition** field, select the triggering condition.
   c. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be greyed out.
   d. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be greyed out.
6. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition. For an explanation of triggered actions, see Appendix F, “Triggered actions for alerts,” on page 801.
7. Click the **Tables** tab; the **Tables** window opens.
8. Select the storage entities to which you want the alert to apply:
   a. In the **Available** field, click the storage entities.
   b. Click >>. The storage entities are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** → **Save**. The Save As window opens.
   b. In the **Specify Alert name** field, type a name for the alert and click **OK**. The alert is displayed in the Navigation Tree pane

Creating fabric alerts

You can use these alerts to monitor fabrics for changes. When the condition is detected during a probe, an alert is generated. These alerts can be applied to fabrics, zones, and zone sets.

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use SNMP traps, IBM Tivoli Enterprise Console events, or e-mail as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** → **Configuration** → **Alert Dispositions**

1. In the Navigation Tree pane, expand **Fabric Manager** → **Alerting**.
2. Right-click **Fabric Alert** and click **Create Fabric Alerts**. The Create Fabric Alerts window opens and the Alert page is displayed.
3. At the far right hand side of the window, click the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert.
6. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition. For an explanation of triggered actions, see Appendix F, "Triggered actions for alerts," on page 801.
7. Click the **Fabrics** tab; the Fabrics page opens.
8. Select the storage entities to which you want the alert to apply:
   a. In the **Available** field, click the storage entities.
   b. Click  . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** → **Save**. The Save As window opens.
   b. In the **Specify Alert name** field, type a name for the alert and click **OK**. The alert is displayed in the Navigation Tree pane

Creating a storage subsystem alerts

You can use these alerts to monitor storage subsystems for changes in configuration or status. When the condition is detected during a probe, an alert is generated. These alerts can be applied to storage subsystem groups and storage subsystems.

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use SNMP traps, IBM Tivoli Enterprise Console events, or e-mail as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** → **Configuration** → **Alert Dispositions**

To create a storage subsystem alert, complete the following steps:

1. In the Navigation tree pane, expand **Disk Manager** → **Alerting**.
2. Right-click **Create Storage Subsystem Alert**, and click **Create Storage Subsystem Alert**. The Create Storage Subsystem Alert window opens and the Alert page is displayed.
3. At the far right hand side of the window, click the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert:

a. In the **Condition** field, select the triggering condition.
b. (If you selected a performance threshold) Specify the thresholds:
   1) In the **Critical Stress** field, type the upper boundary for the threshold. When this boundary is exceeded, a critical condition is recognized.
   2) In the **Warning Stress** field, type the upper boundary for the threshold. When this boundary is exceeded, a warning condition is recognized.
   3) In the **Warning Idle** field, type the lower boundary for the threshold. When the value falls below this boundary, a warning condition is recognized.
   4) In the **Critical Idle** field, type the lower boundary for the threshold. When the value falls below this boundary, a critical is recognized.

   For more information about thresholds, see "Performance thresholds" on page 821

6. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition. For an explanation of triggered actions, see Appendix F, "Triggered actions for alerts," on page 801.
7. Click the **Storage Subsystems** tab; the Storage Subsystems page opens.
8. Select the storage subsystems and storage subsystem groups to which you want the alert to apply:
   a. In the **Available** field, click the storage entities.
   b. Click  . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** → **Save**. The Save As window opens.
   b. In the **Specify Alert name** field, type a name for the alert and click **OK**. The alert is displayed in the Navigation Tree pane

Creating switch alerts

You can use these alerts to monitor switches for changes in configuration or status. When the condition is detected during a probe, an alert is generated. These alerts can only be applied to switches.

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use SNMP traps, IBM Tivoli Enterprise Console events, or e-mail as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** → **Configuration** → **Alert Dispositions**

To create a switch alert, complete the following steps:

1. In the Navigation Tree pane, expand **Fabric Manager** → **Alerting**.
2. Right-click **Switch Alerts** and click **Create Switch Alerts**. The Create Switch Alerts page opens and the Alert page is displayed.
3. At the far right hand side of the window, click the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert:
   a. In the **Condition** field, select the triggering condition.
   b. (If you selected a performance threshold) Specify the thresholds:
      1) In the **Critical Stress** field, type the upper boundary for the threshold. When this boundary is exceeded, a critical condition is recognized.

2) In the **Warning Stress** field, type the upper boundary for the threshold. When this boundary is exceeded, a warning condition is recognized.
3) In the **Warning Idle** field, type the lower boundary for the threshold. When the value falls below this boundary, a warning condition is recognized.
4) In the **Critical Idle** field, type the lower boundary for the threshold. When the value falls below this boundary, a critical is recognized.

For more information about thresholds, see "Performance thresholds" on page 821

6. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition. For an explanation of triggered actions, see Appendix F, "Triggered actions for alerts," on page 801.
7. Click the **Switches** tab; the Switches page opens.

   Note: If you selected a performance threshold as the triggering condition, only those switches for which performance data can be collected are listed.
8. Select the storage entities to which you want the alert to apply:
   a. In the **Available** field, click the storage entities.
   b. Click  . The storage entities are moved to the **Current Selections** field.
9.

   Note: You can select one or more switches from the **Switches** category. For device triggering conditions only, you can select one or more fabrics from the **Switches In Fabric** category. Doing so will include all switches in those fabrics. Or you can select from both categories.
10. Save the alert:
    a. Click **File** → **Save**. The Save As window opens.
    b. In the **Specify Alert name** field, type a name for the alert and click **OK**. The alert is displayed in the Navigation Tree pane

Creating endpoint device alerts

You can use these alerts to monitor endpoint devices. When the condition is detected during a probe, an alert is generated. These alerts can be applied to endpoint devices only.

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use SNMP traps, IBM Tivoli Enterprise Console events, or e-mail as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** → **Configuration** → **Alert Dispositions**

To create an endpoint device alert, complete the following steps:

1. In the Navigation Tree pane, expand **Fabric Manager** → **Alerting**.
2. Right-click **Endpoint Device Alerts** and click **Create Endpoint Device Alerts**. The Create Endpoint Device Alerts window opens and the Alert page is displayed.
3. At the far right hand side of the window, click the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert.

6. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition. For an explanation of triggered actions, see Appendix F, "Triggered actions for alerts," on page 801.
7. Click the **Endpoint Devices** tab; the Endpoint Devices page opens.
8. Select the storage entities to which you want the alert to apply:
   a. In the **Available** field, click the storage entities.
   b. Click  . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
   a. Click **File** → **Save**. The Save As window opens.
   b. In the **Specify Alert name** field, type a name for the alert and click **OK**. The alert is displayed in the Navigation Tree pane

Database connection alerts

Database connection alerts notify the Tivoli Storage Productivity Center administrator when the database connection used by the Tivoli Storage Productivity Center server is lost unexpectedly and when the database connection is restored.

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use SNMP traps, IBM Tivoli Enterprise Console events, or e-mail as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services** → **Configuration** → **Alert Dispositions**

The database connection alerts are always enabled and configured to add the alert to the alert log and send an e-mail to the data administrator.

The database connection alerts are:

- DB connection failed
- DB connection successful

You cannot create, configure, delete, disable, edit, or rename these alerts. You can, however, change the destination for the database connection alerts. Expand **Administrative Services** → **Configuration**. Left-click **Alert Disposition**. In the right pane, under Email, you can change all of the fields.

Configuring alert destinations

Entries in the Alert Disposition window define destination host addresses for SNMP, IBM Tivoli Enterprise Console, and e-mail.

If you want to send SNMP traps, IBM Tivoli Enterprise Console events, or e-mail when an alert is triggered, you must configure IBM Tivoli Storage Productivity Center with the alert destinations and e-mail settings.

To configure alert dispositions, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services** → **Configuration.**
2. Click **Alert Disposition**. The Edit Alert Disposition window opens.
3. In the **SMNP** area, specify the information about where SNMP traps are sent:
   a. In the **Community** field, type the name of the SNMP community. By default, this is set to public.

b. In the **Host field**, type the fully qualified name of the system that you have configured to receive SNMP traps.
c. In the **Port** field, type the port number to which Tivoli Storage Productivity Center will direct SNMP traps. By default, this is set to 162.
d. If you want to configure a second SNMP destination, type the required information in the fields in the **SNMP Destination 2** area.

4. In the **TEC** area, specify information about where Tivoli Enterprise Console events are sent:
   a. In the **Server** field, type the fully qualified name of the Tivoli Enterprise Console server.
   b. In the **Port** field, type the port number to which Tivoli Storage Productivity Center will direct Tivoli Enterprise Console events. By default, this is set to 5529.
5. In the **Email** area, configure e-mail settings:
   a. In the **Mail Server** field, type the name of the mail server.
   b. In the **Mail Port** field, type the port number for the SMTP server.
   c. In the **Default Domain field**, type the domain, for example, MyCompany.com.
   d. In the **Return to field**, type the e-mail address to which undeliverable mail is sent.
   e. In the **Reply to field**, type the e-mail address to which e-mail is sent when someone replies to an e-mail notification about an alert. (It is a best practice to specify the e-mail address of an Tivoli Storage Productivity Center administrator.)
   f. In the **Data Administrator** field, type the e-mail address of the Data Administrator.
6. In the **Alert Log Disposition** field, select the length of time to keep entries in the alert logs. The default is 90 days.

Deleting alerts

You can delete alerts that you no longer want to implement. This ensures that the list of alerts in the Navigation Tree pane is up-to-date.

To delete an alert, complete the following steps:

1. In the Navigation Tree pane, perform one of the following actions:

For computer alerts	Expand **Data Manager → Alerting → Computer Alerts**.
For file system alerts	Expand **Data Manager → Alerting → Filesystem Alerts**.
For directory alerts	Expand **Data Manager → Alerting → Directory Alerts**.
For instance alerts	Expand **Data Manager for Databases → Alerting → Instance Alerts**.
For database-tablespace alerts	Expand **Data Manager for Databases → Alerting → Database-Tablespace Alerts**.
For table alerts	Expand **Data Manager for Databases → Alerting → Table Alerts**.
For storage subsystem alerts	Expand **Disk Manager → Alerting → Storage Subsystem Alerts**.
For fabric alerts	Expand **Fabric Manager → Alerting → Fabric Alerts**.
For switch alerts	Expand **Fabric Manager → Alerting → Switch Alerts**.

For endpoint device alerts	Expand **Fabric Manager** → **Alerting** → **Endpoint Device Alerts**.

The list of existing alerts is displayed.

2. Right-click the alert, and click **Delete**. The Delete Alert window opens.
3. Click **Yes**.

Disabling an alert

You can disable an alert. This retains the alert definition but prevents the alert from being run.

To disable an alert, complete the following steps:

1. In the Navigation Tree pane, complete one of the following actions:

For computer alerts	Expand **Data Manager** → **Alerting** → **Computer Alerts**.
For file system alerts	Expand **Data Manager** → **Alerting** → **Filesystem Alerts**.
For directory alerts	Expand **Data Manager** → **Alerting** → **Directory Alerts**.
For instance alerts	Expand **Data Manager for Databases** → **Alerting** → **Instance Alerts**.
For database-tablespace alerts	Expand **Data Manager for Databases** → **Alerting** → **Database-Tablespace Alerts**.
For table alerts	Expand **Data Manager for Databases** → **Alerting** → **Table Alerts**.
For storage subsystem alerts	Expand **Disk Manager** → **Alerting** → **Storage Subsystem Alerts**.
For fabric alerts	Expand **Fabric Manager** → **Alerting** → **Fabric Alerts**.
For switch alerts	Expand **Fabric Manager** → **Alerting** → **Switch Alerts**.
For endpoint device alerts	Expand **Fabric Manager** → **Alerting** → **Endpoint Device Alerts**.

The list of existing alerts is displayed.

2. Click the alert that you want to disable. The Edit Alert page opens.
3. Clear the **Enabled** check box.
4. Click **File** → **Save**.

Related topics:

"Alerts" on page 62

Editing alerts

You can edit an alert if you want to change the triggering condition, triggered action, or the storage resources against which it is deployed.

To edit an alert, complete the following steps:

1. In the Navigation Tree pane, perform one of the following actions:

For computer alerts	Expand **Data Manager** → **Alerting** → **Computer Alerts**.
For file system alerts	Expand **Data Manager** → **Alerting** → **Filesystem Alerts**.
For directory alerts	Expand **Data Manager** → **Alerting** → **Directory Alerts**.
For instance alerts	Expand **Data Manager for Databases** → **Alerting** → **Instance Alerts**.

For database-tablespace alerts	Expand **Data Manager for Databases → Alerting → Database-Tablespace Alerts**.
For table alerts	Expand **Data Manager for Databases → Alerting → Table Alerts**.
For storage subsystem alerts	Expand **Disk Manager → Alerting → Storage Subsystem Alerts**.
For fabric alerts	Expand **Fabric Manager → Alerting → Fabric Alerts**.
For switch alerts	Expand **Fabric Manager → Alerting → Switch Alerts**.
For endpoint device alerts	Expand **Fabric Manager → Alerting → Endpoint Device Alerts**.

The list of existing alerts is displayed.

2. Click the alert that you want to edit. The Edit Alert page opens.
3. Make any necessary changes.
4. Click **File → Save**.

Renaming alerts

You can rename alerts. This might be useful if the original name is ambiguous or lacks clarity.

To rename an alert, complete the following steps:

1. In the Navigation Tree pane, perform one of the following actions:

For computer alerts	Expand **Data Manager → Alerting → Computer Alerts**.
For file system alerts	Expand **Data Manager → Alerting → Filesystem Alerts**.
For directory alerts	Expand **Data Manager → Alerting → Directory Alerts**.
For instance alerts	Expand **Data Manager for Databases → Alerting → Instance Alerts**.
For database-tablespace alerts	Expand **Data Manager for Databases → Alerting → Database-Tablespace Alerts**.
For table alerts	Expand **Data Manager for Databases → Alerting → Table Alerts**.
For storage subsystem alerts	Expand **Disk Manager → Alerting → Storage Subsystem Alerts**.
For fabric alerts	Expand **Fabric Manager → Alerting → Fabric Alerts**.
For switch alerts	Expand **Fabric Manager → Alerting → Switch Alerts**.
For endpoint device alerts	Expand **Fabric Manager → Alerting → Endpoint Device Alerts**.

The list of existing alerts is displayed.

2. Right-click the alert, and click **Rename**. The Rename Alert window opens.
3. Type the new name, and click **OK**.

Viewing alerts

This section describes different methods for viewing alerts.

To view an alert, choose one of the following methods:

Alert log

Click **IBM Tivoli Storage Productivity Center → Alerting → Alert Log**. Click **All Alerts** or any node that is highlighted in red and preceded by a red circle. Depending on what you select, the Alert History page displays all alerts or only the alerts that you selected. Figure 21 shows the alert history page.

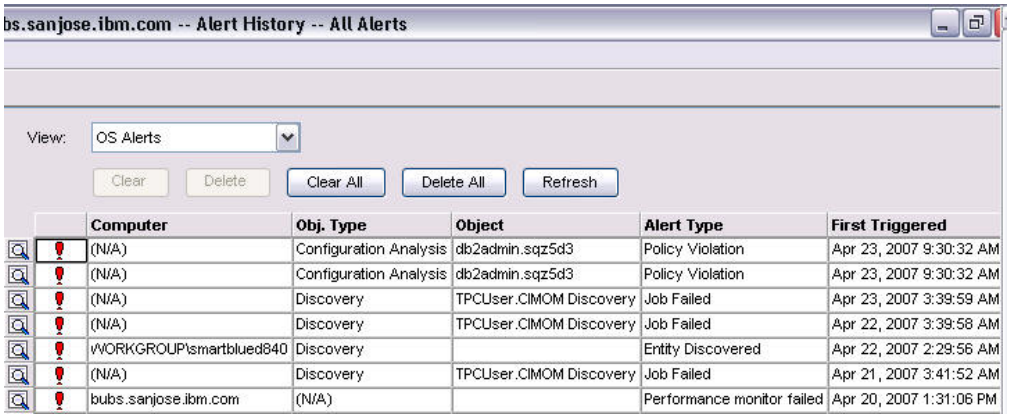

Figure 21. Alert History page

To view details about each alert, click  or double-click  . The Detail for Alert page displays additional information (see Figure 22). You can obtain similar information about any alert by clicking **Topology** to access the topology viewer. Details are available through the *alert overlay* in the graphical view and the Alert tab in the table view. The following figure shows the Detail for Alert page, which lists the name of the job, number of the run, type of alert, state of the alert, timestamp, name of the alert creator, name of the alert, and text about the alert

Figure 22. Detail for Alert page

Alert overlay

The alert overlay represents various types of alerts. You can enable or disable the alert overlay in the graphical view of the topology viewer by right-clicking a blank space in the background of the graphical view,

clicking **Global Settings** from the pop-up menu, and clicking **Alert** under Active overlays (see Figure 23). After you make your selection, click **OK**.

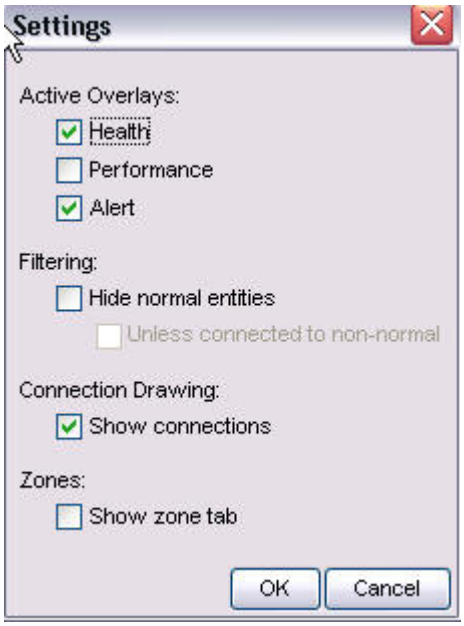

Figure 23. Enabling or disabling the alert overlay. In this example, the alert overlay is enabled.

When the alert overlay is active, entities that have alerts associated with them display a red alert overlay icon  that is located to the right of the entity and below the pinning icon. The icon is similarly located in the title bar of groups and tabs. Note that in the Configuration History page the alert overlay is disabled.

The alert overlay provides rollup, similar to pinning. If alerts are inside a collapsed group, the number of alerts within that group are aggregated and indicated in the collapsed group with an alert icon. All alerts are aggregated into one count, independent of type or entity involved. Hover over the alert icon for a tab or group to determine the total number of alerts that are rolled up into that icon and the time stamp of the newest alert.

In the topology viewer, right-click an entity with an alert icon to display the pop-up menu. Click **Show Alerts in Tabular View** to move to the Alert tab and to highlight alerts that are associated with that entity in the table view (see Figure 24 on page 262). When you select an alert under the Alert tab, entities that are associated with that alert are highlighted in purple in the graphical view. Such an action might highlight more than one entity in the topology viewer.

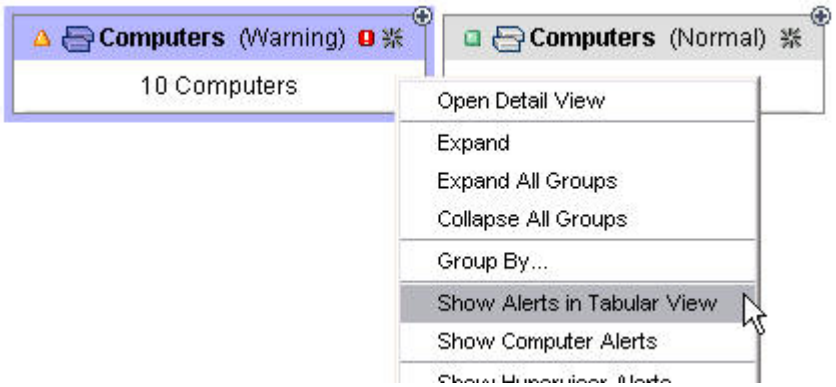

Figure 24. Using the pop-up menu to view alerts in the table view

External alerts do not display in the alert overlay. These alerts are usually Common Information Model (CIM) indications that cannot be tied to specific resources. To view the list of external alerts, click **IBM Tivoli Storage Productivity Center → Alerting → Alert Log → External**.

Alert tab

The Alert tab is located in the table view of the topology viewer (see Figure 25). Click this tab to view a hierarchical list of all current alerts that are associated with entities shown in the topology. For configuration analysis alerts, the top level provides information about which policy was violated during a particular analyzer run. The next level indicates the number of times that the policy was violated. This list contains one node for each policy violation. The Alert tab only lists the policy violations that contain at least one of the affected entities that are displayed in the graphical view. The remaining policy violations are filtered. The only exception is the Overview panel where the Alert tab lists all of the Tivoli Storage Productivity Center alerts. Click each violation to view a list of affected entities. Parent (enclosing) entities display in a separate column.

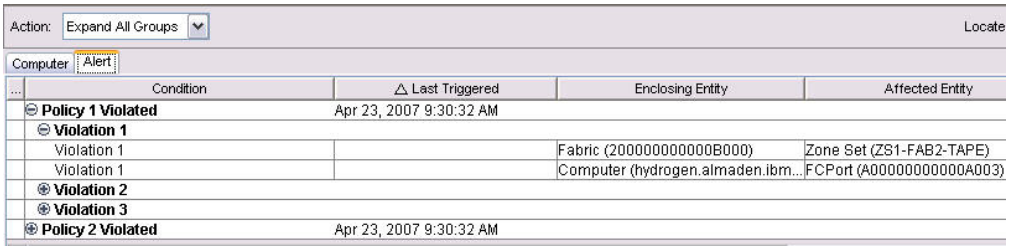

Figure 25. Viewing alerts by using the Alert tab in the table view. Expand the levels to view information about each violation.

For information about fixing policy violations, see "Resolving policy violations" on page 276. You can resort the alert list by clicking the headers in the table. To find entities associated with an alert in the graphical view, click the alert in the table to highlight (in purple) the corresponding entity and any associated entities.

Viewing alert logs

The alert logs record information about the generated alerts. There are two alert logs that you can view:

- For alerts generated by Data Manager for Databases expand **Data Manager for Databases → Alerting → Alert Log**. You can view:
  - All alerts or alerts organized by RDBMS such as by Oracle, SQL Server, Sybase, and UDB. If you do not license for a specific RDBMS, no nodes for that RDBMS will appear under the **Alert Log** node.
  - Alerts by type (such as by Instance, Database, Tablespace, Table, User)
  - Alerts directed to you (the user currently logged into the UI component)
- For alerts generated by all other components, expand **IBM Tivoli Storage Productivity Center → Alerting → Alert Log**. A list of the following alerts is displayed:
  - All
  - Alerts Directed to <user-id> - Alerts directed to the user logged into the interface.
  - Storage Subsystem
  - Computer
  - Disk
  - File System
  - Directory
  - User
  - OS User Group
  - Fabric
  - Switch
  - Endpoint Device
  - External
  - Tape Library
  - Configuration Analysis

Each alert node has an icon next to it that indicates its status:

- Red-outlined circle: a negative alert has occurred, such as job failed, new disk defect.
- Green-outlined square: a neutral alert has been detected, such as a new disk detected.

When you click an alert node, the alert log window opens a table containing information for all the alerts for that node. Click  on a row to drill down for more information.

The following table describes the buttons at the top of the Alert Log window:

Button	Description
Clear	Clear the selected alerts. When you click this button, the icon beside the alert will change to indicate that the alert has been cleared. The alert will continue to appear in the alert log.
Delete	Delete the selected alerts from the alert log.
Clear ALL	Clear all of the alerts in the report. When you click this button, the icon besides all the alerts will change to indicate that the alerts have been cleared. The alerts will continue to appear in the alert log.
Delete ALL	Delete all of the alerts in the current report from the alert log.

Button	Description
Refresh	Regenerate the report you are viewing to pick up any newly generated alerts that have been triggered since you last generated the report.

Viewing tape library alerts

View discovered alert conditions in the tape libraries. Alerts are based on the CIM_AlertIndication specification of SMI-S.

To view these alerts, expand **IBM Tivoli Storage Productivity Center → Alerting → Alert Log** and click the **Tape Library** node.

Checking storage configurations

This section introduces methods that a system administrator can use to assess changes that have occurred in SAN configurations over time and to check whether SAN configurations comply with predefined policies.

Using configuration history

This topic describes how system administrators can capture and analyze historical data that identifies possible problems with a storage area network (SAN) configuration.

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

The configuration history feature takes and displays snapshots of changes that occurred in your SAN configuration over a period of time that you specify. After you set the time period (how often to take snapshots and how long to store them), in a page similar to the topology viewer you can manipulate a snapshot selection panel to show changes that occurred between two or more points in the time period. System administrators can use the configuration history feature to:

Correlate performance statistics with configuration changes
For example, during collection of performance statistics (including volume performance statistics) on an Tivoli Storage Enterprise Storage Server system you might delete a volume. While no new statistics are reported on that volume, the Tivoli Storage Productivity Center Performance Manager would have already collected partial statistical information prior to the deletion. At the end of the data collection task, reporting of the partially collected statistics on the (now) deleted volume would require access to its properties which would not be available. The configuration history feature, with its ability to take and store snapshots of a system's configuration, could provide access to the volume's properties.

Analyze end-to-end performance
You want to know why performance changed on volume A during the last 24 hours. To learn why, it is useful to know what changes were made to the storage subsystem's configuration that might affect the volume's performance, even if performance statistics were not recorded on some of those elements. For example, even if performance statistics on a per-rank basis are not collected, but the number of volumes allocated on a rank is increased from 1 to 100 over time, access to that configuration history information helps with analyzing the volume's degraded performance over time.

Aid in planning and provisioning

The availability of configuration history can enhance the quality of both provisioning and planning. For example, historical data is useful when using the Tivoli Storage Productivity Center Volume Performance Advisor to provision a volume or when using the Tivoli Storage Productivity Center Version 3.3 Integrated Host, Security and Subsystem Planner to plan tasks.

To use the configuration history feature, complete the following steps:

1. In the Navigation tree pane, expand **Administrative Services** → **Configuration** → **Configuration History Settings**. The Configuration History Settings page displays for you to indicate how often to capture SAN configuration data and how long to retain it.
2. Perform the following to collect historical data:
   a. In the **Create snapshot every** field, click the check box to enable this option and type how often (in hours) you want the system to take snapshot views of the configuration.
   b. In the **Delete snapshots older than** field, click the check box to enable this option and type how long (in days) you want the snapshots to be stored.
   c. The page displays the total number of snapshots in the database and the date and time of when the latest snapshot was taken. To refresh this information, click **Update**.
   d. To optionally create and title a snapshot on demand, in the **Title this snapshot** field type a name for the on demand snapshot and click **Create Snapshot Now**. If you do not want to title the on demand snapshot, simply click **Create Snapshot Now**.
   e. To return to the default settings (the default settings are create snapshots every 12 hours and delete snapshots older than 14 days), click **Reset to defaults**.
   f. To save your settings, click **File** → **Save**.
3. In the Navigation tree pane, expand **Tivoli Storage Productivity Center** → **Analytics** → **Configuration History**, and click **Configuration History**. The software loads the snapshot data for the length of time that you specified. The Configuration History page (a variation of the topology viewer) displays the configuration's entities and a floating snapshot selection panel. The panel allows you to define the time periods against which the configuration is compared to determine whether changes have occurred.

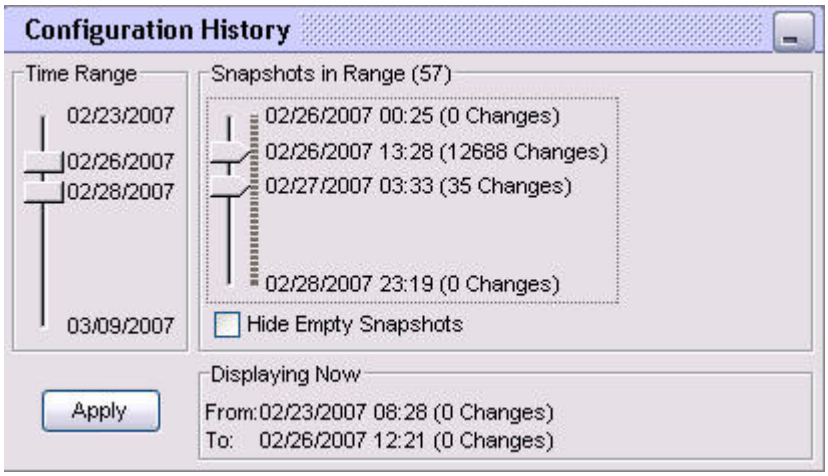

Figure 26. Snapshot selection panel. Use the thumb sliders to establish the time interval that you want to examine.

4. To define the time periods that you want to compare, perform the following:
   a. Using the mouse, drag the two "thumbs" in the left **Time Range** slider to establish the desired time interval. The Time Range slider covers the range of time from the oldest snapshot in the system to the current time. It indicates the date as *mm/dd/yyyy*, where *mm* equals the month, *dd* equals the day, and *yyyy* equals the year.
   b. Drag the two thumbs in the right **Snapshots in Range** slider to indicate the two snapshots to compare. The Snapshots in Range slider allows you to select any two snapshots from the time interval specified by the Time Range slider. The value in parentheses beside the Snapshots in Range slider indicates the total snapshots in the currently selected time range. The Snapshots in Range slider has one tick mark for each snapshot from the time interval that you specified in the Time Range slider. Each snapshot in the Snapshots in Range slider is represented as time stamp *mm/dd/yyyy hh:mm*, where the first *mm* equals the month, *dd* equals the day, *yyyy* equals the year, *hh* equals the hour, and the second *mm* equals the minute. The value in parentheses beside each snapshot indicates the number of changes that have occurred between this and the previous snapshot. Snapshots with zero changes are referred to as empty snapshots. If you provided a title while creating an on demand snapshot, the title displays after the time stamp. If you want to remove empty snapshots, click the check box to display a check mark in **Hide Empty Snapshots**. The **Displaying Now** box indicates the two snapshots that are currently active.
   c. Click **Apply**.
5. Determine the changes that have occurred to the entities by examining the icons and colors associated with them in the graphical and table views: For information about viewing the changes, see "Viewing configuration changes in the graphical view" on page 267 and "Viewing configuration changes in the table view" on page 269.

One single snapshot selection panel applies for all Configuration History views that are open at the same time. Any change that you make in this panel is applied to all of the Configuration History views.

Important: If you manually deleted a missing entity from the Tivoli Storage Productivity Center, the configuration history feature rediscovers and displays the entity in the configuration history viewer as both Missing and in a Normal state.

For example, Entity A (with ID 12) is missing and you decide to manually delete it by using the Remove From Database option from the pop-up menu for Entity A in the topology viewer or configuration history viewer. Later, the Tivoli Storage Productivity Center rediscovers Entity A and assigns it a new ID (ID 22). The configuration history feature treats the same Entity A (with ID 12 and ID 22) as two different entities and displays them in both Missing and Normal states in the configuration history viewer. This scenario can also occur if a missing entity is removed by the Tivoli Storage Productivity Center removed resource retention function and is rediscovered by the Tivoli Storage Productivity Center again at a later time.

Viewing configuration changes in the graphical view

This section describes how to determine the changes that have occurred to a configuration over time by examining the icons and colors of the change overlay in the graphical view of the Configuration History page.

In the Configuration History page, a *change overlay* presents icons and colors that indicate changes in the configuration between the time that a snapshot was taken and the time that a later snapshot was taken:

- The icons display beside the name of the entity in the graphical view and in the change status column of the table view.
- The colors display as background colors for entities in the graphical view and for affected rows in the table view.

Table 15 describes the icons and colors of the change overlay.

Table 15. Icons and colors of the change overlay

Change overlay icon	Change overlay color indicator	Description
Yellow pencil	Blue background	Entity changed between the time that the snapshot was taken and the time that a later snapshot was taken.
 Light gray circle	Dark gray background	Entity did not change between the time that the snapshot was taken and the time that a later snapshot was taken.
 Green cross	Green background	Entity was created or added between the time that the snapshot was taken and the time that a later snapshot was taken.
Red minus sign	Red background	Entity was deleted or removed between the time that the snapshot was taken and the time that a later snapshot was taken.
Not applicable	Light gray background	Entity did not exist at the time that the snapshot was taken or at the time that a later snapshot was taken.

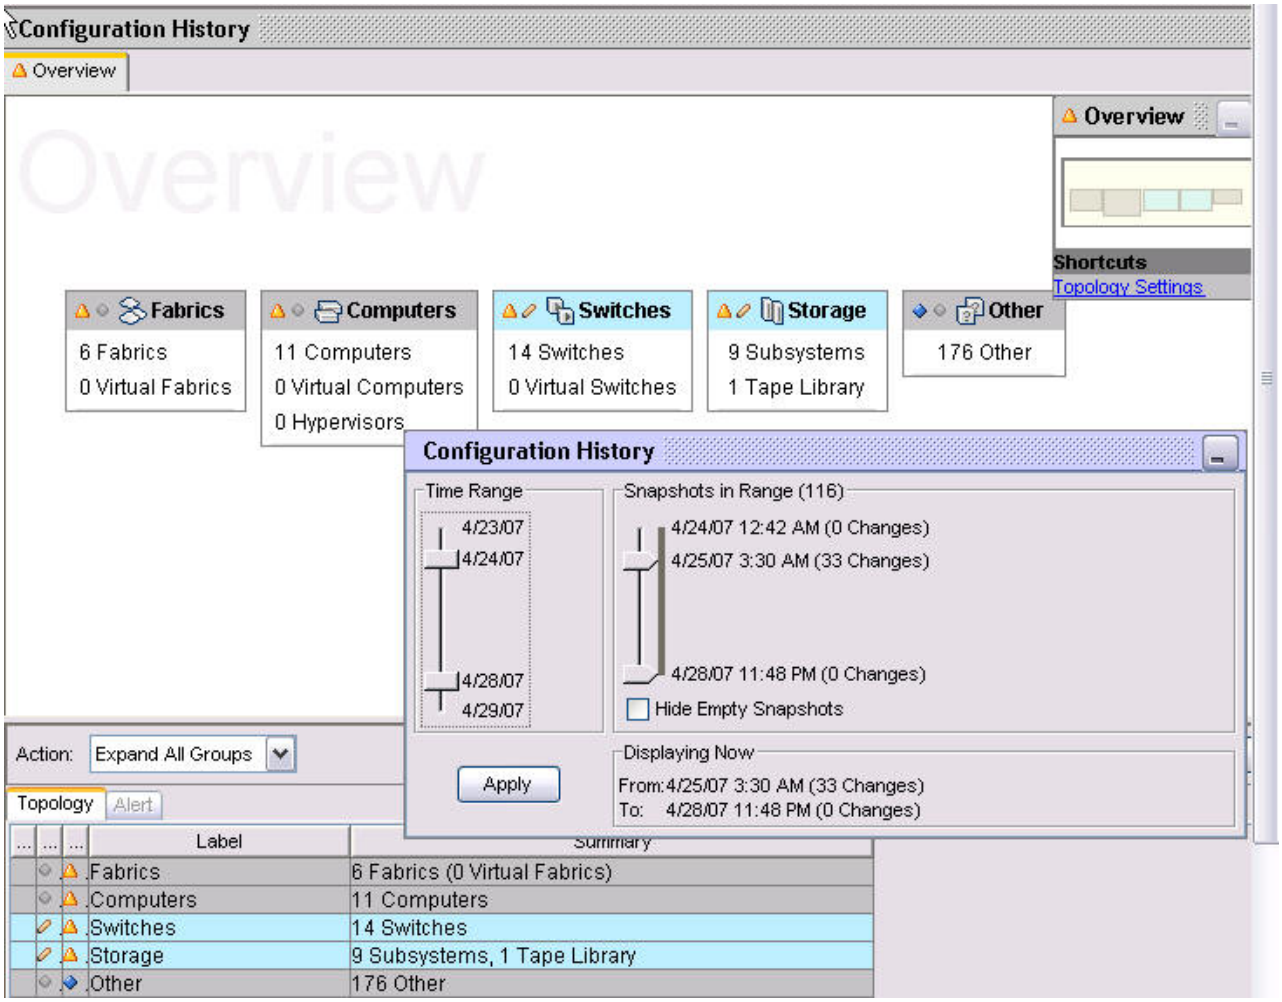

Figure 27. Icons and colors of the change overlay. In the graphical view, the pencil icon beside the switches and storage entities and the blue background color indicate that change occurred to these entities. The pencil icon and blue background also appears for these entities in the table view. In the snapshot selection panel, use the Time Range and Snapshots in Range sliders to determine when the change occurred.

To distinguish them from tabs in the topology viewer page, tabs in the Configuration History page (Overview, Computers, Fabrics, Storage, and Other) have a light gray background and are outlined in orange. The minimap in the Configuration History page uses the following colors to indicate the aggregated change status of groups:

- **Blue** - One or more entities in the group have changed. Note that the addition or removal of an entity is considered a change.
- **Gray** - All of the entities in the group are unchanged.

Entities in the graphical view may be *active* (they existed at one or both snapshots) or *inactive* (not yet created or deleted). *Active entities* act like they normally do in the topology viewer; when you select them all relevant information also appears in the table view. You can adjust a grouping of active entities, but you may not perform actions that change the database, such as pinning. *Inactive entities* do not exist in the selected snapshots, but exist in the other snapshots. They are shown in a light gray background and do not have a change icon associated with them. Inactive entities display to keep the topology layout stable and to make it easier to

follow what has changed (instead of having the entities flicker in and out of existence when you change the snapshot selection). Inactive entities are not listed in the table view.

An entity that is moved from one group to another group appears only once in the new group in the graphical view. For example, if the health status of a computer has changed from Normal to Warning, the Configuration History page displays the computer as changed in the Warning health group (and no longer displays the computer in the Normal health group).

Note: In the Configuration History view, the performance and alert overlays are disabled and the minimap's shortcut to the Data Path Explorer is not available.

Viewing configuration changes in the table view

This section describes how to determine the changes that have occurred to a configuration over time by examining the icons and colors of the change overlay in the table view of the Configuration History page.

In the table view, a changed cell or tab is color coded. The change overlay uses the same colors and icons for the table view as for the graphical view (see Table 15 on page 267). The changed cell or tab displays both the old and new values as **Old Value → New Value**. To display a summary of change for the entire row, hover over the changed row. The summary displays the title of the column. It also displays the old and new value for each cell or tab that changed. For example, if a group contains one new entity but everything else in the group is unchanged, the group as a whole displays as "changed." If you click the group, the one new entity displays as "created" and all other entities display as "unchanged."

The Change column in the table view also indicates changed entities.

An object that is moved from one group to another group appears only once in the new group in the graphical view. For example, if the health status of a computer has changed from Normal to Warning, the Configuration History page displays the computer as changed in the Warning health group (and no longer displays the computer in the Normal health group).

Using configuration analysis

This topic describes how system administrators can determine whether an existing SAN configuration complies with predefined best practices (policies).

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

The configuration analysis feature allows a system administrator to select up to 13 policies and to specify the name of a target SAN (or zone set) for the analysis. The analysis tool reads the specified SAN data and the policy information for the specified policies from the IBM Tivoli Storage Productivity Center database. It outputs the policy violations as Tivoli Storage Productivity Center alerts. To view an alert, you can use one of three methods:

- Click the alert in the Configuration Analysis alert log
- Click the alert overlay icon (a red exclamation mark ) next to the affected entity
- Click the Alert tab in the table view of the topology viewer

System administrators can use configuration analysis to learn of any best-practices violations their actions may have caused. For example, in provisioning a new storage subsystem, an administrator may inadvertently violate a zoning-related policy. If the zoning-related configuration changes persist in the database, configuration analysis can detect this policy violation.

To perform configuration analysis, complete the following steps:

1. Ensure that you have previously run discovery and probe jobs for the computers, fabrics, switches, storage, and other objects of interest.
2. Ensure that the Tivoli Storage Productivity Center Data Server and Device server are running. To do so, in the Navigation tree pane, expand the following:
   - **Administrative Services → Services → Data Server → Server**
   - **Administrative Services → Services → Device Server → Device server**
3. Ensure that the following Tivoli Storage Productivity Center agents are installed and configured:

   Important: Policy violations are not raised if the agent information needed to evaluate a policy is not collected and stored in the Tivoli Storage Productivity Center database.
   - A mix of in-band, out-of-band, and CIMOM fabric agents to collect all fabric topology information (switches, interconnections, and zoning). This information is needed for the evaluation of all configuration analysis policies.
   - CIMOM agents for storage subsystems to collect the storage subsystem information (subsystem details, storage volumes, storage ports and storage volumes to hosts' mappings). This information is needed for the evaluation of configuration analysis policies 2, 3, 5, 7 and 11.
   - Data agents to collect the host information (operating system type and version). This information is needed for the evaluation of configuration analysis policies 3 and 9.
   - In-band fabric agents to collect the attribute information from the host information (HBA and host port details). This information is needed for the evaluation of configuration analysis policies 2, 3, 4, 8 and 9.
   - CIMOM agents for Tape to collect the tape information. This is needed for the evaluation of configuration analysis policy 2.
4. In the Navigation tree pane, expand **IBM Tivoli Storage Productivity Center → Analytics**.
5. Right-click **Configuration Analysis**. A selection list displays.
6. Click **Create Analyzer**. The Create Analyzer window opens (see Figure 28 on page 271).

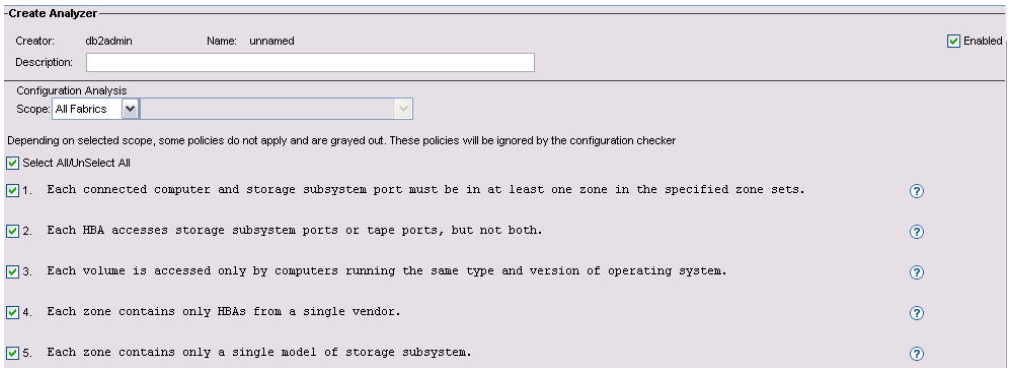

Figure 28. The Create Analyzer window. The figure shows a partial view of the window. Use the Create Analyzer window to select the scope and the policies of the configuration analysis.

7. Perform the following to define the analysis:
   a. Click the check box to display a check mark beside **Enabled**.
   b. In the **Description** field, type a brief description of the of the analysis job.
   c. To specify the scope of the SAN data to be checked, make a selection from the **Configuration Analysis Scope** list. Choices include:
      - All Fabrics
      - One fabric
      - One Zoneset

      If you select One fabric or One Zoneset, an adjacent list displays for you to click a specific fabric or zone set.
   d. Check the SAN data against up to 13 policies by performing one of the following:
      - Choose all of the policies by clicking the check box to display a check mark beside **Select All/UnSelect All**.
      - Choose one or more individual policies by clicking the check box to display a check mark beside each policy.

      Note: For Policies 12 and 13, type values up to a maximum of 9999.

   Table 16 lists the policies and their explanations. You can also access each explanation by clicking ⓘ at the right of the page.

Table 16. Policies used by the configuration analysis feature

Policy	Explanation
1. Each connected computer and storage subsystem port must be in at least one zone in the specified zone sets.	Determines whether an administrator forgot to zone a connected port. Putting connected ports into zones is useful for security and performance reasons. Ports are usually grouped into zones based on applications, server operating systems, or HBA vendors. The Fabric scope is not supported by this policy.
2. Each HBA accesses storage subsystem ports or tape ports, but not both.	Determines whether an HBA accesses both storage subsystem and tape ports. Because HBA buffer management is configured differently for storage subsystems and tape, it is not desirable to use the same HBA for both disk and tape traffic. The Fabric and Zone Set scopes are not supported by this policy because an HBA can be connected to multiple fabrics.

Table 16. Policies used by the configuration analysis feature (continued)

Policy	Explanation
3. Each volume is accessed only by computers running the same type and version of operating system.	Determines whether computers that run different operating systems access the same storage volumes. Use of the same volumes by computers that run different operating systems may corrupt the data that is stored on the volumes. This applies, regardless of whether the computers are in the same zone.
4. Each zone contains only HBAs from a single vendor.	Determines whether HBAs from different vendor types are in their own zone. Receiving a registered state change notification (RSCN) may cause an HBA to lose a zoned device, preventing the HBA from seeing or communicating with other devices in the zone. To avoid losing a zoned device, keep HBAs from different vendor types in their own zone.
5. Each zone contains only a single model of storage subsystem.	Determines whether different storage subsystems are in the same zone. While no technical problem is associated with storage subsystems from different vendors and of different models being in the same zone, an administrator may find them more difficult to organize. When similar storage systems are in the same zone, an administrator can easily group them for different applications.
6. Each zone is part of a zone set.	Determines the presence of orphan zones. Orphan zones are not associated with any zone set. They are not useful because their definitions are not used and they take up switch resources.
7. Each host must be zoned so that it can access all of its assigned volumes.	Determines whether the zones that were configured by the storage administrator allow each computer to access all of the storage volumes that are assigned to it. The administrator specifies the storage subsystem ports through which the computer port accesses volumes, but may forget to configure zones that enable the ports to communicate during volume assignment. This policy also determines whether zoning makes assigned volumes inaccessible to the computer ports. The Fabric scope is not supported by this policy.
8. Each computer has only HBAs of the same model and firmware version.	Checks whether there is only one type of HBA in each computer. Using only one type of HBA minimizes configuration problems. The policy also checks whether firmware upgrades have been done properly for all HBAs in a computer. The Zone Set scope is not supported by this policy.
9. For each host type and operating system, every HBA of a given model must have the same firmware version.	Determines whether all firmware upgrades have been done for the HBAs in the operating system. The Zone Set scope is not supported by this policy.
10. Every SAN switch of a given model must have the same firmware version.	Determines whether firmware upgrades have been done for all switches of the same type. For example, if you have four identical models of SAN switches from the same vendor and you perform a firmware upgrade on one, it is best to perform the upgrade on all of the others. The Zone set scope is not supported by this policy.

Table 16. Policies used by the configuration analysis feature (continued)

Policy	Explanation
11. Every storage subsystem of a given model must have the same firmware version.	Determines whether firmware upgrades have been done for all storage subsystems of the same type. For example, if you have four identical storage subsystems from the same vendor and you perform a firmware upgrade on one, it is best to perform the upgrade on all of the others. The Zone Set scope is not supported by this policy.
12. Each fabric may have a maximum of x zones.	Checks whether the number of zone definitions in the fabric is larger than the number that you entered. In large fabrics, too large a number of zone definitions can become a problem. Fabric zone definitions are controlled by one of the switches in that fabric, and limiting their number ensures that the switch's zoning tables do not run out of space. The Zone Set scope is not supported by this policy. You can enter up to a maximum of 9999 zones.
13. Each zone may have a maximum of x zone members.	Checks whether the number of zone members in a zone is larger than the number that you entered. In large fabrics, too large a number of zone members can become a problem. Fabric zone members are controlled by one of the switches in that fabric, and limiting their number ensures that the switch's zoning tables do not run out of space. You can enter up to a maximum of 9999 zones.

8. To specify how often to run the analysis, scroll to **Scheduling**.
9. In the **How often to run** area, schedule how often you want the job to run. Choose from the following actions:

   Run Now
   Run the job immediately.

   Run Once at
   Run the job once at a specified date and time. Use the list boxes to specify the month, day, and year. Type the hour and minutes over the existing defaults. Click AM or PM.

   Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

   Run Repeatedly
   Run the job repeatedly according to a schedule that you specify:
   - Indicate when to begin the job. Use the list boxes to specify the month, day, and year. Type the hour and minutes over the existing defaults. Click AM or PM.
   - Indicate how often to run the job. Select one of the following:
     – In the **Repeat Every** field, click the desired number of minutes, hours, days, weeks, or months.
     – In the **Run on these days** field, click to display a check mark next to the desired days.
10. In the **How to handle time zones** area, indicate the time zone to be used when scheduling and running the policy. Click one of the following:

Use the time zone that the server runs in
Select this option to use the time zone of the location where the Data server resides.

Use this time zone
Select this option to indicate a specific time zone for when to run the job. From the list box, select the time zone.

11. From the menu bar, select **File → Save**. The Save As pop-up opens.
12. Type a name for the analyzer and click **OK**. The Notice pop-up displays **Analyzer submitted**.
13. Click **OK** to begin the job. In the Navigation tree pane, the newly named analyzer job displays under **IBM Tivoli Storage Productivity Center → Analytics → Configuration Analysis**. The format of the name is *creator.analyzer*, where *creator* is the name that displays near the top of the Create Analyzer window and *analyzer* is name that you assigned in step 12 (for example *db2admin.analyzer2*).
14. Expand the job to view information about the run. The information includes the number of times that the job has been run, as well as the date and time of the run. The color of the type indicates the status:
    - Blue indicates that the run is in progress.
    - Green indicates that the run is finished.

    To refresh the status, right-click the job and click **Update Job Status**.
15. Click the run to view information about it. On a page titled with the name of the job and run number, the following information displays:

Field	Description
Run	Number of the run
Status	The status of the run (Success or Failed), not whether policy violations occurred
Start Time	The time at which the run started
Finish Time	The time at which the run ended
# Jobs	The number of jobs
# failed	The number of jobs that failed
Computer	The IP address of the computer on which the job was run
Status	The status of the run
Start Time	The time at which the run started
Finish Time	The time at which the run ended
Log File Name	The complete path name of the log file

16. To view whether a policy violation occurred, click  . The Job log file page displays the following information:
    - Scope of the run (all fabrics, one fabric, or one zone set)
    - Policies that were checked
    - Total violations
    - Breakdown of violations per policy
    - Whether the run completed successfully or the errors it encountered

    Note: To avoid overloading the Topology Viewer, the analysis job stores and displays only the first 50 policy violations. You must resolve these violations

and run the job again to view any remaining violations, or you can change this policy violation limit by using the **setdscfg** command. The attribute to set is **MaxPolicyViolationsStored**.

17. To view one or more alerts generated by policy violations, in the Navigation tree expand **IBM Tivoli Storage Productivity Center** → **Alerting** → **Alert Log** → **Configuration Analysis**. The Alert History - All Policy Violations page displays a log of job runs that generated alerts for policy violations. A policy violation alert is generated for each policy that was violated during a run. A policy may be violated several times, but only one alert is generated; the text in the alert indicates the number of times the policy was violated. The Alert History - All Policy Violations page includes the following information:

Column	Description
	When clicked, provides details about the alert.
	When clicked, makes the Clear and Delete buttons active. When double-clicked, provides details about the alert.
Computer	Not applicable (N/A)
Configuration Analysis	The name of the configuration analysis job
Alert Type	The type of alert generated during the run
First Triggered	The date and time of when the first alert was generated
Last Triggered	The date and time of when the last alert was generated
# Times	The number of times that an alert was generated for a run
Alert Creator	The userid of the person who created the run
Alert Name	The official name of the alert
Script	Not applicable (N/A)

For more information about the alert log, including the buttons at the top of the page, see "Viewing alert logs" on page 262.

18. To view details about each alert, in the Alert History - All Policy Violations page click the run that you want, then click  or double-click  . The Detail for Alert page displays the following information:

Field	Description
Configuration Analysis	The name and scope of the analyzer job.
Run	The number of run in which the alert occurred.
Alert Type	The type of alert (Policy Violation).
State	The status of the alert. An *active* (unresolved) state displays a red exclamation mark  ; an *acknowledged* (resolved) state displays a green check mark  .
Timestamp	The date and time that the alert was created. Expressed as the month, day, year, hour, minute, second, and AM or PM.
Alert Creator	The userid of the person who created the analyzer job.

Alert Name	The name of the alert.
Alert Text	A description of the policy that was violated and the number of times that it was violated during that analyzer run.

You can also use the topology viewer to view alerts about policy violations. To launch the viewer, click **Topology**. For more information, see "Viewing alerts" on page 259.

Resolving policy violations

When you click the Alert tab in the topology viewer and expand the entries in the table view, information displays about each configuration analysis policy violation. To verify and fix the policy violation, use the information in the following table.

Configuration Analysis Policy	Enclosing Entity	Affected Entity	How to Verify and Fix the Policy Violation
Policy 1. Each connected computer and storage subsystem port must be in at least one zone in the specified zone sets.	Computer, Subsystem, or Other	Port	In the zone tab within the L2:Fabric view, verify that the port is not a member of the zone set. Fix the violation by adding the port to the appropriate zone within the zone set.
	Fabric	Zone set	
Policy 2. Each HBA accesses storage subsystem ports or tape ports, but not both.	Computer	HBA	In the zone tab within the L2:Computer view, verify that HBA ports are accessing both the disk and tape subsystems. Fix the violation by configuring the computer access to disk and tape in two separate zone sets, and ensuring that only one of these is active at a given time.
Policy 3. Each volume is accessed only by computers running the same type and version of operating system.	Subsystem	Volume	In L1:Computers view, one can see that the computers accessing the volume do not have the same OS type and version. This violation can be fixed by mapping or assigning the volume only to the computers with same OS type and version.
	Computer	Port	
	Computer	Port	
Policy 4. Each zone contains only HBAs from a single vendor.	Fabric	Zone	In the zone tab within the L2:Fabric view, find the computers in the zone. Open the respective L2:Computer views and select the HBA tab in the table view to view the vendor of the HBAs on those computers. Fix the violation by editing the zone definition and ensuring that all the HBAs in the zone are from the same vendor.
Policy 5. Each zone contains only a single model of storage subsystem.	Fabric	Zone	In the L2:Fabric view, verify that the models of the subsystems in the zone are different. Fix the violation by editing the zone definition and ensuring that all the subsystems in the zone have the same model.
	Subsystem	Port	

Configuration Analysis Policy	Enclosing Entity	Affected Entity	How to Verify and Fix the Policy Violation
Policy 6. Each zone is part of a zone set.	Fabric	Zone	In the zone tab within the L2:Fabric view, verify that the zone does not belong to any of the zone sets and is listed under Orphan Zones. Fix the violation by adding the zone to a zone set or by deleting the zone.
Policy 7. Each host must be zoned so that it can access all of its assigned volumes.	Computer or Other	Port	None of the zones in the L2:Computer or L2:Subsystem views has both the computer and subsystem port. Fix the violation by adding or editing the zone definition, and by ensuring that both the computer and subsystem port are in the same zone.
	Subsystem	Port	
	Subsystem	Volume	
Policy 8. Each computer has only HBAs of the same model and firmware version.	Computer	Computer	Open the L2:Computer view and select the HBA tab in the table view to verify that the model and firmware version for the HBAs on that computer are different. Fix the violation by installing the HBAs of the same model or upgrading the firmware version of the HBA on that computer.
Policy 9. For each host type and operating system, every HBA of a given model must have the same firmware version.	Computer	HBA	Open the respective L2:Computer views and select the HBA tab in the table view to verify that the firmware and driver version for the HBAs on those computers are different. Fix the violation by upgrading the firmware and driver versions of the HBAs on those computers.
	Computer	HBA	
Policy 10. Every SAN switch of a given model must have the same firmware version.	Switch	Switch	Open the L1:Switches view and select the switch tab to verify that the firmware versions on those switches are different. Fix the violation by upgrading the firmware versions on those switches.
	Switch	Switch	
Policy 11. Every storage subsystem of a given model must have the same firmware version.	Subsystem	Subsystem	Open the L1:Subsystems view and select the subsystem tab to verify that the firmware versions on those subsystems are different. Fix the violation by upgrading the firmware versions on those subsystems.
	Subsystem	Subsystem	
Policy 12. Each fabric may have a maximum of *x* zones.	Fabric	Fabric	In the zone tab within the L2:Fabric view, verify that the number of zones present in the fabric is greater than *x* (the value entered by the user in this policy). Fix the violation by reducing the number of zones to a value less than *x*.

Configuration Analysis Policy	Enclosing Entity	Affected Entity	How to Verify and Fix the Policy Violation
Policy 13. Each zone may have a maximum of *x* zone members.	Fabric	Zone	Search for the zone in the zone tab in the L2:Fabric view. Verify that the number of zone members present in that zone is greater than *x* (the value entered by the user in this policy). Fix the violation by reducing the number of zone members to a value less than *x*.

Planning and modifying storage configurations

This section describes SAN Planner and how to use it to plan and modify changes to the system configuration.

SAN Planner overview

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

The SAN Planner consists of three panels.

1. The configuration panel contains:
   - The **Create Plan** pane.
   - The **Planner Selection** pane.
   - Three planners: **Volume Planner**, **Path Planner**, and **Zone Planner**.
   - Access to the selection panel with the **Add** button.
   - Access to the Recommendation panel with the **Get Recommendation** button.
2. The **Planner Selection Topology Viewer** panel contains:
   - The topology viewer in the graphical and table view of the configuration from which configuration elements can be selected.
   - The **Selected Elements** pane where the selected elements are moved into. These elements are used later by the planners during the planning and implementation stages.
   - Effective Provisioning Profile will list the single effective profile that is being used to pre-select the planner inputs. It will vary depending on the profiles associated with the storage resource groups input to the plan.
3. The recommendation panel contains:
   - The **Create Plan** pane.
   - The **Plan involves the following changes** pane displaying the planner recommendations.
   - The **When to run** task scheduler pane.
   - The **How to handle time zones** pane.
   - Four buttons:
     a. The **Return to Planner Input** button returns the user to the initial planner configuration panel.
     b. The **Save** button saves the plan for later modification.
     c. The **Execute Plan** button creates a task to execute the plan.

d. The **Cancel** button cancels the SAN plan effort and exits out of the SAN Planner.

Usage combinations for the three planners

You can use the three planners together or independently of one another. The following are examples of the combinations in which you can use the three planners:

1. Using all three planners provide end-to-end planning on fabrics (one or more) that will create the storage, zone the new storage to the host, and configure the multipath settings on the host.
2. Using Path and Zone Planners for cases such as adding a new path between the host and volume. Select a host and volume to be used as input to the Path and Zone planners.
3. Using Volume and Path Planners for cases where the host is already zoned to the storage subsystem and you want to add more volumes.
4. Using the Volume and Zone Planners for cases where additional paths are added between the switch and storage subsystem. Select a host to be used as input to Volume and Zone planners.
5. Using the Volume Planner by itself to create a best choice volume for later use.
6. Using the Path Planner by itself to modify existing multipath settings on a host. Select an assigned volume to be used as input to the Path planner.

Supported Subsystems

The planners support the Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 subsystems. Unsupported subsystems cannot be selected by the user.

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 then the planner will not be able generate the plan or execute portions of the plan and issues an error message.

For volume creation,Tivoli Storage Productivity Center must be managing the subsystem. You need an active CIMOM to do a CIMOM discovery and a subsystem probe.

For volume 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 zoning planner can be used to configure the zoning.

For a zoning configuration, Tivoli Storage Productivity Center must manage the fabric. For Brocade fabrics, use the Storage Management Initiative Specification (SMI-S) agent to obtain zone information from Brocade SAN. For an out-of-band fabric agent, if you select the **Advanced Brocade Discovery** check box a message informs you that by selecting this option probe failures might occur. The system prompts you to respond. Yes or No. If your response is no, the message dialog is closed and the **Advanced Brocade Discovery** check box is cleared. If your response is yes, the **Advanced Brocade Discovery** check box remains selected and can enter a user ID and password. For other fabrics, you must connect an in-band fabric agent to the fabric for zone control operations.

A storage controller performance monitor must be run in order to select the **Workload Profile** options. If a performance monitor has not been run, the **Space Only** workload profile option is allowed for planning on 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.

Creating a SAN plan

You can create a SAN plan by right clicking the **IBM Tivoli Storage Productivity Center->Analytics->SAN Planner** node and selecting **Create Plan**. The planner configuration panel is displayed.

You can also create a SAN plan from the **IBM Tivoli Storage Productivity Center->Disk Manager->SAN Planner** node and selecting **Create Plan**. The planner configuration panel is displayed.

The SAN Planner Configuration panel

The Planner selection pane displays a tree view of the hosts, fabrics, storage resource groups and subsystems which are used during the planning process. This pane is populated in 2 ways:

1. Clicking the **Add** button on the **Planner Selection** pane. This brings the user into the Planner Selection **Topology Viewer** panel. This viewer displays the configuration and what has been selected for the planning. The selections are moved from the graphical view into the **Select Elements** pane. Click **OK** to return to the **Planner Selection** pane on the configuration panel. The pane displays the selected elements.
2. Through the **Launch Planner** selection from the Tivoli Storage Productivity Center Topology Viewer. In the graphical view of the **Topology Viewer**, highlight a supported element (or several), right click, and select **Launch Planner** from the context menu. (The supported elements for **Launch Planner** are fabrics, storage, storage pools, storage volumes, and computers). The selection is placed into the **Planner Selection** pane of the SAN Planner configuration panel. If a SAN plan has not been created, clicking **Launch Planner** opens a new SAN Planner configuration panel and places the selection into the **Planner Selection** pane. If another SAN plan already exists and is active, clicking **Launch Planner** opens a new SAN Planner configuration panel and places the selection into the **Planner Selection** pane. The already existing SAN plan is not updated unless you click the **Add** button on that panel.

Both ways provide a topology viewer from which you can select a variety of elements. The elements allowed for selection are fabrics, storage subsystems, storage pools, volumes, computers, and storage resource groups. You can select a storage subsystem, not select any storage subsystems, or select a storage subsystem and one or more storage pools. You might also select one or more fabrics for use in path and zone planning.

Volume Planner

Use the Volume Planner to plan and select appropriate storage controllers, storage pools, and storage volumes (when using unassigned volumes) that satisfy your inputs. It allows you to select controller type preference,

whether storage request can be satisfied by multiple controller types, and select RAID level. VPA 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 Volume Planner will use 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.

Select workload profiles an inputs to the Volume Planner's analysis.

Path Planner

The Path Planner allows setup of multipath options. The supported multipath driver on a host is the IBM Subsystem Device Driver (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 Path Planner enables system administrators and SAN administrators to plan and implement storage provisioning for hosts and storage subsystems 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 subsystems are supported by the host multipath driver and the multipath modes which are supported by the driver and storage subsystem. Planning the paths between the host(s) and storage controller requires designing paths between hosts and storage subsystems which will be implemented through zones in the fabric.

The Path Planner is used for specifying multiple paths options between selected hosts and storage subsystems. This setting requires each host to have at least two fibre channel ports connected to the SAN. Path Planner assists the administrator in the multipath tuning process through the selection of these policies:

- The **Multipath Option** specifies how the driver uses the paths between the host and the storage subsystem.

  Note: For DM_Multipath, only the round robin option is supported.

  The options are:
  - **Load Balancing** sends Input/Output on all paths
  - **Round Robin** sends Input/Output on one path until a time interval expires (set in an SDD setting at the host) or stops to use another path.
  - **Fail-Over** sends Input/Output on one path until a failure occurs and fails over (switches) to another path.
- **Specify number of paths** option specifies the number of paths between each host and the storage subsystem.
- **Use fully redundant paths** option causes the Path Planner to check for redundant fabrics between each host and storage subsystem and to create paths in each fabric. This requires at least two fabrics.

The Path planner does not directly interact with the Zone planner. It provides the path information which the Zone planner uses. Each path is represented by a host port WWPN, target port WWPN, and a Volume ID for the volume on the target which is mapped to the host port. These paths are created when the Zone planner is implemented.

Zone Planner

Zone Planner specifies automatic zoning between ports on the selected hosts and subsystems in a fabric. All zoning is done with WWPN based zoning. It plans zoning configuration for new storage that is provisioned for use by a host. For example, it can be used when a new storage volume is created and assigned to a host. It can also be used with a volume that has already been created and is assigned to a host needing more storage. In these cases, the Path and Volume Planner determines which host and storage need to be zoned together, providing the Zone planner the exact set of ports which need to be zoned together. The Zone planner then uses the zoning inputs for the planning. In the case where Volume and Path planners are not used, manually select the host and storage ports and then invoke the Zone planner.

The Zone planner expects a list of host port and storage port pairs as input. If the Path planner has been invoked prior to the Zone planner, its output is used as input to the Zone planner. If the subsystem/host are within the same fabric and **Zone Planner** is not checked, then existing zones or zone sets are used. If **Zone Planner** is checked, this creates a new zone or zone set.

In order to do any zoning, the host and subsystem being zoned by the Zone planner must be within the same fabric. For the case where the host and subsystem reside in more than one of the same fabrics, you are given two options. The first option is to create identical zones in all of the fabrics. The second option is to select specific fabrics to create identical zones in.

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:

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

Policy two in the preceding list has a higher priority than policy three.

Planner and user inputs interdependencies

The planner options depend on the inputs you select. For example, selecting a host that doesn't have an agent installed would limit the kinds of planning that could 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 will make 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.

Printing

To print plan information, click**File → Print**, or click the **Print data** icon. On the planner configuration panel, the print output includes the selected

hosts, fabrics, storage subsystems, and planner options. On the recommendation panel, the print output includes the details of the recommended plan implementation.

Using the SAN Planner

This topic describes how to use the SAN Planner to select configuration information for a plan and modify a configuration.

You must run a fabric, disk and data discovery, and subsystem and fabric configuration probe on storage controllers that will be used in planning. You must also run, a performance monitor on the storage controllers . The SAN Planner is available for IBM Tivoli Storage Enterprise Storage Server, IBM System Storage DS6000, and IBM System Storage DS8000 systems. Unsupported subsystems cannot be selected.

A data agent is required to obtain host information used in volume assignments. An inband fabric agent is optional if the fabric information is available through another inband agent in the same fabric, an outbound agent, or through a CIMOM managed switch. Run the inband host agents to get the host operating system information. If the inband host agent has not been run and the storage controller does not have information about the host operating system, an asterisk is displayed next to the host port. Ensure that the operating system-imposed limits on the number of volumes that can be accessed are not exceeded.

1. Expand the **IBM Tivoli Storage Productivity Center → Analytics** or **Disk Manager** node in the navigation tree.
2. Right-click on **SAN Planner** and select **Create Plan**. The Configuration panel opens. The Configuration panel contains five panes: Create Plan, Planner Selection, Volume Planner, Path Planner, and Zone Planner. This panel provides the interface for user inputs to put together a plan for volume, path, and zone planning. It also provides the ability to get configuration recommendations based on the current storage configurations and the user inputs.

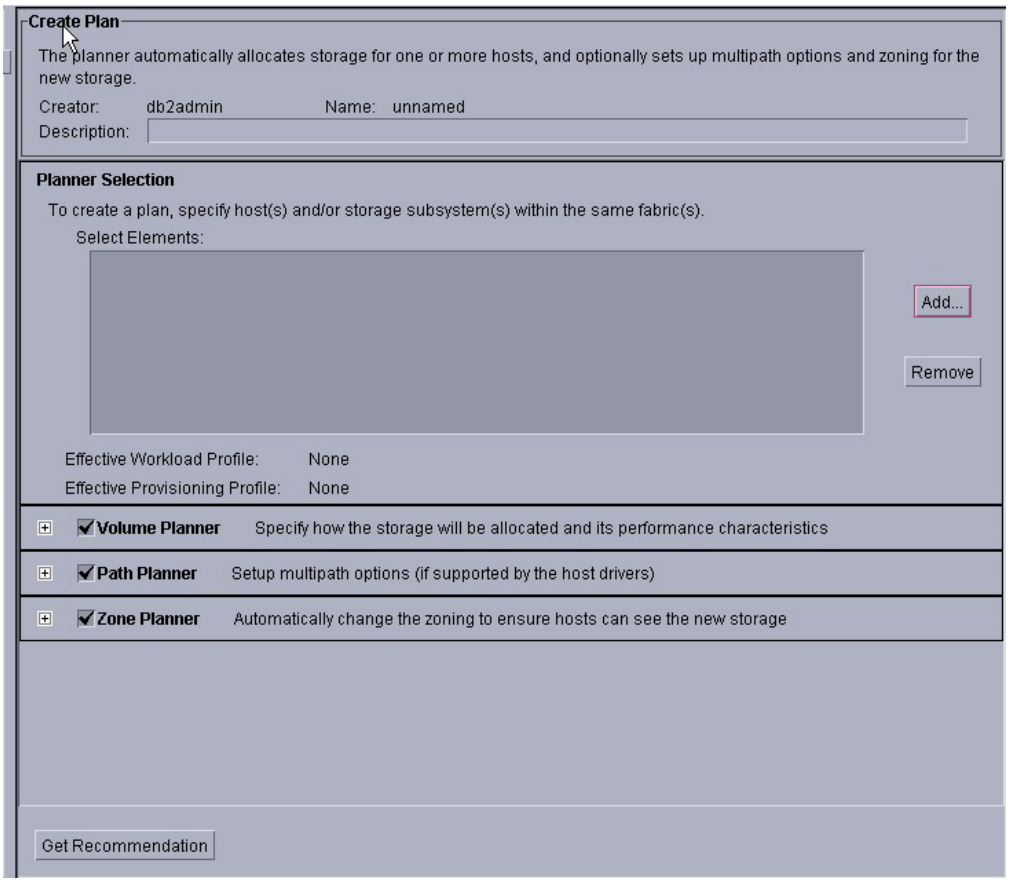

Figure 29. SAN Planner Configuration Panel

Note: Clicking on a saved plan (located under **SAN Planner** in the Navigation tree pane) opens the Configuration panel with the **Edit Plan TPCuserid.plan_name** pane instead of the Create Plan pane, where **TPCuserid** is the Tivoli Storage Productivity Center userid and **plan_name** is the name the plan was given when it was previously saved.

3. Optional: In the **Create Plan** pane, enter a short description of the plan into the **Description** field.
4. The **Planner Selection** pane lists any configuration elements that have been added to the plan. You can add configuration elements to the plan by either using the **Add** button in this pane (see step5) or through **Topology** (see step 6 on page 286
5. To use the **Add** button to add elements to the plan:
   a. Click the **Add** button. The **Planner Selection Topology Viewer** panel opens. This provides a topology view of the current system configuration. Select the elements to be used for consideration in the SAN planning. Select the storage resource groups, storage controllers, hosts, fabrics, and storage pools to be used for volume provisioning considerations. Currently, only Tivoli Storage Enterprise Storage Server, DS6000, or DS8000 storage controllers can be selected.

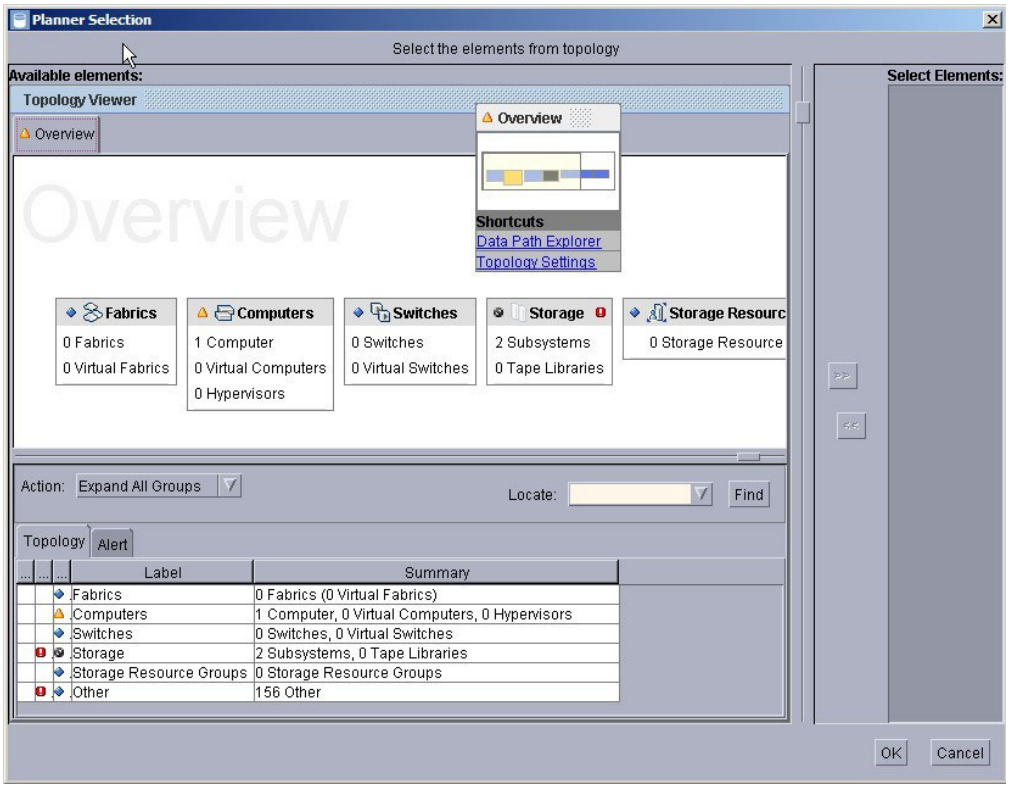

Figure 30. SAN Planner Topology Viewer panel

Note: The Topology Viewer pane displays the configuration, the plan selections, and the configuration from which selections are made. The Select Elements pane displays the selections that will be used in the plan.

b. Double click the box title for a subsystem to expand the storage view into the L0 level view of the available subsystems. Before selecting subsystems, you might want to view the relationship between the subsystems and switch connectivity. To do this, double click on the Subsystems header at the L0 Storage tab view. This opens the L1 Subsystems view, which displays the Subsystems and the Connectivity (Switches) boxes. Double clicking on a specific subsystem displays the L2 view of that subsystem and everything it is connected to.

c. Click any supported element to select it for the plan. To select multiple elements, press and hold the **CTRL** key and click on each element icon.

d. Click >> to move the selected elements into the **Select Elements** pane. If you decide to remove the selected elements from the **Select Elements** pane, click <<.

Note: The **Select Elements** pane lists the storage resource groups, fabrics, computers, subsystems, pools, and volumes if selected. All selections must be located within the same fabric for SAN Planner to work. The **Fabrics** section lists the fabrics (by WWN) and the corresponding selected subsystems and selected hosts within each fabric. A subsystem is listed multiple times if it has been configured in different fabrics. The **Subsystems** section contains the selected storage subsystems. The number of fabrics the subsystem belongs to is shown next to each subsystem. Selected pools are listed under their storage subsystems. Selected volumes are listed under their pools. The **Computers** section lists the selected hosts

and the number of fabrics the host belongs to. The hosts is used in path planning and zone planning if **Path Planner** and **Zone Planner** are selected in the configuration panel.

e. When you are satisfied with all your selections, click **OK**. You return to the SAN Planner Configuration panel with the selections displayed in the **Planner Selection** pane. Continue with the instructions at step 7

6. To use the**Topology Viewer** pane to add configuration elements to the plan:
   a. Expand **IBM Tivoli Storage Productivity Center → Topology** and click on **Topology**.
   b. Double-click the box title for a subsystem to expand the storage view into the L0 level view of the available subsystems. Before selecting subsystems, you might want to view the relationship between the subsystems and switch connectivity. To do this, double click on the Subsystems header at the L0 Storage tab view. This opens the L1 Subsystems view, which displays the Subsystems and the Connectivity (Switches) boxes. Double clicking on a specific subsystem displays the L2 view of that subsystem and everything it is connected to.
   c. In the L0 view, click each SAN Planner supported element you want to include in the SAN plan. To select multiple elements, press and hold the **CTRL** key and click on each element icon. Right click to open the pop-up menu and select **Launch Planner**.
7. When you have returned to the configuration panel, the **Planner Selection** pane will contain the selected elements. If you decide to remove an element from the plan, click on the item in the **Planner Selection** pane and click the **Remove** button. The element is no longer listed in the **Planner Selection** pane and will not be used for the SAN plan.

   Note: To save your configuration inputs to review later, click the **Save** icon on the **Tivoli Storage Productivity Center** toolbar. The **Planner Selection** pane will contain all the elements as they were at the time the plan was saved, even if any hardware configurations were modified.
   - If an element appearing in the **Planner Selection** pane has been physically removed, clicking the **Get Recommendation** button causes an error. To remove the element, select it in the **Planner Selection** pane and click the **Remove** button.
   - If more fabrics, storage subsystems, storage pools, volumes, or computers were added, they are taken into consideration when you click the **Get Recommendation** button. If they meet your input requirements, they will be included in the recommendation panel.
8. To start the volume planning, click **+** to expand the **Volume Planner** options. If you do not want to do volume planning, clear the **Volume Planner** check box to disable the fields within the pane.

   Note: The **Volume Planner** uses your inputs to select appropriate storage controllers, storage pools, and storage volumes when using unassigned volumes.
   a. Enter the total capacity of the volumes to be used for provisioning in the **Total Capacity** field.
   b. Click **Divide capacity between** to specify how to divide the **Total Capacity** among volumes. For example, enter **1 and 1 volumes** if you want the total capacity to be on one volume. Enter **1 and 5 volumes** if you want the total capacity to be divided among one to five volumes.

c. Click **Divide capacity among volumes of size** to specify the size of the volumes. For example, you can keep the capacity of each volume between **x.x GB and y.y GB**.

d. For **Workload Profile**, select the workload profile that represents how the new volumes will be used.

   Note: All of the workload profile options except **Space Only** require a performance monitor to have been run on the storage controller.
   The predefined workload profiles are as follows:

   OLTP Standard
   For typical online transaction processing.

   OLTP High
   For very active online transaction processing.

   Data Warehouse
   For applications with inquiries into large data warehouses.

   Batch Sequential
   For batch applications involving large volumes of data.

   Document Archival
   For document archival applications.

   Space only
   When making the pool selection, the consideration is made towards storage pool space information only on a storage controller. It is not based on storage controller performance data. Use this option for a storage subsystem that has not been monitored for a performance data collection.

   You can create your own **Workload Profile** by clicking the **Disk Manager → Profile Management → Workload Profiles** node. The selection list is the same as listed above.

e. For **RAID Level**, select the possible RAID levels on the available storage subsystems. The predefined RAID levels are as follows:

   <system selected>
   This is the default value. The best possible RAID level of the volume will be selected based on user input and the available performance and capacity of the underlying storage subsystems.

   RAID 1
   The format type of the volume will be RAID 1 format.

   RAID 5
   The format type of the volume will be RAID 5 format.

   RAID 10
   The format type of the volume will be RAID 10 format.

f. For **Volume Name Prefix**, enter a name which will be used as the prefix in the volume name.

g. Select the check box for **Use existing unassigned volumes (if available)** if you want to use existing unassigned volumes.

   Note: An unassigned volume must meet the performance requirements selected in **Workload Profiles** and **RAID Level**.

h. Click the **Suggest Storage Pools** button to obtain a selection of storage pools.

Note:

- If you specified storage controllers in the Planner Selection pane, the SAN Planner lists the storage pools from those storage controllers that have the storage capacity you requested.
- If you specified storage controllers and pools in the Planner Selection pane, the SAN Planner lists additional pools from storage controllers that have the storage capacity you requested.

Suggest Storage Pools selects a set of storage pools to create the volumes from. **Suggest Storage Pools** cannot be used if volumes were selected during the selection process and are visible in the Planner Selection pane. Storage pools are used if they are not full or are not visible from all hosts. Only fixed block storage pools can be used by the SAN Planners.

9. Click the **Get Recommendation** button if you want the SAN Planner to select one or more volumes based on your entries. Follow the instructions in step 12 on page 289

   Note:

   - If you did not select any storage pools, the SAN Planner considers all the pools on the selected storage subsystems.
   - If you selected any storage controllers and storage pools, the SAN Planner considers only those pools visible in the Planner Selection pane.

10. You are ready to start the path planning. Click **+** to expand the **Path Planner** options. If you do not want to do any path planning, clear the **Path Planner** check box.

    Note: Based on your inputs, **Path Planner** determines the appropriate number of paths that are necessary from the host to the storage subsystem. The Path Planner allows setup of multipath options if supported by the host drivers. Currently, the IBM Subsystem Device Driver (SDD) is the only supported multipath driver and must be installed on the hosts. After you select the **Path Planner** options, you can continue with the **Zone Planner** selections or click the **Get Recommendation** button. The **Get Recommendation** button causes the SAN Planner to analyze the current configuration, performance values, and storage space to recommend a potential configuration change. You can accept the plan or change the inputs and get another plan recommendation.

    a. Select the **Multipath Option** to determine how IO will be distributed across all paths. The predefined IO options are as follows:

       Load Balancing
       Sends IO on all paths.

       Round Robin
       Sends IO on one path until a time interval expires (set in an SDD setting at the host), then switches to send IO on another path.

       Fail-Over
       Sends IO on one path until a failure occurs, then fails over (switches) to another path.

    b. Select the **Specify number of paths** check box to enter the number of paths you wish to configure. Enter the number of paths in the box on the right.

    c. Select the **Use Fully redundant paths (requires 2 fabrics)** check box to use the paths from host to storage subsystems through at least two fabrics.

d. If you are not going to continue with **Zone Planning** then go to step 12 to click on the **Get Recommendation** button so that SAN Planner can start the plan using your inputs.

11. You are ready to start the zone planning. Click **+** to expand the **Zone Planner** options. If you do not want to do any zone planning, clear the **Zone Planner** check box.

    Note: Based on your entries, the **Zone Planner** allows the zoning to be changed to ensure that hosts can recognize the new storage. After you select the **Zone Planner** options, click the **Get Recommendation** button. The **Get Recommendation** button causes the SAN Planner to analyze the current configuration, performance values, and storage space to recommend a potential configuration change. You can accept the plan or change the inputs and get another plan recommendation.

    a. Select the **Automatically create zone** to indicate where zoning will be done. The predefined options are as follows:

       <auto-zone>
       This is the default. The plan is generated creating the maximum number of zones without grouping the data paths based on host, host ports, or host bus adapters (HBAs).

       ...for each host
       Creates a zone for each host.

       ...for each HBA
       Creates a zone for each HBA.

       ...for each host port
       Creates a zone for each host port.

    b. Select the **Specify maximum number zones** check box. Enter the maximum number of zones in the box to the right.

    c. Select the **Specify maximum zone members per zone** check box. Enter maximum number of zone members (per zone) in the box to the right.

    d. If you want any zone set that is available to be selected, select the **Use active zone set** check box.

    e. If you want to set a prefix name for each zone, select the **Append zone name prefix** check box and enter a zone name prefix in the box to the right.

12. Click the **Get Recommendation** button to validate the settings. If there are no errors, a plan is generated and the recommended configuration is displayed in a new panel. The Planner Recommendation panel displays the plan recommendations once the planner takes all the user inputs into consideration. It also provides the user interface to specify when the plan will be run and the time zone preference for the task execution.

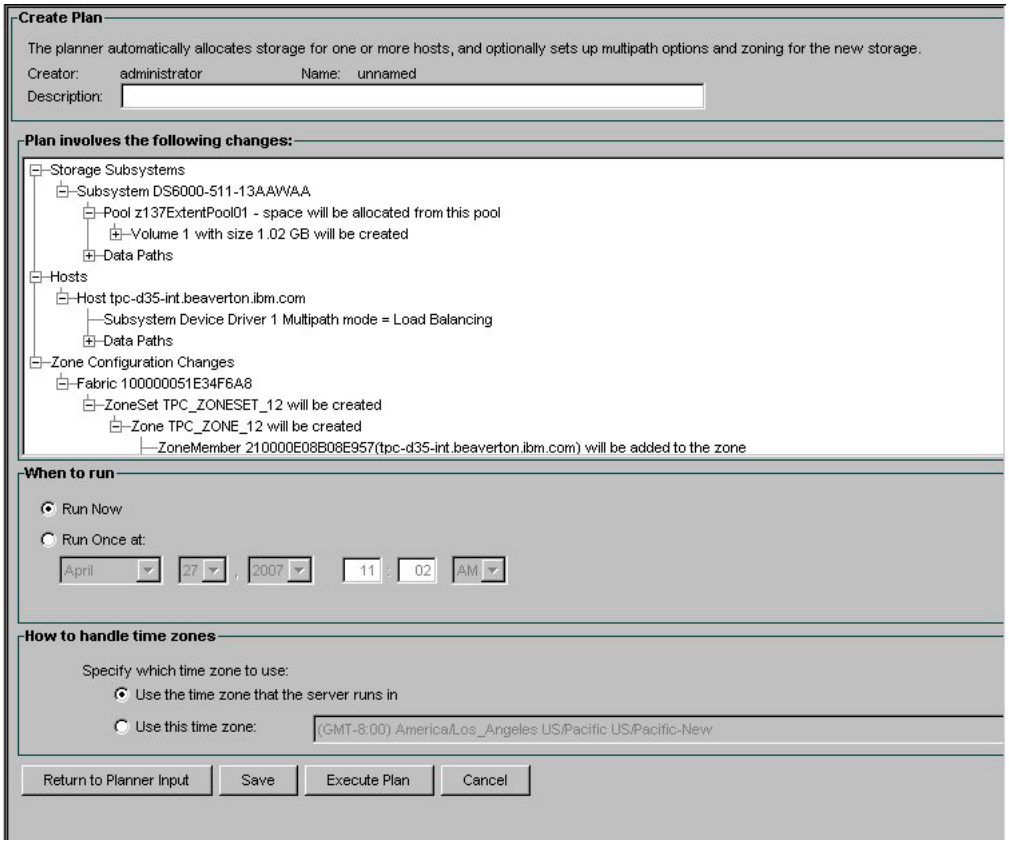

Figure 31. Planner Recommendation panel

a. Optional: In the **Create Plan** pane, enter a description in the **Description** field.

b. The pane titled **Plan involves the following changes** contains the planner recommendations. Review the configuration.

c. If you are satisfied with your selections, go to the **When to run** pane and click either **Run Now** to start the task immediately when it will be submitted or click **Run Once at** to start the task at the date and time you select.

d. The Tivoli Storage Productivity Center server and the console can be in different time zones. The **How to handle time zones** pane controls running the task at the specified time zone. Click on **Use the time zone that the server runs in** if you want to use the time zone where the Tivoli Storage Productivity Center server resides. Click on **Use this time zone** and select the time zone from the list box. This will start the task at the specified date and time entered in the **When to run** pane and use the time zone selected here.

e. Click the **Execute Plan** button to save the plan and execute it. This starts the job at the date, time, and time zone specified in the **When to run** and the **How to handle time zones** panes. The executed task will have a job status with the ability to view the job log(s).

f. Click the **Save** button to save the plan for later inspection. Enter a plan name in the **Specify Plan name** field. This creates and saves a task under **SAN Planner**. To return into this plan, go to **SAN Planner** and click on the task. This will be displayed with all the selections that were made prior to saving.

Note: Tasks are typically saved and executed by clicking the Save icon on the Tivoli Storage Productivity Center toolbar. For SAN Planner, clicking the Save icon saves the plan but does not execute (run) the task. Clicking on the Save icon is the same as clicking the **Save** button on the planner recommendation panel.

If hardware configurations were modified before you reactivate the saved plan, the **Plan involves the following changes** pane lists all the elements as they were when the plan was saved.

- If an element appearing in the **Plan involves the following changes** pane was physically removed from the system, continuing with **Execute Plan** causes an error. Click the **Return to Planner Input** button to return to the configuration input panel. Select the element in the **Planner Selection** pane and click on the **Remove** button. Click the **Get Recommendation** button to refresh the recommendation panel with the latest selections.
- Any fabrics, storage subsystems, storage pools, volumes, or computers that were added to the system are taken into consideration when you click the **Get Recommendation** button. If the additions meet your input requirements, they are included in the recommendation panel.

g. Click the **Return to Planner Input** button to adjust the plan inputs. You return to the Configuration panel.

h. If you decide to completely stop planning and discard your inputs, click the **Cancel** button. You return to will be taken completely out of the SAN Planner and return to the location from where you entered the SAN Planner.

Optimizing storage configurations

This section describes how to use the Storage Optimizer to analyze storage subsystems to identify performance bottlenecks and create a recommendations report that describes improvements you can make to subsystem performance, as well as recommendations for migration and consolidation.

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
- DS4000
- Tivoli Storage Enterprise Storage Server
- SAN Volume Controller

Note: Non-IBM subsystems, including SAN Volume Controllers that use non-IBM back-end subsystems, are not supported by Storage Optimizer. Solid state drives (SSDs) are not supported at this time. Subsystems that contain solid state drives may produce unexpected results in the Storage Optimizer heat maps. Storage Optimizer analysis of subsystems containing solid state drives should be avoided until support is available.

Before running Storage Optimizer, you must set up performance monitors and collect performance monitoring data for all storage subsystems that you want Storage Optimizer to analyze. You must also collect performance monitoring data for a SAN Volume Controller's back-end subsystems in order to produce the most accurate Storage Optimizer analysis.

The analysis will be 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. See "Creating storage subsystem performance monitor" on page 242 for more information.

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. See "Creating an analysis report" for more information.
2. Review the Storage Optimizer analysis report. See "Viewing an analysis report" on page 294 for more information.
3. Create an optimization report based on this analysis. See "Creating an optimization report" on page 296 for more information.
4. View the optimization report. See "Viewing an optimization report" on page 297 for more information.
5. At your discretion, implement the recommendations suggested in the optimization report, or re-run the analysis using different inputs.

Using the Storage Optimizer

This topic describes how to use the Storage Optimizer to analyze storage subsystem performance and create an optimization report that includes recommendations to improve performance.

Before running Storage Optimizer, you must set up performance monitors and collect performance monitoring data for all storage subsystems that you want Storage Optimizer to analyze. Storage Optimizer analyzes performance monitoring data for the time interval you specify. The analysis will be less accurate if there are any gaps in the data collection for the time interval being analyzed. For best results, 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. See "Creating storage subsystem performance monitor" on page 242 for more information.

Tip: The Storage Optimizer does not 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.

Creating an analysis report

Creating a performance analysis report is the first step in using Storage Optimizer.

To create the analysis, complete the following steps:

1. Navigate to one of the following nodes:
   - **IBM Tivoli Storage Productivity Center > Analytics**
   - **IBM Tivoli Storage Productivity Center> Disk Manager**.
2. Right-click **Storage Optimizer** and select **Create Analysis** to open the Create Analysis panel. Use this panel to specify the storage subsystems that you want to analyze.
3. Optional: In the **Create Analysis** panel, enter a short description of the analysis job in the **Description** field.
4. The **Select Storage** pane lists the storage subsystems that you want to include in the analysis. Use the **Add** button to add elements to the plan:

a. Click **Add**. The **Optimizer Selection** panel opens. This provides a topology view of the current system configuration. Select the storage subsystems that you want to include in the analysis.

   Tip: The Topology Viewer pane displays the configuration from which selections are made. The Select Elements pane displays the selections that will be used in the analysis.

b. Double-click the Storage box title to expand the storage view into the L0 level view of the available subsystems.

c. Click an element to select it. To select multiple elements, press and hold the **CTRL** key and click on each element icon.

d. Click **>>** to move the selected elements into the **Select Elements** pane. If you decide to remove the selected elements from the **Select Elements** pane, click **<<**.

   Note: The **Select Elements** pane lists the subsystems to be included in the analysis. All selections must be located within the same fabric. The **Fabrics** section lists the fabrics and the corresponding selected subsystems and selected hosts within each fabric. A subsystem is listed multiple times if it has been configured in different fabrics. The **Subsystems** section contains the selected storage subsystems. The number of fabrics the subsystem belongs to is shown next to each subsystem. Selected pools are listed under their storage subsystems. Selected volumes are listed under their pools. The **Computers** section lists the selected hosts and the number of fabrics the host belongs to.

e. When you are satisfied with all your selections, click **OK**. You return to the Storage Optimizer Create Analysis panel with the selections displayed in the **Select Storage** pane.

5. Optional: select an element to remove it by clicking **Remove**.
6. In the **Performance Time Interval** section, select the **Start Date** and **End Date** to define the time interval that the Storage Optimizer uses to extract performance monitoring data from the IBM Tivoli Storage Productivity Center database. Performance monitoring data must exist in the database for the entire time interval you specify. The default time interval is the start and end date of the current date.
7. Click **Run Now** to begin the analysis as soon as you save the analysis job using **File** > **Save**. Alternatively, click **Run Once at** to begin the analysis at a time you specify. When you save the analysis job using **File** > **Save**, the analysis job appears under the **Storage Optimizer** node in the navigation tree, but the job status is not displayed until the job has started running.
8. For **How to handle time zones**, specify which time zone to use.
9. Click **File** > **Save** to save and submit the analysis report job. After the analysis report job starts, you can view the analysis job status under the **Storage Optimizer** node in the navigation tree.
10. When the analysis report job is complete, select it to view the analysis report.

Viewing an analysis report

The analysis report displays a performance analysis for storage systems, and lets you generate an optimization report that lists storage migration and consolidation recommendations. You can experiment with different migration and consolidation scenarios to help you achieve the desired performance improvements for your storage infrastructure.

To view the analysis report:

1. Navigate to one of the following nodes:
   - **IBM Tivoli Storage Productivity Center > Analytics**
   - **IBM Tivoli Storage Productivity Center> Disk Manager**.
2. Expand the **Storage Optimizer** node and select a completed (green) analysis job to display the analysis report. By default, the heat maps display the performance utilization of all components.
3. To display the performance utilization for a specific component, select the component from the **Heat map based on** list.
4. Drag the **Performance Threshold** slider to change the performance threshold for all the selected subsystems. The heat maps are automatically updated to reflect changes in the performance threshold. This allows you to see the performance threshold at which different subsystems or components become bottlenecked.

There is one heat map for each storage subsystem included in the analysis. The heat map uses different colors to represent the actual performance of a storage subsystem, as measured against the default performance threshold of 80%. If you change the performance threshold, the heat maps are automatically updated to display the changes.

Each cell represents a storage pool. If there aren't enough storage pools to fill an entire row, the remaining cells in that row will be gray. Select a cell to display more information about that storage pool. Green, blue, yellow, orange, and red cells indicate storage pool performance as follows:

Green storage pools that are performing at less than or equal to 25% of the performance threshold.

Blue storage pools that are performing at less than or equal to 50% of the performance threshold.

Yellow
storage pools that are performing at less than or equal to 75% of the performance threshold.

Orange
storage pools that are performing at less than or equal to 100% of the performance threshold.

Red storage pools that exceed the performance threshold.

The Performance tables represent the same information that is displayed in each heat map, but in more detail. There is one table for each storage subsystem. Each row represents a storage pool. You can sort the table by clicking a column. Column descriptions are as follows:

Pool Name
the name of the storage pool.

Utilization
the aggregated performance utilization of all components, as measured against the performance threshold.

Host Adaptor
the performance utilization of the host adapter ports on the storage subsystem.

Hard Disk
the performance utilization of the disk arrays.

Controller
processor and memory performance utilization.

Device Adapter
the performance utilization of the device adapter which connects the controller to the disk arrays.

Space The physical disk space utilization.

Creating an optimization report

You can create an optimization report that lists storage migration and consolidation recommendations. You can experiment with different migration and consolidation scenarios to help you achieve the desired performance improvements for your storage infrastructure.

To create the optimization report, follow these steps:

1. Navigate to one of the following nodes:
   - **IBM Tivoli Storage Productivity Center > Analytics**
   - **IBM Tivoli Storage Productivity Center> Disk Manager**.
2. Select a completed analysis job to display the analysis report.
3. To include a storage pool or storage subsystem as a source entity that you want to migrate or consolidate to improve performance:
   a. In the heat map or table, select a storage pool or storage subsystem. Select **Select all pools for storage subsystem** to select all storage pools in a storage subsystem so that you can include them as source entities.
   b. Click >> to the left of **Source Entities** to add the selected item to the source entities. The same storage pool cannot be added as both a source and target entity.
   c. For the list of source entities, if you select **Make retirement recommendations for the selected entity**, the optimization report will make retirement recommendations for all source entities.
   d. To remove an item from the list of source entities, select it and click << to the left of **Source Entities**.
4. To include a storage pool or storage subsystem as a target entity that you want to keep after storage migration and consolidation:
   a. In the heat map or table, select a storage pool or storage subsystem. Select **Select all pools for storage subsystem** to select all storage pools in a storage subsystem so that you can include them as target entities.
   b. Click >> to the left of **Target Entities** to add the selected item to the target entities. The same storage pool cannot be added as both a source and target entity.
   c. To remove an item from the list of target entities, select it and click << to the left of **Target Entities**.

   When the Storage Optimizer considers potential pools as targets of a migration, it takes into consideration whether the source and target pools have the same RAID level, format (CKD or FB), and whether or not the potential target has enough available capacity.
5. Choose when to run the optimization report using one of the following options:
   - **Run Now**: creates the optimization report as soon as you save your report settings by clicking **File** > **Save**. After the report job begins, you can view the report and job status under the **View Previously Run Optimization Reports** pane.

- **Run Once at**: begins the report job at a specified date and time. When you save the report job by clicking **File** > **Save**, the report job is displayed under the **View Previously Run Optimization Reports** pane, but the job status will not appear until the job has started running.

6. Drag the **Performance Threshold** slider to change the performance threshold for the selected subsystems and pools. The heat maps are automatically updated to reflect changes in the performance threshold. This allows you to see the performance threshold at which different subsystems and pools become bottlenecked. Keep in mind that if you select the **Make retirement recommendations for the selected entity** option, the threshold you choose using the **Performance Threshold** slider determines if the move to the target subsystem will keep the target subsystem below the chosen threshold.
7. Click **File** > **Save** to save and submit the optimization report job.
8. To update the job status for all optimization reports, click **Refresh Job Status**.
9. To view a completed optimization report, click  .
10. To see how the optimization recommendations change if you use a different performance threshold, select a different performance threshold and run another optimization report. You can continue creating optimization reports using different performance thresholds until you achieve the expected performance improvements.

Viewing an optimization report

The optimization report displays a storage migration and consolidation report that lists recommendations for improving storage subsystem performance.

The Storage Optimizer does not actually perform any migrations or make any modifications to subsystem configurations. Its primary purpose is to provide you with 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 view the optimization report:

1. Navigate to one of the following nodes:
   - **IBM Tivoli Storage Productivity Center > Analytics**
   - **IBM Tivoli Storage Productivity Center> Disk Manager**.
2. Select a completed analysis job.
3. Optional: Under the **View Previously Run Optimization Reports** pane, click **Refresh Job Status** to update the job status for all optimization reports.
4. Click  to open an optimization report.
5. To view the results of implementing only some of the recommendations, remove the check mark next to **Select all recommendations** and select a subset of the recommendations from the table of recommendations. Only the selected recommendations will be included in the printed report.
6. To display the performance utilization for a specific component, select the component next to **Heat map based on**.
7. To print the report, click the **Print** button in the optimization report. If a PDF printer driver is installed, you can print the report to a PDF file. Alternatively,

you can print the report to a printer. The printed report includes an explanation of each recommendation, and sample migration scripts for the SAN Volume Controller.

The optimization recommendations include the following information:

Recommendation number
The sequential number associated with the recommendation.

Source subsystem
The name of the source subsystem.

Target subsystem
The name of the target subsystem.

Source volume
The name of the source volume.

Target volume
The name of the target volume.

Source pool
The name of the source pool.

Target pool
The name of the target pool.

Reason
The reason why the storage pool utilization has exceeded the performance threshold.

Port information
The recommendation on which ports need to be configured on the target subsystem for the target volume. This information only appears in the printed report.

Zone information
The recommended zoning changes necessary so that the target volume is visible to the host. This information only appears in the printed report.

SVC pseudoscript
The sample migration script that serves as a guide to the commands that you must enter using the SAN Volume Controller command-line interface. This information only appears in the printed report.

After you implement the recommendations, if you want to verify that you have achieved the expected performance improvements, you must first collect more performance monitoring data before running the Storage Optimizer again.

Deleting analysis jobs

You can delete analysis job definitions or individual analysis jobs if you no longer need them.

To delete an analysis job definition or job run:

1. Navigate to one of the following nodes:
   - **IBM Tivoli Storage Productivity Center > Analytics**
   - **IBM Tivoli Storage Productivity Center> Disk Manager**.
2. Expand the **Storage Optimizer** node.
3. Right-click an analysis job definition or an analysis job.
4. Click **Delete**.

Keep in mind that deleting an analysis job definition also deletes all analysis jobs and optimization reports that are associated with that job definition. If you delete an analysis job, all optimization reports associated with that job are also deleted.

Deleting an optimization report

You can delete optimization reports if you no longer need them.

To delete an optimization report job:

1. Navigate to one of the following nodes:
   - **IBM Tivoli Storage Productivity Center > Analytics**
   - **IBM Tivoli Storage Productivity Center> Disk Manager**.
2. Expand the **Storage Optimizer** node.
3. Select a completed analysis job to display the list of optimization reports that are associated with that analysis job.
4. Under **View Previously Run Optimization Reports**, right-click a completed report and choose **Delete**.

Retrieving an optimization report job definition

You can retrieve the report job definition that was used to generate an optimization report.

To retrieve the optimization report job definition associated with an optimization report:

1. Navigate to one of the following nodes:
   - **IBM Tivoli Storage Productivity Center > Analytics**
   - **IBM Tivoli Storage Productivity Center> Disk Manager**.
2. Expand the **Storage Optimizer** node.
3. Select a completed analysis job to display the list of optimization reports that are associated with that analysis job under **View Previously Run Optimization Reports**.
4. Right-click an optimization report and choose **Retrieve**.

The **Generate Optimization Report** pane displays the inputs that were used to create the selected optimization report.

Understanding zoning recommendations

This topic describes the zoning recommendations provided by Storage Optimizer in the printed optimization report.

Note: Storage Optimizer only provides zoning recommendations in the printed optimization report. The report lists the zoning recommendations for all volumes under each volume.

Where possible, Storage Optimizer provides zoning recommendations for each migration recommendation. The zoning recommendations will be to put the target subsystem's port(s) into the same zones as that of the source subsystem's ports, regardless of which zones the host is also a member of.

For example, suppose you want Storage Optimizer to provide zoning recommendations using the following source and target subsystems.

Source subsystem A includes the following:

- Port 1 (member of zones: host1_zone, host2_zone, host3_zone)

- Port 2 (member of zones: host4_zone, host5_zone, host6_zone)
- Volume X (mapped to host1 using Port 1)
- Volume Y (mapped to host4 using Port 2)

Target subsystem B includes the following:

- Port 1
- Port 2

Suppose that Storage Optimizer recommends migrating Volume X from Source subsystem A to Target subsystem B. The zoning recommendations that accompany this migration recommendation will be to add Target subsystem B's Port 1 to zones host1_zone, host2_zone, and host3_zone.

Enforcing storage policies

Enforcing policies for storage usage can be critical to ensuring the proper and most cost-effective use of storage devices. Policies are enforced when conditions that you defined are detected by a monitoring job. Use quotas, constraints, filesystem extensions, and scheduled actions to manage and enforce your storage policies.

Use the **Policy Management** node in Data Manager and Data Manager for Databases to manage and enforce storage policies.

For example, if you want to enforce a policy that limits the amount of space a group of users can consume on the network, use a Data Manager quota. That quota will be triggered: 1) after a scan is run and 2) the scan detected that a group of users is consuming more storage space than was indicated in the quota.

The following table describes the jobs you can use to enforce policies for storage usage in your environment:

Table 17. Descriptions of policy types and their triggering conditions

Policy Type / Manager	What can I do with this job?	What conditions can trigger an alert?
Quotas (Data Manager, Data Manager for Databases) • **Policy Management → Quotas → User** • **Policy Management → Quotas → OS User Group**	Use quotas to define limits on and monitor the amount of storage that a user or a group of users can consume. You can set quotas to specify limits for storage usage at different levels, which enables you to control precisely how much storage a user can consume on specific storage entities: • Define quotas in Data Manager to set limits on the amount of storage that a user or a group of users can consume at a network, filesystem, or computer level. • Define quotas in Data Manager for Databases to set limits on the amount of storage that a user or a group of users can consume at a network, instance, database, or tablespace level. Monitoring job that detects conditions in quotas: scan	**User consumes more than *n***, where *n* represents a number measured in kilobytes, megabytes, or gigabytes. **OS User Group consumes more than *n***, where *n* represents a number measured in kilobytes, megabytes, or gigabytes.

Table 17. Descriptions of policy types and their triggering conditions (continued)

Policy Type / Manager	What can I do with this job?	What conditions can trigger an alert?
Network Appliance Quotas (Data Manager) • **Policy Management → Network Appliance Quotas → Schedules** • **Policy Management → Network Appliance Quotas → Imported User Quotas** • **Policy Management → Network Appliance Quotas → Imported OS User Group Quotas**	Use Network Appliance Quotas to import, view, and report on the NetApp Quotas defined on NAS filers. Monitoring job that detects conditions in Network Appliance quotas: scan	**% of hard limit**, where *n* represents a percentage of the hard limit defined in the NetApp Quota.
Constraints (Data Manager) **Policy Management → Constraints**	Use constraints to: • define the acceptable and unacceptable file types, file sizes, and file owners for computers in your environment. For example, you can use constraints to alert you when users store certain files (such as MP3 or .avi files) on a monitored computer. • Request an IBM Tivoli Storage Productivity Center archive and backup of the largest violating files identified by a constraint. IBM Tivoli Storage Manager protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage. Monitoring job that detects conditions in constraints: scan	**Violating files consume more than n** , where *n* represents a number measured in kilobytes, megabytes, or gigabytes.
Filesystem Extension (Data Manager) **Policy Management → Filesystem Extension**	Use file system extension to create additional space in the local file systems of managed hosts. You can extend file systems manually, or set up policy to do it automatically. You can extend file systems at a specified time, or when utilization reaches a specified threshold. Monitoring job that detects conditions in file system extensions: scan	**A file system extension action started automatically**
Scheduled Actions (Data Manager, Data Manager for Databases) **Policy Management → Scheduled Actions → Scripts**	Use scheduled actions to schedule scripts to run against computers and computer groups. Scheduled actions are script-based, enabling you to use any third-party tools for actions, such as recovery or provisioning. Monitoring job that detects conditions in scheduled actions: n/a	Not applicable. The Scheduled Action facility enables you to run scripts according to the schedule that *you define*. You do not have to associate the script with an alert.

Table 17. Descriptions of policy types and their triggering conditions (continued)

Policy Type / Manager	What can I do with this job?	What conditions can trigger an alert?
Archive/Backup (Data Manager) **Policy Management → Archive/Backup**	View and edit archive and backup jobs that were created based on files selected from reports in the Reporting facility. Monitoring job that detects conditions in archive/backups: scan	Not applicable. Any archive/backup jobs displayed under this node were created based on files shown in the following reports: • Usage > Files > Largest Files • Usage > Files > Most Obsolete Files • Usage > Files > Orphan Files • Usage > Files > Duplicate Files • Backup > Most At Risk Files • Usage Violations > Constraint Violations

Setting storage usage quotas

Use quotas to set limits on the amount of storage space that a user or group of users can consume. You can use Data Manager and Data Manager for Databases to set quotas at the filesystem, computer, network, database, tablespace, and instance levels.

Use quotas in the to set limits on the amount of storage that a user or a group of users can consume.

The following steps provide an overview of how quotas work:

1. A scan gathers extensive information about the storage resources within an enterprise. This information is stored in the database repository.
2. When a quota job runs, it checks the information in the repository to determine if any quota violations exist.
3. If there are any violations, an alert is triggered and the appropriate action is taken.

Note: Because the information analyzed by quotas is provided by scans, you should schedule a quota to run after the scan to collect that information. This ensures that quotas are using the latest information about an enterprise to determine violations.

You can define two different kinds of quotas: User and OS User Group.

- **User Quota.** Select users and groups of users for a quota that have been defined within Data Manager or Data Manager for Databases.
- **OS User Group Quota.** Select OS User Groups for a quota. OS User Groups are groups of users that have been defined in the operating system outside of Data Manager. You can create Groups that contain many OS User Groups. When creating a quota of this type, you can select individual OS User Groups and groups that consist of multiple OS User Groups.

By providing you with the ability to define quotas at different levels, you can control precisely how much storage a user or user group can consume on specific filesystems, computers, network, databases, tablespaces, and instances:

Level	Description
Filesystem	Set storage usage limits on specific filesystems and groups of filesystems that are monitored by Data Manager. This means that you will be alerted if a user or group of users has exceeded the storage usage limit you defined for the specific filesystem or groups of filesystems in the quota.
Computer	Set storage usage limits on specific computers and groups of computers that are monitored by Data Manager. This means that you will be alerted if a user or group of users has exceeded the storage usage limit you defined for the specific computers or groups of computers in the quota.
Database-Tablespace	Set storage usage limits on specific databases, tablespaces, groups of databases, and groups of tablespaces that are monitored by Data Manager for Databases. This means that you will be alerted if a user or group of users has exceeded the storage usage limit you defined for those storage resources in the quota.
Instance	Set storage usage limits on specific instances, computers, and groups of computers that are monitored by Data Manager for Databases. This means that you will be alerted if a user or group of users has exceeded the storage usage limit you defined for the specific instances, computers, and groups of computers in the quota.
Network	Set network-level storage usage limits for *multiple* computers, groups of computers, filesystems, databases, and tablespaces within your network that are monitored by Data Manager and Data Manager for Databases.

Creating a network-wide quota

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of storage space across the filesystems and computers on a network.

To create a network-wide quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** → **Policy Management** → **Quotas**.
2. Expand the **User** or **OS User Group** node depending upon which entity you want to apply a quota.
3. Select the level at which you want to apply the quota by right-clicking **Network**.
4. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
5. In the **Description** field, type a brief description of the quota.
6. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP_BY_OWNER type profiles shown. Make sure that the selected profile is used in scans of all the filesystems chosen on the **Filesystems** tab. This ensures that the quota is fairly enforced for all users.
7. Select users or user groups against which you want to apply the quota:
   a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
   b. Click 

. The users or groups are moved to the **Current Selections** field.

8. Click the **Filesystems** tab; the Filesystems page opens.
9. Select filesystems, filesystem groups, computers, and computers groups on which you want to limit the storage usage of users and user groups.
   a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager** → **Monitoring** → **Groups**.
   b. Click  . The resources are moved to the **Current Selections** field.
10. Click the **When to Run** tab; the When to Run page opens.
11. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

    Run Repeatedly
    Check the quota repeatedly according to a schedule you specify. You can:
    - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
    - Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
    - Indicate which days on which run the job. A check mark will appear next to the days you select.
12. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
    a. 9:00 AM of the time zone where the Data server resides
    b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

    Use the timezone that the server runs in
    Select this option to use the time zone of the location where the Data server resides.

    Use this timezone
    Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.
13. Click the **Alert** tab; the Alert page opens.
14. In the **Triggering Condition** area, specify the event that will trigger the alert:

a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where ***n*** represents a value that you define in the **Value** field. No other conditions are available.

b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota will be triggered.

c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.

15. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

SNMP trap
An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

TEC event
A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

Login notification
The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who will receive alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who will receive alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that will be recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog will be sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

16. Save the quota:

a. Click **File** → **Save**. The Save As window opens.

b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the Navigation Tree pane

Creating a network-wide quota for databases

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of RDBMS storage space across the databases, tablespaces, groups of database-tablespaces, computers, and groups of computers on a network.

To create a network-wide quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager for Databases** → **Policy Management** → **Quotas**.
2. Select the level at which you want to apply the quota by right-clicking **Network**.
3. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
4. In the **Description** field, type a brief description of the quota.
5. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP_BY_OWNER type profiles shown. Make sure that the selected profile is used in scans of all the databases and tablespaces chosen on the **Databases-Tablespaces** tab. This ensures that the quota is fairly enforced for all users.
6. Select users or user groups against which you want to apply the quota:
   a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
   b. Click >>. The users or groups are moved to the **Current Selections** field.
7. Click the **Databases-Tablespaces** tab; the Databases-Tablespaces page opens.
8. Select the databases, tablespaces, database-tablespaces groups, computers, and computer groups on which you want to limit the storage usage of users and user groups.
   a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager for Databases** → **Monitoring** → **Groups**.
   b. Click >>. The resources are moved to the **Current Selections** field.
9. Click the **When to Run** tab; the When to Run page opens.
10. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

Run Repeatedly
Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark will appear next to the days you select.

11. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
    a. 9:00 AM of the time zone where the Data server resides
    b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

    Use the timezone that the server runs in
    Select this option to use the time zone of the location where the Data server resides.

    Use this timezone
    Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.
12. Click the **Alert** tab; the Alert page opens.
13. In the **Triggering Condition** area, specify the event that will trigger the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where ***n*** represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota will be triggered.
    c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
14. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

    TEC event
    A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

    Login notification
    The alert is sent to a Tivoli Storage Productivity Center user. The user

receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who will receive alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who will receive alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that will be recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog will be sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

15. Save the quota:
    a. Click **File** → **Save**. The Save As window opens.
    b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the Navigation Tree pane

Creating a computer quota

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of storage space across computers and groups of computers.

To create a computer quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** → **Policy Management** → **Quotas**.
2. In the **Description** field, type a brief description of the job.
3. Expand the **User** or **OS User Group** node depending upon which entity you want to apply a quota.
4. Select the level at which you want to apply the quota by right-clicking **Computer**.
5. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
6. In the **Description** field, type a brief description of the quota.
7. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP_BY_OWNER type profiles shown.
8. Select users or user groups against which you want to apply the quota:
    a. In the **Available** field, click the users or groups of users against which you want to apply the quota.

b. Click  . The users or groups are moved to the **Current Selections** field.

9. Click the **Computers** tab; the Computers page opens.
10. Select computers and computers groups on which you want to limit the storage usage of users and user groups.
    a. In the **Available** field, click the computers and computers groups to which you want to apply the quota. The groups displayed in this field were created using **Data Manager** → **Monitoring** → **Groups**.
    b. Click  . The resources are moved to the **Current Selections** field.
11. Click the **When to Run** tab; the When to Run page opens.
12. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

    Run Repeatedly
    Check the quota repeatedly according to a schedule you specify. You can:
    - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
    - Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
    - Indicate which days on which run the job. A check mark will appear next to the days you select.
13. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
    a. 9:00 AM of the time zone where the Data server resides
    b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

    Use the timezone that the server runs in
    Select this option to use the time zone of the location where the Data server resides.

    Use this timezone
    Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.
14. Click the **Alert** tab; the Alert page opens.

15. In the **Triggering Condition** area, specify the event that will trigger the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where ***n*** represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota will be triggered.
    c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
16. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

    TEC event
    A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

    Login notification
    The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

    Login ID
    The ID of the user who will receive alert notification upon logging in to the system.

    Windows Event Log, UNIX Syslog
    The ID of the user who will receive alert notification upon logging in to the system.

    Event Type (Windows Event Log only)
    The type of event that will be recorded to the OS log.

    Facility (UNIX Syslog only)
    The location where the UNIX Syslog will be sent. You can select User or Local.

    Level (UNIX Syslog only)
    The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

    Run script
    A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

    E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.
17. Save the quota:

a. Click **File** → **Save**. The Save As window opens.
b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the Navigation Tree pane

Creating a filesystem quota

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of storage space across filesystems, groups of filesystems, computers, and groups of computers.

To create a filesystem quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** → **Policy Management** → **Quotas**.
2. Expand the **User** or **OS User Group** node depending upon which entity you want to apply a quota.
3. Select the level at which you want to apply the quota by right-clicking **Filesystem**.
4. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
5. In the **Description** field, type a brief description of the job.
6. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP_BY_OWNER type profiles shown. Make sure that the selected profile is used in scans of all the filesystems chosen on the **Filesystems** tab. This ensures that the quota is fairly enforced for all users.
7. Select users or user groups against which you want to apply the quota:
   a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
   b. Click 

. The users or groups are moved to the **Current Selections** field.
8. Click the **Filesystems** tab; the Filesystems page opens.
9. Select filesystems, filesystems, computers, and computer groups on which you want to limit the storage usage of users and user groups.
   a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager** → **Monitoring** → **Groups**.
   b. Click 

. The resources are moved to the **Current Selections** field.
10. Click the **When to Run** tab; the When to Run page opens.
11. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

Run Repeatedly
Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark will appear next to the days you select.

12. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
    a. 9:00 AM of the time zone where the Data server resides
    b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

    Use the timezone that the server runs in
    Select this option to use the time zone of the location where the Data server resides.

    Use this timezone
    Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.
13. Click the **Alert** tab; the Alert page opens.
14. In the **Triggering Condition** area, specify the event that will trigger the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where ***n*** represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota will be triggered.
    c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
15. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

    TEC event
    A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

    Login notification
    The alert is sent to a Tivoli Storage Productivity Center user. The user

receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who will receive alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who will receive alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that will be recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog will be sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

16. Save the quota:
    a. Click **File** → **Save**. The Save As window opens.
    b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the Navigation Tree pane

Creating an instance quota

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of RDBMS storage space across instances, computers, and groups of computers.

To create an instance quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager for Databases** → **Policy Management** → **Quotas**.
2. Select the level at which you want to apply the quota by right-clicking **Instance**.
3. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
4. In the **Description** field, type a brief description of the quota.
5. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP_BY_OWNER type profiles shown.
6. Select users or user groups against which you want to apply the quota:
    a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
    b. Click 

. The users or groups are moved to the **Current Selections** field.

7. Click the **Instances** tab; the Instances page opens.
8. Select the instances, computers, and computer groups on which you want to limit the storage usage of users and user groups.
   a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager for Databases** → **Monitoring** → **Groups**.
   b. Click  . The resources are moved to the **Current Selections** field.
9. Click the **When to Run** tab; the When to Run page opens.
10. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

    Run Repeatedly
    Check the quota repeatedly according to a schedule you specify. You can:
    - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
    - Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
    - Indicate which days on which run the job. A check mark will appear next to the days you select.
11. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
    a. 9:00 AM of the time zone where the Data server resides
    b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

    Use the timezone that the server runs in
    Select this option to use the time zone of the location where the Data server resides.

    Use this timezone
    Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.
12. Click the **Alert** tab; the Alert page opens.
13. In the **Triggering Condition** area, specify the event that will trigger the alert:

a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where ***n*** represents a value that you define in the **Value** field. No other conditions are available.

b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota will be triggered.

c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.

14. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

SNMP trap
An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

TEC event
A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

Login notification
The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who will receive alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who will receive alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that will be recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog will be sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

15. Save the quota:

a. Click **File** → **Save**. The Save As window opens.

b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the Navigation Tree pane

Creating a database-tablespace quota

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of RDBMS storage space across databases, tablespaces, groups of databases, and groups of tablespaces.

To create a databases-tablespaces quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager for Databases** → **Policy Management** → **Quotas**.
2. Select the level at which you want to apply the quota by right-clicking **Databases-Tablespaces**.
3. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
4. In the **Description** field, type a brief description of the quota.
5. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP_BY_OWNER type profiles shown. Make sure that the selected profile is used in scans of all the databases and tablespaces chosen on the **Databases-Tablespaces** tab. This ensures that the quota is fairly enforced for all users.
6. Select users or user groups against which you want to apply the quota:
   a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
   b. Click >>. The users or groups are moved to the **Current Selections** field.
7. Click the **Databases-Tablespaces** tab; the Databases-Tablespaces page opens.
8. Select the databases, tablespaces, database-tablespaces groups, computers, and computer groups on which you want to limit the storage usage of users and user groups.
   a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager for Databases** → **Monitoring** → **Groups**.
   b. Click >>. The resources are moved to the **Current Selections** field.
9. Click the **When to Run** tab; the When to Run page opens.
10. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

Run Repeatedly
Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark will appear next to the days you select.

11. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
    a. 9:00 AM of the time zone where the Data server resides
    b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

    Use the timezone that the server runs in
    Select this option to use the time zone of the location where the Data server resides.

    Use this timezone
    Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.
12. Click the **Alert** tab; the Alert page opens.
13. In the **Triggering Condition** area, specify the event that will trigger the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where ***n*** represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota will be triggered.
    c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
14. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

    TEC event
    A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

    Login notification
    The alert is sent to a Tivoli Storage Productivity Center user. The user

receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who will receive alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who will receive alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that will be recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog will be sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

15. Save the quota:
    a. Click **File** → **Save**. The Save As window opens.
    b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the Navigation Tree pane

Editing a storage usage quota

Learn how to change the settings for a quota, including the users or user groups to which the quota applies, the resources on which the quota is set, and the amount set as the limit for storage usage.

To edit a quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** → **Policy Management** → **Quotas**.
2. Expand the **User** or **OS User Group** node depending upon the type of quota you want to edit.
3. Expand the level of the quota you want to edit. For example, if you want to edit a computer quota, expand the **Computer** node.
4. Select the quota you want to edit. Information about the quota you want to edit appears in the right content pane.
5. In the **Description** field, type a brief description of the quota.
6. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP_BY_OWNER type profiles shown.
7. Select additional users or user groups against which you want to apply the quota:

a. In the **Available** field, click the users or groups of users against which you want to apply the quota. The groups displayed in this field were created using **Data Manager** → **Monitoring** → **Groups**

b. Click >>. The users or groups are moved to the **Current Selections** field.

8. Remove users or user groups previously included in the quota:
   a. In the **Current Selections** field, click the users or groups of users that you want to remove from the quota.
   b. Click <<. The users or groups are moved back to the **Available** field.
9. Click the **Computers** or **Filesystems** tab depending on quota you are editing; the Computers or Filesystems page opens.
10. Select additional filesystems, filesystem groups, computers, or computers groups on which you want to limit the storage usage of users and user groups.
    a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager** → **Monitoring** → **Groups**.
    b. Click >>. The resources are moved to the **Current Selections** field.
11. Remove any filesystems, filesystem groups, computers, or computers groups that you do not want included in the quota.
    a. In the **Current Selections** field, click the storage resources to remove.
    b. Click <<. The resources are moved to the **Current Selections** field.
12. Click the **When to Run** tab; the When to Run page opens.
13. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

Run Now
Check the quota immediately.

Run Once at
Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

Run Repeatedly
Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark will appear next to the days you select.

14. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
    a. 9:00 AM of the time zone where the Data server resides
    b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

    Use the timezone that the server runs in
    Select this option to use the time zone of the location where the Data server resides.

    Use this timezone
    Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.
15. Click the **Alert** tab; the Alert page opens.
16. In the **Triggering Condition** area, specify the event that will trigger the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where ***n*** represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota will be triggered.
    c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
17. Click **File** → **Save** to save your changes.

Editing a storage usage quota for databases

Learn how to change the settings for a quota, including the users or user groups to which the quota applies, the resources on which the quota is set, and the amount set as the limit for storage usage.

To edit a quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager for Databases** → **Policy Management** → **Quotas**.
2. Expand the level of the quota you want to edit. For example, if you want to edit a instance quota, expand the **Instance** node.
3. Click the quota you want to edit. Information about the quota you want to edit appears in the right content pane.
4. In the **Description** field, type a brief description of the quota.
5. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP_BY_OWNER type profiles shown.
6. Select additional users or user groups against which you want to apply the quota:
    a. In the **Available** field, click the users or groups of users against which you want to apply the quota. The groups displayed in this field were created using **Data Manager** → **Monitoring** → **Groups**
    b. Click 

. The users or groups are moved to the **Current Selections** field.

7. Remove users or user groups previously included in the quota:
   a. In the **Current Selections** field, click the users or groups of users that you want to remove from the quota.
   b. Click << . The users or groups are moved back to the **Available** field.
8. Click the **Databases-Tablespaces** or **Instances** tab depending on quota you are editing; the Databases-Tablespaces or Instances page opens.
9. Select additional databases, tablespaces, instances, computers, or computers groups on which you want to limit the storage usage of users and user groups.
   a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager for Databases** → **Monitoring** → **Groups**.
   b. Click >> . The resources are moved to the **Current Selections** field.
10. Remove any databases, tablespaces, instances, computers, or computers groups that you do not want included in the quota.
    a. In the **Current Selections** field, click the storage resources to remove.
    b. Click << . The resources are moved to the **Current Selections** field.
11. Click the **When to Run** tab; the When to Run page opens.
12. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

    Run Now
    Check the quota immediately.

    Run Once at
    Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

    Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

    Run Repeatedly
    Check the quota repeatedly according to a schedule you specify. You can:
    - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
    - Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
    - Indicate which days on which run the job. A check mark will appear next to the days you select.
13. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
    a. 9:00 AM of the time zone where the Data server resides

b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

Use the timezone that the server runs in
Select this option to use the time zone of the location where the Data server resides.

Use this timezone
Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.

14. Click the **Alert** tab; the Alert page opens.
15. In the **Triggering Condition** area, specify the event that will trigger the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where ***n*** represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota will be triggered.
    c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
16. Click **File** → **Save** to save your changes.

Storage usage for Network Appliance quotas

Use Network Appliance quotas to import quotas from Network Appliance file servrs and determine how and when you will be alerted to the hard limits defined in those quotas.

Restriction: You can only import quotas if the NetApp device is configured as a file server. It is not available if the NetApp device was added only as a CIMOM.

You can use the NAS support within IBM Tivoli Storage Productivity Center to work with quotas defined using the software provided with Network Appliance (NetApp) File Servers. Using the **Network Appliance** node under **Policy Management**, you can do the following:

- Use the **Schedules** node to define a Data Manager quota that will import NetApp quotas and alert you to when those quotas are close to being exceeded.
- Use the **Imported User Quotas** node to view the definitions of User quotas stored on the NetApp file servers against which you ran a Data Manager quota (defined in the **Schedules** node).
- Use the **Imported OS User Group Quotas** node to view the definitions of OS User Group quotas stored on the NetApp file servers against which you ran a Data Manager quota (defined in the Schedules node).

For example, you can define to be alerted when:

- A user or user group is close to reaching the hard limit defined within the corresponding NetApp Quota
- The space consumed on a QTree is approaching the hard limit defined within the corresponding NetApp Quota

Specifically, the Data Manager NAS component enables you to work with the following types of Network Appliance quotas:

- Quotas that limit the space that QTree directories can consume on a volume
- Quotas that limit the space consumed by specific users on a volume
- Quotas that limit the space consumed by specific user groups on a volume

Quick Steps for Working with Network Appliance quotas

The following describes the general steps you must perform to work with Network Appliance Quotas:

1. Run a Data Manager scan against the Network Appliance file servers whose quotas you want to check. This ensures that Data Manager is using the latest statistics when determining quota violations.
2. Expand **Policy Management → > Network Appliance Quotas**.
3. Right click the **Schedules** node and select Create NetApp Quota Job to define a Quota job. The NetApp Quota window appears.
4. Use the tabs in the NetApp Quota window to specify:
   - the Network Appliance file servers whose quotas you want to import
   - when and how often to run the quota
   - the condition (% of hard limit) that will trigger an alert for the Network Appliance quotas and how you will be notified of that alert. For example, if you enter 90% for the alert condition, Data Manager will notify you whenever any of the quotas you are importing are within 10% of being violated.
5. Run the quota job from the **Schedules** node to import the Network Appliance quota definitions from the target NAS filer. When quotas are imported, the following occurs:
   - any User or OS User Group quotas violations are shown in the Quota Violation reports
   - any directory quota (Qtree) violations and definitions are displayed in **Alert Log → > Directory**.

Creating a Network Appliance storage usage quota

Learn how to create a Network Appliance storage usage quota that imports Network Appliance quotas from filers and alert you to violations of those quotas.

Complete the following steps to:

- define the NetApp filers whose quotas you want to import
- determine when and how often to run the quota that checks for user-defined violations (% of hard limit) of NetApp Quotas
- specify the condition (% of hard limit) that will trigger an alert violation for a quota and how you will be notified of that alert

1. In the Navigation tree pane, expand **Data Manager → Policy Management → Network Appliance Quotas**.
2. Right click the **Schedules** node.
3. Select **New NetApp Quota Job** from the pop-up menu. The content pane displays the What to PROBE page.
4. In the **Description** field, type a brief description of the quota.
5. Select the NetApp filers against which you want to run the quota job:
   a. In the **Available** field, click the NetApp filers or NetApp filer groups against which you want to apply the quota.

b. Click >>. The NetApp filers or NetApp filer groups are moved to the **Current Selections** field.

6. Click the **When to Run** tab; the When to Run page opens.
7. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

   Run Now
   Check the quota immediately.

   Run Once at
   Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

   Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

   Run Repeatedly
   Check the quota repeatedly according to a schedule you specify. You can:
   - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
   - Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
   - Indicate which days on which run the job. A check mark will appear next to the days you select.

8. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
   a. 9:00 AM of the time zone where the Data server resides
   b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

   Use the timezone that the server runs in
   Select this option to use the time zone of the location where the Data server resides.

   Use this timezone
   Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.

9. Click the **Alert** tab; the Alert page opens.
10. In the **Triggering Condition** area, specify the event that will trigger the alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where ***n*** represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, enter the value that represents the percentage of a NetApp's hard limit at which the alert will be triggered. The hard limit represents the most amount of space a user or directory (Qtree) can consume on a NAS filer. For example, if the hard limit defined within a

NetApp Quota for a user group is 150MB, and you enter **90** in the **Value** field, the Data Manager quota will report a violation when that user group consumes 135 or more megabytes.

c. In the **Value Units** field, the only value units available is the % of Hard Limit.

11. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

SNMP trap
An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

TEC event
A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

Login notification
The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who will receive alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who will receive alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that will be recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog will be sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

12. Save the quota:
    a. Click **File** → **Save**. The Save As window opens.
    b. In the **Specify NetApp Quota job name** field, type a name for the quota and click **OK**. The quota is displayed in the Navigation Tree pane

Viewing user quotas imported from Network Appliance filers

Learn how to view the definitions of user quotas stored on the NetApp filers against which you ran a Data Manager quota.

Before you can import Network Appliance user quotas and view them through **Data Manager → Policy Management → Network Appliance Quotas → Imported User Quotas**, you must:

1. Run a scan against the Network Appliance filers whose user quotas you want to import.
2. Use **Policy Management → Network Appliance Quotas → Schedules** to define and run a quota checking job against the NetApp filers whose user quotas you want to import.

Complete the following steps to view imported Network Appliance user quotas:

1. In the Navigation tree pane, expand **Data Manager → Policy Management → Network Appliance Quotas → Imported User Quotas**.
2. Click the NetApp Quota whose definition you want to view. The content pane displays information about the NetApp user quota.
3. In the Users page of the Edit *quota_name* window, view the following information about imported Network Appliance user quotas:

   Hard Limit
   The hard limit for used space defined within the NetApp quota.

   Current Selections
   The users against which the NetApp quota applies.
4. In the Filesystems page of the Edit *quota_name* window, view the filesystems to which the NetApp Quota applies.

Viewing OS user quotas imported from Network Appliance filers

Learn how to view the definitions of OS User Quotas stored on the NetApp filers against which you ran a Data Manager quota.

Before you can import Network Appliance OS user quotas and view them through **Data Manager → Policy Management → Network Appliance Quotas → Imported OS User Group Quotas**, you must:

1. Run a scan against the Network Appliance filers whose OS user quotas you want to import.
2. Use **Policy Management → Network Appliance Quotas → Schedules** to define and run a quota checking job against the NetApp filers whose OS user quotas you want to import.

Complete the following steps to view imported Network Appliance OS user quotas:

1. In the Navigation tree pane, expand **Data Manager → Policy Management → Network Appliance Quotas → Imported OS User Quotas**.
2. Click the NetApp Quota whose definition you want to view. The content pane displays information about the NetApp OS user quota.
3. In the OS Users page of the Edit *quota_name* window, view the following information about imported Network Appliance OS user quotas:

   Hard Limit
   The hard limit for used space defined within the NetApp quota.

Current Selections

The OS users against which the NetApp quota applies.

4. In the Filesystems page of the Edit *quota_name* window, view the filesystems to which the NetApp Quota applies.

Viewing directory (QTree) quotas and violations

Learn how to view the definitions of user quotas stored on the Network Appliance filers against which you ran a Data Manager quota.

Before you can view QTree quotas and violations, perform the following steps:

1. Use **Policy Management** → **Network Appliance Quotas** → **Schedules** to define and run a quota checking job against the Network Appliance filers containing the QTree quotas you want to import.
2. Create and run a scan job that includes the following objects:
   - the NetApp filers containing the QTrees whose quotas you want to import
   - the directory group named TPCUser.NetApp QTree Group
3. Use **Policy Management** → **Network Appliance Quotas** → **Schedules** to rerun the quota checking job you defined in step 1.

Complete the following steps to view a QTree quota violation and definition:

1. In the Navigation tree pane, expand **Alerting** → **Alert Log**.
2. Click the **Directory** node. The Alert History — All Directories window appears. This window lists all the alerts triggered by directory and QTree quota violations.
3. To view details about a violation you can do either of the following:
   - Double click the entry representing the QTree violation to view details about that violation and the QTree quota definition.
   - Right click the entry representing the QTree violation and select **Edit Definition** from the pop-up menu to view the QTree quota definition.

     Note: You cannot edit a QTree quota definition using Data Manager. To edit a QTree quota definition you must use the software provided with your Network Appliance filer.

Setting file constraints

Use constraints to define the acceptable and unacceptable file types, file sizes, and file owners for a NAS filer or a set of NAS filers in your environment. You can also use constraints to request an IBM Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint.

Creating a file constraint

Learn how to create a constraint that defines the acceptable and unacceptable file types, file sizes, and file owners for a computer or a set of computers in your environment.

To create a constraint, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** → **Policy Management**.
2. Right-click **Constraints**, and click **Create Constraint**. The Create Constraint window opens.
3. In the **Description** field, type a brief description of the job.
4. Select filesystems, filesystem groups, computers, and computer groups against which to apply the constraint:

a. In the **Available** field, click the storage resources against which you want to apply the constraint. Select individual filesystems and computers, as well as file system and computer groups that you have previously defined using the group facility. Note that cluster resource groups in an HACMP or MSCS clustered environment are listed as computers.

b. Click  . The storage resources are moved to the **Current Selections** field.

5. Click the **File Types** tab; the File Types page opens. Use this page to select the file types that you want to allow or forbid.
6. Click one of the following:

   Forbid file matching these patterns
   Do *not* allow files of the specified type on the storage resources selected on the Filesystems page.

   Allow ONLY file matching these patterns
   Allow files *only* of the specified type on the storage resources selected on the Filesystems page.
7. Click the file types that you would like to allow or forbid in the **Choose a file type** list box. The file types listed in this box are predefined in the system.
   a. In the **Choose a file type:** field, click the file types that you would like to allow or forbid.
   b. Click  . The file types are moved to the **Forbidden files** field.
8. Optionally, create a custom file type:
   a. Enter the pattern in the **Or enter a pattern** field.
   b. Click  . The file type is moved to the **Forbidden files** field.
9. Click the **Users** tab; the Users page opens. Use this page to allow or forbid files owned by specific users
10. Click one of the following:

    Forbid file owned by selected users
    Do *not* allow files owned by the specified users.

    Allow ONLY file owned by the specified owner
    Allow files owned by the specified users only.
11. Select the users whose files you want to allow or forbid:
    a. In the **Available Users** field, click the users you want to include.
    b. Click  . The users are moved to the **Forbidden Users** field.
12. Click the **Options** tab; the Options page opens. Use this page to define alerts for files based on characteristics other than file type and file owner, and to view a where clause that will be used as a filter for locating files that violate the constraint definition. This page contains the following information:

    Violating File Limits
    Indicate the maximum number of violating filenames that should be kept per agent.

    number of days since last access greater than
    Click this box and enter a length of time (in days) in which files have not been accessed. When a constraint job locates a file that has not been accessed in the length of time you indicated in this field, an alert is triggered.

bigger than ##
Check this box and enter a file size that should trigger an alert. When a constraint job locates a file of the size (or greater) that you indicated in this field, an alert is triggered. Select the unit of measurement for the value you entered in the list box to the right. The valid choices are: bytes, **KB** (kilobytes)**MB** (megabytes), **GB** (gigabytes), and **TB** (terabytes).

set-UID root
Check this box trigger alerts on files that have the SET-UID root bit on (because this can be a security consideration).

not backed up (Windows only) and number of days since last modification greater than
Click this box and enter the number of days since a file was last modified and not backed up. When the constraint job locates a file that was modified and not backed up in the amount of days that you indicate in this field, an alert is triggered.

File Filter text
This section displays a *where* clause that will be used as a filter for locating files that violate the Constraint definition. The choices you make on this tab and the choices you made on the Users and File Types pages are reflected in the content of this *where* clause. To edit the *where* clause filter text, click **Edit Filter**. This will allow you to create new where conditions and groupings using the where clause interface.

13. Click the **Alert** tab; the Alert page opens. Use the Alert tab to define the alerting criteria for a constraint.
14. In the **Triggering Condition** area, specify the event that will trigger the constraint alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than** ***n***, where ***n*** represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the amount of space a file must consume to violate the constraint.
    c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.

    Once Data Manager has found enough files that violate the condition (i.e., that exceed the total size amount you specified in this section), the constraint is considered violated and the alert is triggered. When the alert is triggered, the actions that you specify in the **Triggered Actions** section are taken.
15. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

    TEC event
    A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also

performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

Login notification
The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who will receive alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who will receive alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that will be recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog will be sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

16. Save the constraint:
    a. Click **File** → **Save**. The Save As window opens.
    b. In the **Specify Constraint name** field, type a name for the constraint and click **OK**. The constraint is displayed in the Navigation Tree pane. The constraint is automatically checked whenever an appropriate scan job is run.

Editing a constraint

Edit a constraint to: change the acceptable and unacceptable file types, file sizes, and file owners, add or remove users and user groups to which the constraint applies, edit the constraint options, and change the alert settings.

To edit a constraint, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** → **Policy Management** → **Constraints**.
2. Click the constraint you want to edit. Information about the constraint you want to edit appears in the right content pane.
3. Add or remove the filesystems, filesystem groups, computers, and computer groups against which to apply the constraint:
    a. In the **Available** field, click the storage resources that you want add. Click 

. The users or groups are moved to the **Current Selections** field.

b. In the **Current Selections** field, click the storage resources that you want remove from the constraint. Click  . The users or groups are moved to the **Available** field.

4. Click the **File Types** tab; the File Types page opens. Use this page to edit the file types that you want to allow or forbid.
5. Click one of the following:

   Forbid file matching these patterns
   Do *not* allow files of the specified type on the storage resources selected on the Filesystems page.

   Allow ONLY file matching these patterns
   Allow files *only* of the specified type on the storage resources selected on the Filesystems page.
6. Click the file types that you would like to allow or forbid in the **Choose a file type** list box. The file types listed in this box are predefined in the system.
   a. In the **Choose a file type:** field, click the file types that you would like to allow or forbid.
   b. Click  . The file types are moved to the **Forbidden files** field.
7. Optionally, create a custom file type:
   a. Enter the pattern in the**Or enter a pattern** field.
   b. Click  . The file type is moved to the **Forbidden files** field.
8. Click the **Users** tab; the Users page opens. Use this page to allow or forbid files owned by specific users
9. Click one of the following:

   Forbid file owned by selected users
   Do *not* allow files owned by the specified users.

   Allow ONLY file owned by the specified owner
   Allow files owned by the specified users only.
10. Select the users whose files you want to allow or forbid:
    a. In the **Available Users** field, click the users you want to include.
    b. Click  . The users are moved to the **Forbidden Users** field.
11. Click the **Options** tab; the Options page opens. Use this page to edit alerts for files based on characteristics other than file type and file owner, and to view a where clause that will be used as a filter for locating files that violate the constraint definition. This page contains the following information:

    Violating File Limits
    Indicate the maximum number of violating filenames that should be kept per agent.

    number of days since last access greater than
    Click this box and enter a length of time (in days) in which files have not been accessed. When a constraint job locates a file that has not been accessed in the length of time you indicated in this field, an alert is triggered.

    bigger than ##
    Check this box and enter a file size that should trigger an alert. When a constraint job locates a file of the size (or greater) that you indicated in this field, an alert is triggered. Select the unit of measurement for

the value you entered in the list box to the right. The valid choices are: bytes, **KB** (kilobytes)**MB** (megabytes), **GB** (gigabytes), and **TB** (terabytes).

set-UID root
Check this box trigger alerts on files that have the SET-UID root bit on (because this can be a security consideration).

not backed up (Windows only) and number of days since last modification greater than
Click this box and enter the number of days since a file was last modified and not backed up. When the constraint job locates a file that was modified and not backed up in the amount of days that you indicate in this field, an alert is triggered.

File Filter text
This section displays a *where* clause that will be used as a filter for locating files that violate the Constraint definition. The choices you make on this tab and the choices you made on the Users and File Types pages are reflected in the content of this *where* clause. To edit the *where* clause filter text, click **Edit Filter**. This will allow you to create new where conditions and groupings using the where clause interface.

12. Click the **Alert** tab; the Alert page opens. Use the Alert tab to define the alerting criteria for a constraint.
13. In the **Triggering Condition** area, specify the event that will trigger the constraint alert:
    a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where ***n*** represents a value that you define in the **Value** field. No other conditions are available.
    b. In the **Value** field, type a numeric value that represents the amount of space a file must consume to violate the constraint.
    c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.

    Once Data Manager has found enough files that violate the condition (i.e., that exceed the total size amount you specified in this section), the constraint is considered violated and the alert is triggered. When the alert is triggered, the actions that you specify in the **Triggered Actions** section are taken.
14. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

    TEC event
    A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

    Login notification
    The alert is sent to a Tivoli Storage Productivity Center user. The user

receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who will receive alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who will receive alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that will be recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog will be sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

15. Click **File** → **Save** to save your changes.

Requesting an archive or backup

Request anIBM Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint. Tivoli Storage Manager protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage.

Data Manager provides a method within the alerting function of a constraint where you can automatically invoke a Tivoli Storage Manager archive or backup to run against the files that violate that constraint. For example, you can do the following:

- Define a constraint to identify the .vbs files within your organization that are larger than two megabytes. Other file types include .avi, .dll, .doc, .exe, .gif, .jpg, .mp3, .rtf, .txt., wsh, and .zip. You can also define a pattern to use when identifying files that need to be archived or backed up.
- Define a constraint that will invoke Tivoli Storage Manager to archive and then delete those files

Before you begin

Before you can use the Tivoli Storage Manager archive / backup functionality with Data Manager constraints, you must ensure the following:

- Your user ID has Administrator authority within IBM Tivoli Storage Productivity Center.

- Tivoli Storage Manager client software is installed on each of the hosts where a Data Manager scanning agent will be used to perform archive and backup functions.
- (for SAN attached storage only) Tivoli Storage Manager client software must be installed on the Data agent host machine that is configured to act as the scanning agent for NAS storage.
- A Tivoli Storage Manager client must be configured to run without a password. This is typically done by either including the password in the client options (dsm.opt) file, or using the options PASSWORDACCESS = GENERATE in the client options file.
- the Tivoli Storage Manager client must be version 5.1 or higher

Defining an archive or backup in a constraint:

Learn how to request an IBM Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint. Request a Tivoli Storage Manager archive and backup on the alert page of a constraint definition.

Create a new constraint:

1. In the Navigation tree pane, expand **Data Manager** → **Policy Management**.
2. Right-click **Constraints**, and click **Create Constraint**. The Create Constraint window opens.
3. Select filesystems, filesystem groups, computers, and computer groups against which to apply the constraint:
   a. In the **Available** field, click the storage resources against which you want to apply the constraint. Select individual filesystems and computers, as well as file system and computer groups that you have previously defined using the group facility. Note that cluster resource groups in an HACMP or MSCS clustered environment are listed as computers.
   b. Click  . The users or groups are moved to the **Current Selections** field.
4. Click the **File Types** tab; the File Types page opens. Use this page to select the file types that you want to allow or forbid.
5. Click one of the following:

   Forbid file matching these patterns
   Do *not* allow files of the specified type on the storage resources selected on the Filesystems page.

   Allow ONLY file matching these patterns
   Allow files *only* of the specified type on the storage resources selected on the Filesystems page.
6. Click the file types that you would like to allow or forbid in the **Choose a file type** list box. The file types listed in this box are predefined in the system.
   a. In the **Choose a file type:** field, click the file types that you would like to allow or forbid.
   b. Click  . The file types are moved to the **Forbidden files** field.
7. Optionally, create a custom file type:
   a. Enter the pattern in the **Or enter a pattern** field.
   b. Click  . The file type is moved to the **Forbidden files** field.

8. Click the **Users** tab; the Users page opens. Use this page to allow or forbid files owned by specific users
9. Click one of the following:

   Forbid file owned by selected users
   Do *not* allow files owned by the specified users.

   Allow ONLY file owned by the specified owner
   Allow files owned by the specified users only.
10. Select the users whose files you want to allow or forbid:
    a. In the **Available Users** field, click the users you want to include.
    b. Click 

. The users are moved to the **Forbidden Users** field.
11. Click the **Options** tab; the Options page opens. Use this page to define alerts for files based on characteristics other than file type and file owner, and to view a where clause that will be used as a filter for locating files that violate the constraint definition. This page contains the following information:

    Violating File Limits
    Indicate the maximum number of violating filenames that should be kept per agent.

    number of days since last access greater than
    Click this box and enter a length of time (in days) in which files have not been accessed. When a constraint job locates a file that has not been accessed in the length of time you indicated in this field, an alert is triggered.

    bigger than ##
    Check this box and enter a file size that should trigger an alert. When a constraint job locates a file of the size (or greater) that you indicated in this field, an alert is triggered. Select the unit of measurement for the value you entered in the list box to the right. The valid choices are: bytes, **KB** (kilobytes)**MB** (megabytes), **GB** (gigabytes), and **TB** (terabytes).

    set-UID root
    Check this box trigger alerts on files that have the SET-UID root bit on (because this can be a security consideration).

    not backed up (Windows only) and number of days since last modification greater than
    Click this box and enter the number of days since a file was last modified and not backed up. When the constraint job locates a file that was modified and not backed up in the amount of days that you indicate in this field, an alert is triggered.

    File Filter text
    This section displays a *where* clause that will be used as a filter for locating files that violate the Constraint definition. The choices you make on this tab and the choices you made on the Users and File Types pages are reflected in the content of this *where* clause. To edit the *where* clause filter text, click **Edit Filter**. This will allow you to create new where conditions and groupings using the where clause interface.
12. Click the **Alert** tab; the Alert page opens. Use the Alert tab to define the alerting criteria for a constraint.
13. In the **Triggering Condition** area, specify the event that will trigger the constraint alert:

a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where ***n*** represents a value that you define in the **Value** field. No other conditions are available.
b. In the **Value** field, type a numeric value that represents the amount of space a file must consume to violate the constraint.
c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.

Once Data Manager has found enough files that violate the condition (i.e., that exceed the total size amount you specified in this section), the constraint is considered violated and the alert is triggered. When the alert is triggered, the actions that you specify in the **Triggered Actions** section are taken.

14. In the **Triggered Actions** area, select **Archive/Backup** and click **Define**. The Archive / Backup Options window appears. Click one of the following:

Archive
Perform an archive on the files that violate a constraint. Archives are copies of data from online storage copies to offline storage. If you select the Archive option, you can also check the **Delete After Successful Archival** box to automatically delete the files that violated the constraint after the files have been successfully archived.

Selective Backup
Perform a selective backup of the files that violate a constraint. A selective backup enables you to back up specified files. These files are not excluded in the include-exclude list and meet the requirement for serialization in the backup copy group of the management class assigned to each file.

Incremental Backup
Perform an incremental backup of the files that violate a constraint. An incremental backup enables you to back up new or changed files or directories from a client domain or from specified directories or files. These directories or files are not excluded in the include-exclude list and meet the requirements for frequency, mode, and serialization as defined by a backup copy group of the management class assigned to each file.

Tivoli Storage Manager Command Options
Enter additional that will be passed on to the Tivoli Storage Manager dsmc command. You can enter up to 512 bytes of text.

15. Save the constraint:
a. Click **File** → **Save**. The Save As window opens.
b. In the **Specify Constraint name** field, type a name for the constraint and click **OK**. The constraint is displayed in the Navigation Tree pane. The constraint is automatically checked whenever an appropriate scan job is run.

See the *IBM Tivoli Storage Manager Client User's Guide* for more information about each of the TSM actions you can select when defining triggered actions for a constraint.

Using filesystem extension

Use filesystem extensions to create additional space in the local filesystems of managed hosts. You can extend filesystems manually, or set up policies to do it automatically.

Use **Data Manager → Policy Management → Filesystem Extension** to create additional space in the local filesystems of managed hosts either manually or automatically through a policy. Policies can be configured to extend filesystems at a specified time, or when utilization reaches a specified threshold. For managed hosts that have access to an Tivoli Storage Enterprise Storage Server, you can also allocate additional LUNs when there is not enough space to extend file systems in the local volume group.

Filesystem extension is supported for JFS file systems running on AIX 5.1 and VxFS file systems running on Sun Solaris 2.8. The AIX JFS2 file system is not supported.

Note: Filesystem extension/LUN provisioning jobs are not supported under Solaris 9.

How Filesystem Extension works

Filesystem extension is defined as a policy. The policy specifies the following:

- Which filesystems will be eligible for extension
- How the filesystems will be extended
- Whether LUN provisioning will be enabled, and how it will work
- When to enforce the policy
- How alerting will be handled

You can also choose to have no extension performed when the policy is applied. Instead, any policy actions that would have been performed are written to a log file. This feature, called LOG ONLY mode, can be used to preview the results of a policy before extending any file systems.

LUN Provisioning

By default, IBM Tivoli Storage Productivity Center will attempt to extend a filesystem in its local volume group. If there is not enough space in the local volume group and an Tivoli Storage Enterprise Storage Server subsystem is available, additional LUNs can be provisioned for filesystem extension. Provisioning consists of creating and formatting a new LUN according to the parameters specified in the filesystem extension policy, and assigning the new LUN to the managed host. You can enable and configure LUN provisioning using the Provisioning tab in the Filesystem Extension window.

The Data server uses the Tivoli Storage Enterprise Storage Server Common Information Model/Object Manager (CIMOM) to interact with Tivoli Storage Enterprise Storage Server subsystems. The CIMOM is installed as an agent on a host server or workstation in your network. In many cases, CIMOMs can be automatically discovered, using Service Location Protocol (SLP). The ability to discover a CIMOM depends on its location and the operating system on which it is installed.

When a CIMOM is discovered or manually added, you must provide access information for it, including the username, password, and location of the TrustStore certificate file. The certificate file is used to ensure secure communication between the Data Manager server and the CIMOM. The certificate file is created when you install and set up the CIMOM. The certificate file for each CIMOM must be copied to the Data Manager server machine. For information about supported versions of the CIMOM, see the Tivoli Storage Productivity

Center support website at http://www.ibm.com/servers/storage/support/software/tpc/ For more information and detailed instructions for setting up CIMOM access, see the CIMOM Logins section in the *IBM Tivoli Storage Productivity Center for Data Installation and Configuration Guide*.

Note: To use a CIMOM with Data Manager, eFix 1.1.0.2 must be installed on the CIMOM machine. This patch is available from the Tivoli Storage Productivity Center Technical Support Website.

Creating a file system extension

Learn how to create a file system extensions policy that automatically creates additional space in the local file systems of managed hosts.

To create a file system extension policy, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** → **Policy Management**.
2. Right-click **Filesystem Extension**, and click **Create Filesystem Extension Rules**. The Create Filesystem Extension Rules window opens and the Filesystems page is displayed. Use this page to select the file systems to which this policy will apply. Filesystem extension does not currently support the use of predefined file system groups. Only individual file systems can be selected when defining extension policy. However, multiple file systems can be specified for the same policy.

   A file system can only have one policy assigned to it. If a file system is already associated with a policy, the policy name will be displayed next to the file system name. If you select this file system for the extension policy, the previous policy assignment will be overridden when you save the new policy.

   Note: When a new file system is added, it will appear in the Available list box only after a Probe job is run.

   The file systems tree view includes a file system extension preview function. File systems that would be extended under the currently defined policy are displayed with a green arrow icon. If a file system would qualify for extension but has reached the maximum size specified by the policy, it is displayed with a red arrow icon. Click the Refresh button to update the tree view.

   Note: File system extension supports the use of directly attached Serial Storage Architecture (SSA) devices, such as the IBM Tivoli Storage 7133 Serial Disk System. This support is provided for AIX versions 5.1, 5L, and 5.2. For operating systems that connect to SSA devices by emulating SCSI, Data Manager will recognize the SSA device as a SCSI drive.
3. In the **Description** field, type a brief description of the job.
4. Select the file systems for the policy:
   a. In the **Available** field, click the file systems you want to include.
   b. Click 

. The file systems are moved to the **Current Selections** field.
5. Click the **Extension** tab; the Extension page opens. Use this page to specify how file systems will be extended. You can choose to extend file systems:
   - By a specific amount of space
   - By a percentage of file system capacity
   - By increasing free space to a specific size
   - By increasing free space to a percentage of file system capacity
   - By increasing capacity to a specific size

Select the option you want to use for this policy. You can also set a maximum capacity size if necessary. If you specify a maximum size, a file system extension job that would increase a file system beyond that size will fail.

Note: The actual amount of file system extension can vary, depending on your environment. For example, if 5GB physical partitions are set up for a volume group, its file systems will be extended in 5GB increments. If you specify a smaller amount, it will be rounded up to the physical partition size.

As another example, software striping and software RAID will claim part of the space created when file systems are extended. In this case, file system extension will appear to create slightly less space than specified in the policy.

If you are using LUN provisioning, you can specify the size of new LUNs by using the Create LUNs that are at least option on the Provisioning tab and one of the Amount to Extend options on the Extension tab. If the Amount to Extend value is greater than or equal to the preferred minimum LUN size, 1% more space will be added to the LUN to cover overhead requirements.
If you choose to extend file systems regardless of remaining free space, the policy options you specify will be applied to all selected file systems when the policy is enforced. To use policy to automatically manage file system capacity, you must do the following:

- Use the **When to Run** page to set the policy to run after every probe or scan
- Use the **Extension** page to specify an amount of free space to be used as a condition for extension. You can do this two ways:
  - Select **Make Freespace** and specify an amount. Then select Extend filesystems regardless of remaining free space. When free space falls below the amount you specify, the difference in space will be added. This option can be used to maintain file system free space at a specific level.
  - Select either **Add** or **Make Freespace** and specify an amount. Then select **Extend filesystems** when free space is less than and specify an amount. The amount you specify will be used to trigger the **Amount to Extend** action you selected. This option can be used to add extra space to a file system if free space reaches a specified level. For example, if free space reaches 10% or less, it could be increased to 20%, or 5GB could be added to the total file system capacity.

  Note: The following restrictions apply when you create or edit a file system extension policy:
  - If you select both **Make Freespace** and **Extend filesystems when freespace is less than**, you must choose the same type of extension for each. If you specify a percentage amount for one, you must specify a percentage amount for the other. If you specify an amount expressed in megabytes for one, you must specify an amount in either megabytes or gigabytes for the other.
  - If you select both **Make Freespace** and **Extend filesystems when freespace is less than**, the extension amount you specify for **Make Freespace** must be larger than the amount you specify for **Extend filesystems when freespace is less than**.
  - If you select **Make Capacity**, you cannot select **Extend filesystems when freespace is less than**. Because the **Make Capacity** option extends capacity regardless of actual utilization, it is not intended to be used for the automated maintenance of file system capacity. You can use this

option to extend file system capacity when necessary by selecting **Run Now** or **Run Once** at on the **When to Run** page.

You can also choose to have no file systems extended when the policy is applied. Instead, any policy actions that would have been performed are written to the log file. This feature, called LOG ONLY mode, can be used to preview the results of a policy over time before actually extending any file systems.

6. Click the **Provisioning** tab; the Provisioning page opens. Use the Provisioning tab to enable LUN provisioning and specify how it will be performed. Provisioning is an optional feature that is not required for file system extension. If you enable LUN provisioning, two things will happen when there is not enough space in a volume group to extend a file system:
   - A new LUN is created to add space to the volume group, according to the policy specified on the **Provisioning** tab.
   - The file system is extended within the volume group, according to the policy specified on the **Extension** tab.

   For example, a volume group might first be extended by 100Mb, and then one or more file systems within it might be extended by 10MB each.

   Data Manager relies on information stored by a CIM/OM to interact with Tivoli Storage Enterprise Storage Server subsystems. To implement provisioning, CIM/OM access information must be defined to Data Manager. See the CIM/OM Logins section in the Configuration and Getting Started Guide for more information. After CIM/OM access is configured, a Discovery job must be run to identify Tivoli Storage Enterprise Storage Server subsystems.

   Note: LUNs can be provisioned for file system hosts running Sun Solaris, but the hosts must be specially configured to avoid a reboot after provisioning. See the appendix on configuring LUN provisioning for Sun Solaris in the IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide for instructions.

   If you choose to enable automatic LUN provisioning, you must choose a model and location for new LUNs. Data Manager models new LUNs on existing LUNs. Modeled characteristics include CIM/OM access parameters, as well as connectivity and RAID type. You must specify how far Data Manager will look for a model LUNs, as well as limit where it will create new LUNs on the Tivoli Storage Enterprise Storage Server.

   Note: To perform LUN provisioning, at least one Tivoli Storage Enterprise Storage Server LUN must be currently assigned to the volume group of the file system you want to extend, or to the managed host associated with the file system you want to extend. Select one of the following options to limit the scope of LUN modeling:
   - Model new LUNs on others in the volume group of the file system being extended. If you select this option, Data Manager will only search for a model within existing LUNs in the volume group of the file system being extended.
   - Model new LUNs on others on the same host as the file system being extended. If you select this option, Data Manager will first search for a model within existing LUNs in the volume group of the file system being extended, before looking further on the same host.

New LUNs can be created in the same Storage Pool as the file system being extended, or in any Storage Pool in the Tivoli Storage Enterprise Storage Server of the file system being extended. Select one of the following options to limit where new LUNs will be created:

- Provision new LUNs within the same Storage Pool as the file system being extended. For Tivoli Storage Enterprise Storage Server, a Storage Pool is a Volume Space. If you are using RAID, a Volume Space is equal to one RAID rank, which consists of a Disk Group (8 physical disks).
- Provision new LUNs within the same Storage System as the file system being extended. If you select this option, new LUNs can be created in any Storage Pool in the Tivoli Storage Enterprise Storage Server. Data Manager will first attempt to provision a new LUN in the same Storage Pool as the file system being extended, before looking further in the same Storage System.

You can optionally specify a preferred minimum size for new LUNs by selecting the **Create LUNs that are at least** option. If you select this option, LUNs will not be created any smaller than the size you specify. If you do not specify a preferred minimum size, the Amount to Extend specified for the policy will be used. For Tivoli Storage Enterprise Storage Server subsystems, the default minimum LUN size increment is 100 MB. New LUNs will be rounded up to satisfy this requirement. If the Amount to Extend value is greater than or equal to the preferred minimum LUN size you specify, 1% more space will be added to the LUN to cover overhead requirements. See Table 18.

Note: Data Manager will only provision new LUNs. Existing unassigned LUNs will not be reused. Similarly, if provisioning fails after LUNs have been created, the new LUNs will not be deleted or reused. This limitation protects against accidentally overwriting existing data that might be associated with currently unassigned LUNs.
Table 18 shows the possible outcomes when a new LUN is provisioned.

Table 18. LUN Provisioning Results

Description	Amount to Extend	Tivoli Storage Enterprise Storage Server minimum for new LUNs	Create LUNs that are at least (preferred minimum LUN size)	Actual LUN size that will be created	Actual amount by which the file system will be extended
Minimum LUN size exceeds amount to extend	20Mb	100 Mb	200Mb	200Mb	20Mb
Amount to extend equals preferred minimum LUN size	200Mb	100 Mb	200Mb	202Mb (1% added for overhead)	200Mb
Minimum LUN size exceeds amount to extend	100Mb	100 Mb	200Mb	200Mb	100Mb

Table 18. LUN Provisioning Results (continued)

Description	Amount to Extend	Tivoli Storage Enterprise Storage Server minimum for new LUNs	Create LUNs that are at least (preferred minimum LUN size)	Actual LUN size that will be created	Actual amount by which the file system will be extended
Amount to extend exceeds preferred minimum LUN size, but both are below Tivoli Storage Enterprise Storage Server minimum	75Mb	100 Mb	50Mb	100Mb	75Mb
Minimum LUN size not specified	200Mb	100 Mb	Not specified	202Mb (1% added for overhead)	200Mb

Note: IBM Tivoli Storage Enterprise Storage Server and theTivoli Storage Enterprise Storage Server CIM/OM use multiples of 1000 bytes to represent kilobytes, megabytes, and gigabytes. Data Manager uses multiples of 1,024 bytes. Data Manager automatically converts this difference when it discovers Tivoli Storage Enterprise Storage Server capacity and utilization. If you are using IBM Tivoli Storage Tivoli Storage Enterprise Storage Server Expert or Tivoli Storage Enterprise Storage Server Specialist, the storage information they show will differ from what is shown by Data Manager.

Due to AIX Volume Manager limitations on AIX Striped Volumes, provisioning of new LUNs cannot be done. If the policy definition for a striped file system indicates to Enable Automatic LUN Provisioning on the Provisioning tab, then any LUNs that are created will be added to the Volume Group but will not be usable for Extension. If provisioning is attempted on an AIX Striped Volume, then you will see the message: `STA0301W: To prevent runaway condition, we have removed the file system file system_name from the policy, please fix the problem and add the file system back to the policy.`

To avoid this problem, make sure the **Enable Automatic LUN Provisioning** option is disabled for Filesystem Extension policy definitions that have striped file systems on AIX.

7. Click the **When to Run** tab; the When to Run page opens. Use this page to specify when to apply the file system extension policy to selected file systems.
8. In the **How often to run** area, schedule how often you want to enforce the policy. You can choose from the following actions:

   Enforce policy after every Probe or Scan
   Enforce the policy after every probe or scan that is run against the selected filesystems. If you want to automate filesystem extension, you should choose to apply the policy after every new probe or scan.

   Run Now
   Enforce the policy immediately.

   Run Once at
   Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

Note: TheTivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

Note: If the policy applies to multiple filesystems and you select **Run Now** or **Run Once at**, messages specific to each filesystem will be mixed together in the log file. This can make it difficult to troubleshoot errors. To avoid this, you can select one policy at a time and then select **Run Now**.

9. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
   a. 9:00 AM of the time zone where the Data server resides
   b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

   Use the timezone that the server runs in
   Select this option to use the time zone of the location where the Data server resides.

   Use this timezone
   Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.
10. Click the **Alert** tab; the Alert page opens. Use this page to define the alerting criteria for a file system extension.
11. In the **Triggering Condition** area, specify the alerting criteria for the filesystem extension. In the **Condition** field, the default setting is **A filesystem extension action started automatically**. No other conditions are available. When a filesystem is extended by the policy, an alert will be triggered. When the alert is triggered, the actions you specified in the Triggered Actions section are taken.
12. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

    TEC event
    A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

    Login notification
    The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who will receive alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who will receive alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that will be recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog will be sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

13. Save the filesystem extension:
    a. Click **File** → **Save**. The Save As window opens.
    b. In the **Specify Filesystem Extension Rule name** field, type a name for the job and click **OK**. The filesystem extension is displayed in the Navigation Tree pane.

Editing a file system extension

Learn how to edit a file system extensions policy to change the settings for that policy.

To edit a file system extension policy, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** → **Policy Management** → **Filesystem Extension**.
2. Click the file system extension you want to edit. Information about the file system extension you want to edit appears in the right content pane.
3. In the **Description** field, type a brief description of the job.
4. Select the file systems for the policy:
   a. In the **Available** field, click the file systems you want to include.
   b. Click 

. The file systems are moved to the **Current Selections** field.
5. Click the **Extension** tab; the Extension page opens. Use this page to edit how file systems will be extended. You can choose to extend file systems:
   - By a specific amount of space
   - By a percentage of file system capacity
   - By increasing free space to a specific size
   - By increasing free space to a percentage of file system capacity
   - By increasing capacity to a specific size

Select the option you want to use for this policy. You can also set a maximum capacity size if necessary. If you specify a maximum size, a file system extension job that would increase a file system beyond that size will fail.

Note: The actual amount of file system extension can vary, depending on your environment. For example, if 5GB physical partitions are set up for a volume group, its file systems will be extended in 5GB increments. If you specify a smaller amount, it will be rounded up to the physical partition size.

As another example, software striping and software RAID will claim part of the space created when file systems are extended. In this case, file system extension will appear to create slightly less space than specified in the policy.

If you are using LUN provisioning, you can specify the size of new LUNs by using the Create LUNs that are at least option on the Provisioning tab and one of the Amount to Extend options on the Extension tab. If the Amount to Extend value is greater than or equal to the preferred minimum LUN size, 1% more space will be added to the LUN to cover overhead requirements.
If you choose to extend file systems regardless of remaining free space, the policy options you specify will be applied to all selected file systems when the policy is enforced. To use policy to automatically manage file system capacity, you must do the following:

- Use the **When to Run** page to set the policy to run after every probe or scan
- Use the **Extension** page to specify an amount of free space to be used as a condition for extension. You can do this two ways:
  - Select **Make Freespace** and specify an amount. Then select Extend filesystems regardless of remaining free space. When free space falls below the amount you specify, the difference in space will be added. This option can be used to maintain file system free space at a specific level.
  - Select either **Add** or **Make Freespace** and specify an amount. Then select **Extend filesystems** when free space is less than and specify an amount. The amount you specify will be used to trigger the **Amount to Extend** action you selected. This option can be used to add extra space to a file system if free space reaches a specified level. For example, if free space reaches 10% or less, it could be increased to 20%, or 5GB could be added to the total file system capacity.

  Note: The following restrictions apply when you create or edit a file system extension policy:
  - If you select both **Make Freespace** and **Extend filesystems when freespace is less than**, you must choose the same type of extension for each. If you specify a percentage amount for one, you must specify a percentage amount for the other. If you specify an amount expressed in megabytes for one, you must specify an amount in either megabytes or gigabytes for the other.
  - If you select both **Make Freespace** and **Extend filesystems when freespace is less than**, the extension amount you specify for **Make Freespace** must be larger than the amount you specify for **Extend filesystems when freespace is less than**.
  - If you select **Make Capacity**, you cannot select **Extend filesystems when freespace is less than**. Because the **Make Capacity** option extends capacity regardless of actual utilization, it is not intended to be used for the automated maintenance of file system capacity. You can use this

option to extend file system capacity when necessary by selecting **Run Now** or **Run Once** at on the **When to Run** page.

You can also choose to have no file systems extended when the policy is applied. Instead, any policy actions that would have been performed are written to the log file. This feature, called LOG ONLY mode, can be used to preview the results of a policy over time before actually extending any file systems.

6. Click the **Provisioning** tab; the Provisioning page opens. Use the Provisioning tab to enable LUN provisioning and specify how it will be performed. Provisioning is an optional feature that is not required for file system extension. If you enable LUN provisioning, two things will happen when there is not enough space in a volume group to extend a file system:
   - A new LUN is created to add space to the volume group, according to the policy specified on the **Provisioning** tab.
   - The file system is extended within the volume group, according to the policy specified on the **Extension** tab.

   For example, a volume group might first be extended by 100Mb, and then one or more file systems within it might be extended by 10MB each.

   Data Manager relies on information stored by a CIM/OM to interact with Tivoli Storage Enterprise Storage Server subsystems. To implement provisioning, CIM/OM access information must be defined to Data Manager. See the CIM/OM Logins section in the Configuration and Getting Started Guide for more information. After CIM/OM access is configured, a Discovery job must be run to identify Tivoli Storage Enterprise Storage Server subsystems.

   Note: LUNs can be provisioned for file system hosts running Sun Solaris, but the hosts must be specially configured to avoid a reboot after provisioning. See the appendix on configuring LUN provisioning for Sun Solaris in the IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide for instructions.

   If you choose to enable automatic LUN provisioning, you must choose a model and location for new LUNs. Data Manager models new LUNs on existing LUNs. Modeled characteristics include CIM/OM access parameters, as well as connectivity and RAID type. You must specify how far Data Manager will look for a model LUNs, as well as limit where it will create new LUNs on the Tivoli Storage Enterprise Storage Server.

   Note: To perform LUN provisioning, at least one Tivoli Storage Enterprise Storage Server LUN must be currently assigned to the volume group of the file system you want to extend, or to the managed host associated with the file system you want to extend. Select one of the following options to limit the scope of LUN modeling:
   - Model new LUNs on others in the volume group of the file system being extended. If you select this option, Data Manager will only search for a model within existing LUNs in the volume group of the file system being extended.
   - Model new LUNs on others on the same host as the file system being extended. If you select this option, Data Manager will first search for a model within existing LUNs in the volume group of the file system being extended, before looking further on the same host.

New LUNs can be created in the same Storage Pool as the file system being extended, or in any Storage Pool in the Tivoli Storage Enterprise Storage Server of the file system being extended. Select one of the following options to limit where new LUNs will be created:

- Provision new LUNs within the same Storage Pool as the file system being extended. For Tivoli Storage Enterprise Storage Server, a Storage Pool is a Volume Space. If you are using RAID, a Volume Space is equal to one RAID rank, which consists of a Disk Group (8 physical disks).
- Provision new LUNs within the same Storage System as the file system being extended. If you select this option, new LUNs can be created in any Storage Pool in the Tivoli Storage Enterprise Storage Server. Data Manager will first attempt to provision a new LUN in the same Storage Pool as the file system being extended, before looking further in the same Storage System.

You can optionally specify a preferred minimum size for new LUNs by selecting the **Create LUNs that are at least** option. If you select this option, LUNs will not be created any smaller than the size you specify. If you do not specify a preferred minimum size, the Amount to Extend specified for the policy will be used. For Tivoli Storage Enterprise Storage Server subsystems, the default minimum LUN size increment is 100 MB. New LUNs will be rounded up to satisfy this requirement. If the Amount to Extend value is greater than or equal to the preferred minimum LUN size you specify, 1% more space will be added to the LUN to cover overhead requirements. See Table 19.

Note: Data Manager will only provision new LUNs. Existing unassigned LUNs will not be reused. Similarly, if provisioning fails after LUNs have been created, the new LUNs will not be deleted or reused. This limitation protects against accidentally overwriting existing data that might be associated with currently unassigned LUNs.
Table 19 shows the possible outcomes when a new LUN is provisioned.

Table 19. LUN Provisioning Results

Description	Amount to Extend	Tivoli Storage Enterprise Storage Server minimum for new LUNs	Create LUNs that are at least (preferred minimum LUN size)	Actual LUN size that will be created	Actual amount by which the file system will be extended
Minimum LUN size exceeds amount to extend	20Mb	100 Mb	200Mb	200Mb	20Mb
Amount to extend equals preferred minimum LUN size	200Mb	100 Mb	200Mb	202Mb (1% added for overhead)	200Mb
Minimum LUN size exceeds amount to extend	100Mb	100 Mb	200Mb	200Mb	100Mb

Table 19. LUN Provisioning Results (continued)

Description	Amount to Extend	Tivoli Storage Enterprise Storage Server minimum for new LUNs	Create LUNs that are at least (preferred minimum LUN size)	Actual LUN size that will be created	Actual amount by which the file system will be extended
Amount to extend exceeds preferred minimum LUN size, but both are below Tivoli Storage Enterprise Storage Server minimum	75Mb	100 Mb	50Mb	100Mb	75Mb
Minimum LUN size not specified	200Mb	100 Mb	Not specified	202Mb (1% added for overhead)	200Mb

Note: IBM Tivoli Storage Enterprise Storage Server and theTivoli Storage Enterprise Storage Server CIM/OM use multiples of 1000 bytes to represent kilobytes, megabytes, and gigabytes. Data Manager uses multiples of 1,024 bytes. Data Manager automatically converts this difference when it discovers Tivoli Storage Enterprise Storage Server capacity and utilization. If you are using IBM Tivoli Storage Tivoli Storage Enterprise Storage Server Expert or Tivoli Storage Enterprise Storage Server Specialist, the storage information they show will differ from what is shown by Data Manager.

Due to AIX Volume Manager limitations on AIX Striped Volumes, provisioning of new LUNs cannot be done. If the policy definition for a striped file system indicates to Enable Automatic LUN Provisioning on the Provisioning tab, then any LUNs that are created will be added to the Volume Group but will not be usable for Extension. If provisioning is attempted on an AIX Striped Volume, then you will see the message: `STA0301W: To prevent runaway condition, we have removed the file system file system_name from the policy, please fix the problem and add the file system back to the policy.`

To avoid this problem, make sure the **Enable Automatic LUN Provisioning** option is disabled for Filesystem Extension policy definitions that have striped file systems on AIX.

7. Click the **When to Run** tab; the When to Run page opens. Use this page to edit when to apply the file system extension policy to selected file systems.
8. In the **How often to run** area, schedule how often you want to enforce the policy. You can choose from the following actions:

   Enforce policy after every Probe or Scan
   Enforce the policy after every probe or scan that is run against the selected filesystems. If you want to automate filesystem extension, you should choose to apply the policy after every new probe or scan.

   Run Now
   Enforce the policy immediately.

   Run Once at
   Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

Note: TheTivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

Note: If the policy applies to multiple filesystems and you select **Run Now** or **Run Once at**, messages specific to each filesystem will be mixed together in the log file. This can make it difficult to troubleshoot errors. To avoid this, you can select one policy at a time and then select **Run Now**.

9. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
   a. 9:00 AM of the time zone where the Data server resides
   b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

   Use the timezone that the server runs in
   Select this option to use the time zone of the location where the Data server resides.

   Use this timezone
   Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.
10. Click the **Alert** tab; the Alert page opens. Use this page to edit the alerting criteria for a filesystem extension.
11. In the **Triggering Condition** area, specify the alerting criteria for the filesystem extension. In the **Condition** field, the default setting is **A filesystem extension action started automatically**. No other conditions are available. When a filesystem is extended by the policy, an alert will be triggered. When the alert is triggered, the actions you specified in the Triggered Actions section are taken.
12. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

    TEC event
    A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

    Login notification
    The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

Login ID
The ID of the user who will receive alert notification upon logging in to the system.

Windows Event Log, UNIX Syslog
The ID of the user who will receive alert notification upon logging in to the system.

Event Type (Windows Event Log only)
The type of event that will be recorded to the OS log.

Facility (UNIX Syslog only)
The location where the UNIX Syslog will be sent. You can select User or Local.

Level (UNIX Syslog only)
The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

13. Click **File** → **->Save** to save the filesystem extension.

Scheduling script-based actions

Use scheduled actions to run user-defined scripts against selected computers and computer groups.

You can access the scheduled actions facility using either Data Manager or Data Manager for Chargeback. The actions that you schedule to run are based on user-defined scripts, allowing you to use any third-party tools for actions, such as recovery or provisioning. You can schedule scripts to run:

- Immediately
- Once at a specified time and date
- Repeatedly according to a schedule you specify

Select the option you want and then edit the time and date information accordingly. You can also select the time zone for the schedule you specify. You can either select the time zone that the Data server is in or any other time zone.

Note: Scheduled actions should not be confused with triggered actions. Triggered actions enable you to run scripts based upon conditions and alerts that are detected within your storage environment. Scheduled actions enable you to run scripts according to the schedule that you define, without having to associate the script with an alert.

Running scripts using scheduled actions

Use the scheduled action facility to run a script according to a schedule that you define. The scheduled action facility is available in Data Manager and Data Manager for Databases.

Use scheduled actions to run scripts against selected computers and computer groups. The scheduled actions facility is script-based, allowing you to use any third-party tools for actions, such as recovery, or provisioning. You can schedule scripts to run:

- Immediately
- Once at a specified time and date
- Repeatedly according to a schedule you specify

Scheduled actions should not be confused with triggered actions. Triggered actions enable you to run scripts based upon conditions and alerts that are detected within your storage environment. Scheduled actions enable you to run scripts according to the schedule that you define, without having to associate the script with an alert.

To create a scheduled action, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** → **Policy Management** → **Scheduled Actions** or **Data Manager for Databases** → **Policy Management** → **Scheduled Actions**.
2. Right-click **Scripts**, and click **Create Script**. The Create Script window opens.
3. In the **Description** field, type a brief description of the job.
4. Select computers and computer groups against which to run the action:
   a. In the **Available** field, click the computers and computer groups against which you want to run the action. You can select individual computers as well as computer groups that you have previously defined using the group facility. Note that cluster resource groups in an HACMP or MSCS clustered environment are listed as computers.
   b. Click  . The computers and computer groups are moved to the **Current Selections** field.
5. Click the **Script Options** tab; the Script Options page opens. Use this page to select the script you want to run against the selected computers and computer groups.
6. Select a script from **Script Name**. The scripts that appear in **Script Name** are stored in the `\<TPC_installation_directory>\scripts` directory on the Data server. The default installation directory where scripts are located is:
   - Windows: `\program files\IBM\TPC\data\scripts`
   - UNIX, Linux: `/opt/IBM/TPC/Data/scripts` or `/usrs/IBM/TPC/Data/scripts`

   Note: The scripts that appear in the `Script Name` list are stored on the computer where the Data server component is located. If you want to run a script that is stored on a Data agent, you must type the name of that script in the `Script Name` field.
7. Click the **When to Run** tab; the When to Run page opens. Use the When to Run page to specify when the scheduled action will run.
8. In the **How often to run** area, schedule how often you want to run the script associated with the action. You can choose from the following options:

   Run Now
   Run the script immediately.

   Run Once at
   Run the script once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

Note: TheTivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

Run Repeatedly
Run the script repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark will appear next to the days you select.

9. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
   a. 9:00 AM of the time zone where the Data server resides
   b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

   Use the timezone that the server runs in
   Select this option to use the time zone of the location where the Data server resides.

   Use this timezone
   Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.
10. Click the **Alert** tab; the Alert page opens. Use the Alert tab to define the alerting criteria for a scheduled action.
11. In the **Triggering Condition** area, the only condition is **Script Failed**. This indicates that an alert will automatically be triggered if the scheduled action fails to run.
12. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition. For an explanation of triggered actions, see Appendix F, "Triggered actions for alerts," on page 801.
13. Save the scheduled action:
    a. Click **File** → **Save**. The Save As window opens.
    b. In the **Specify Script name** field, type a name for the job and click **OK**. The scheduled action is displayed in the Navigation Tree pane.

Editing a scheduled action

Edit a scheduled action to change: the computers and groups against which the action is run, the script that is run, when the script is run, and the alerting conditions if the action fails to run.

To edit a scheduled action, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** → **Policy Management** → **Scheduled Actions** → **Scripts** or **Data Manager for Databases** → **Policy Management** → **Scheduled Actions** → **Scripts**.

2. Select the scheduled action that you want to edit. Information about the action appears in the right content pane.
3. In the **Description** field, type a brief description of the job.
4. Select computers and computer groups against which to run the action:
   a. In the **Available** field, click the computers and computer groups against which you want to run the action. Select individual computers, as well as computer groups that you have previously defined using the group facility. Note that cluster resource groups in an HACMP or MSCS clustered environment are listed as computers.
   b. Click 

. The computers and computer groups are moved to the **Current Selections** field.
5. Click the **Script Options** tab; the Script Options page opens. Use this page to change the script you want to run against the selected computers and computer groups.
6. Select a script from **Script Name**. The scripts that appear in **Script Name** are stored in the \<*TPC_installation_directory*>\scripts directory on the Data server. The default installation directory where scripts are located is:
   - Windows: \program files\IBM\TPC\data\scripts
   - UNIX, Linux: /opt/IBM/TPC/Data/scripts or /usrs/IBM/TPC/Data/scripts

   Note: The scripts that appear in the Script Name list are stored on the computer where the Data server component is located. If you want to run a script that is stored on a Data agent, you must type the name of that script in the Script Name field.
7. Click the **When to Run** tab; the When to Run page opens. Use this page to change when the scheduled action is set to run.
8. In the **How often to run** area, schedule how often you want to run the script associated with the action. You can choose from the following options:

   Run Now
   Run the script immediately.

   Run Once at
   Run the script once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

   Note: TheTivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

   Run Repeatedly
   Run the script repeatedly according to a schedule you specify. You can:
   - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
   - Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
   - Indicate which days on which run the job. A check mark will appear next to the days you select.

9. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
   a. 9:00 AM of the time zone where the Data server resides
   b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

   Use the timezone that the server runs in
   Select this option to use the time zone of the location where the Data server resides.

   Use this timezone
   Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.
10. Click the **Alert** tab; the Alert page opens. Use the Alert tab to change the alerting criteria for a scheduled action.
11. In the **Triggering Condition** area, the only condition is **Script Failed**. This indicates that an alert will automatically be triggered if the scheduled action fails to run.
12. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

    SNMP trap
    An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

    TEC event
    A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

    Login notification
    The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

    Login ID
    The ID of the user who will receive alert notification upon logging in to the system.

    Windows Event Log, UNIX Syslog
    The ID of the user who will receive alert notification upon logging in to the system.

    Event Type (Windows Event Log only)
    The type of event that will be recorded to the OS log.

    Facility (UNIX Syslog only)
    The location where the UNIX Syslog will be sent. You can select User or Local.

    Level (UNIX Syslog only)
    The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

13. Click **File** → **Save** to save your changes.

Archive and backup functions

Data Manager provides a method within the alerting function of a constraint by which you can automatically invoke a IBM Tivoli Storage Manager archive or backup to run against the files that violate that constraint.

An archive is a ″snapshot in time″. Each time an archive is done, a new ″snapshot″ is taken. When an archive is performed, everything that has been specified is captured. So, for example, if an archive was performed on a C: drive, everything on C: will be archived. If, the next day, another archive of C: was performed, everything on the C: drive would be archived again. Files and file systems that have been archived as part of the Archive/Backup job can be deleted after archiving, but you must select **Delete After Successful Archival** specifically.

Lets say that, as a result of running a report, you find a large number of files that haven't been accessed in over two years. Looking at the cumulative sizes of these files you determine they are taking up space that would be better utilized elsewhere. You can archive and delete these files to reclaim that space. If, at some time in the future those files are needed, they can be retrieved from the archive. The result of the archive job is more free space without data loss.

A backup can be incremental (**Selective Backup**) or full. Usually a full backup is done first and subsequent backups are incremental. For example, say that a backup is performed on C: drive. All the files on the C: drive are then backed up, which is the full backup operation. The next day we need to back up C: drive again. Usually, this will be an incremental backup since a full backup was done the day before. In the incremental backup, only the files that have changed on the drive since the last backup will be backed up.

The archived or backed up files are sent to theTivoli Storage Manager server and are stored on that machine. Files are put into storage groups that are controlled through a Tivoli Storage Manager policy that the Tivoli Storage Manager server administrator usually creates.

Constraints

Data Manager provides a method within the alerting function of a constraint by which you can automatically invoke a Tivoli Storage Manager archive or backup to run against files that violate constraints you have defined. You can do the following:

- Define a constraint that will invoke Tivoli Storage Manager to archive and (if that option is selected) delete those files
- Define a constraint to run an incremental or full backup

See the scenario *Archive and delete files* to see an example of how constraints are developed.

Before you begin

Before you can use the Tivoli Storage Manager archive/backup functionality with Data Manager constraints, you must ensure the following:

- Your user ID has Administrator authority within IBM Tivoli Storage Productivity Center.
- Tivoli Storage Manager client software is installed on each of the hosts where a Data Manager scanning agent will be used to perform archive and backup functions.
- (for SAN attached storage only) Tivoli Storage Manager client software must be installed on the Data agent host machine that is configured to act as the scanning agent for NAS storage.
- A Tivoli Storage Manager client must be configured to run without a password. This is typically done by either including the password in the client options (dsm.opt) file, or using the options PASSWORDACCESS= GENERATE in the client options file.
- The Tivoli Storage Manager client must be version 5.1 or higher.

Create an archive/backup job

Create a IBM Tivoli Storage Manager archive/backup job for files in a report

You follow these general steps when you create an archive or backup.

1. Make sure you have followed the prerequisite steps
2. Decide whether to create a new constraint or use one of the default constraints
3. Generate the report using the constraint you decided to use
4. Create an archive or backup job from files in the report

To create an Archive/Backup job on files from a report, complete the following steps:

1. Generate and view a Largest Files, Most Obsolete Files, Orphan Files, Duplicate Files, Most At Risk, Constraint Violations report, or a report that you have created.
2. Click on the files you want to include in the archive or backup job. Use ctrl + click to select multiple, non-consecutive files; use shift + click to select multiple, consecutive files.
3. Right click on the reports and select **Create a new archive/backup job** from the pop-up menu. The **Create File Archive/Backup Action** window is displayed.
4. In the **Description** field, type a brief description of the job.
5. In the **Archive/Backup Options** tab, select the type of the job and its TSM command options, if any:

   Archive
   Use this option to archive the files that violate a constraint. Archives are copies of data from online storage copies to offline storage.

   Delete After Successful Archival
   Automatically delete the files that violated the constraint after they have been successfully archived. See the *IBM Tivoli Storage Manager Client User's Guide* for more information about Tivoli Storage Productivity Center archives.

Selective Backup
Select this option to perform a selective backup of the files that violate a constraint. A selective backup enables you to back up specified files. These files are not excluded in the include-exclude list and meet the requirement for serialization in the backup copy group of the management class assigned to each file. See the *IBM Tivoli Storage Manager Client User's Guide* for more information about selective backups.

Incremental Backup
Select this option to perform an incremental backup of the files that violate a constraint. An incremental backup enables you to back up new or changed files or directories from a client domain or from specified directories or files. These directories or files are not excluded in the include-exclude list and meet the requirements for frequency, mode, and serialization as defined by a backup copy group of the management class assigned to each file. See the *IBM Tivoli Storage Manager Client User's Guide* for more information about incremental backups.

IBM Tivoli Storage Manager Command Options
Enter additional Tivoli Storage Manager command line options in this field that will be passed on to the Tivoli Storage Manager dsmc command. You can enter up to 512 bytes of text. See the *IBM Tivoli Storage Manager Client User's Guide* for more information about valid Tivoli Storage Manager command-line options.

6. Click the **Files** tab; the **Files** page opens. Use this page to view or remove the files included in an archive/backup job.
7. Click the **When to Run** tab. Use the **How often to run** area to specify when the archive or backup job runs.
8. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

Run Now
Check the quota immediately.

Run Once at
Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

Note: The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

Run Repeatedly
Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark will appear next to the days you select.

9. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
   a. 9:00 AM of the time zone where the Data server resides
   b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

   Use the timezone that the server runs in
   Select this option to use the time zone of the location where the Data server resides.

   Use this timezone
   Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.
10. Click the **Alert** tab; the Alert page opens. Use the Alert page to define the alerting criteria for an archive or backup job.
    a. In the **Triggering Condition** area, the only condition is **File Archive/Backup Failed**. This indicates that an alert will automatically be triggered if the archive or backup fails to run.
    b. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition You can choose from the following actions:

       SNMP trap
       An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

       TEC event
       A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, etc.

       Login notification
       The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

       Login ID
       The ID of the user who will receive alert notification upon logging in to the system.

       Windows Event Log, UNIX Syslog
       The ID of the user who will receive alert notification upon logging in to the system.

       Event Type (Windows Event Log only)
       The type of event that will be recorded to the OS log.

       Facility (UNIX Syslog only)
       The location where the UNIX Syslog will be sent. You can select User or Local.

       Level (UNIX Syslog only)
       The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

Run script
A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window will be passed to the script when the alert is triggered

E-mail An e-mail is sent. Click **Add** or **Delete** to edit the list of e-mail addresses to which an e-mail is sent. Click **Edit e-mail** to change the actual e-mail message.

11. Save the archive or backup:
    a. Click **File** → **Save**. The **Save As** window opens.
    b. In the **Specify File Archive/Backup Action name** field, type a name for the job and click **OK**. The new job will be listed in **Data Manager** → **Policy Management** → **Archive/Backup** as **<creator>.<job name>**.

View and edit an Archive/Backup job

Learn how to edit IBM Tivoli Storage Manager archive and backup jobs that were defined to run against the files selected from Data Manager reports or from reports you have generated yourself.

Each archive and backup that you have created appears as a node under **Archive/Backup**.

To view and edit a Tivoli Storage Manager archive or backup job, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager** → **Policy Management** → **Archive/Backup**. Select the archive or backup job that you want to edit. Information about the job appears in the right content pane. All the user-defined information can be edited.
2. Change the description for the archive/backup job.
3. In the **Options** tab, you can select the Archive/Backup options. You can change whether the job is an archive or a backup, and add or change any additional Tivoli Storage Manager commands that should be run with the job.
4. In the **Files** tab the files that were included in the job are listed.
5. In the **When to Run** tab, select how often to run and how to handle time zones. You can change the **How often to run** and **How to handle time zones** settings.
6. In the **Alert** tab, **File Archive/Backup Failed** is the only **Triggering-Condition**. In **Triggered-Actions**, change the action you wish to occur if the job fails. You do not need to select any of the actions. Notification will be written to the error log by default.
7. Select **File** → **Save**. Name the Archive/Backup job. After you save a the job, an entry for it will appear under the **Data Manager** → **Policy Management** → **Archive/Backup** → **<Creator>.<Name>**.

Add additional files to an archive/backup job

Add files to an existing IBM Tivoli Storage Manager archive or backup job from a report.

To add files to an existing Archive/Backup job, complete the following steps:

1. Add files to an existing Tivoli Storage Manager archive or backup job from a report

2. Click on the files you want to add to an archive or backup job. Use ctrl + click to select multiple, non-consecutive files; use shift + click to select multiple, consecutive files.
3. Right click on the report and select **Add to an existing archive/backup job** from the pop-up menu. The Archive/Backup Job Selection window appears.
4. Select the archive or backup job to which you want to add the files from the list box.
5. Click **OK** to add the files. The Archive/Backup - Files window appears and shows the file(s) you added to the selected job.

Working with storage resource groups

Use storage resource groups to organize logically related storage entities into named groups. For example, a storage resource group that represents a business critical email application might include all the hosts that participate in the email application cluster, the storage subsystems that provide storage to the application, and the switches and fabrics through which the application's data is configured to travel.

You can use storage resource groups as input to the SAN Planner for planning and provisioning considerations. For example, if a storage resource group represents a business critical email application, you can use that group as input into the SAN planner to help consistently determine when and where to add capacity to the application while minimizing the manual entry of storage entities.

Storage resource group management

Use the storage resource group management list panel to view, create, and delete storage resource groups.

This panel contains the following information about each storage resource group:

Storage resource group name
Specifies a user-defined name of the storage resource group.

State Specifies the state of the storage resource group:

Normal
The storage resource group is operating normally.

Warning
At least one part of the storage resource group is not operating or has serious problems.

Critical
The storage resource group is either not operating or has serious problems.

Missing
The storage resource group was previously discovered by IBM Tivoli Storage Productivity Center, but was not detected in the most recent discovery.

Unknown
The storage resource group was discovered but is not recognized by IBM Tivoli Storage Productivity Center.

Note: The state of a storage resource group might appear as **Unknown** after it is created. This occurs because the status of a

new storage resource group might not yet be updated when the Storage Resource Group Management panel is re-displayed. To resolve this issue, exit the Storage Resource Group Management panel, wait for the processing of the group to complete, and open the Storage Resource Group Management panel again to view the updated **State**.

Lea

Description (optional)
Displays the user defined description for the storage resource group.

Create Creates a storage resource group that you selected.

Delete
Deletes a storage resource group that you selected.

Create or Edit a Storage Resource Group

Use the Create or Edit Storage Resource group panel to create or modify storage resource groups.

To create a storage resource group:

- In the navigation tree, expand **IBM Tivoli Storage Productivity Center**, click **Storage Resource Group Management**, then click **Create...** in the content pane.
- In the navigation tree, expand **IBM Tivoli Storage Productivity Center** > **Topology**, click a node that represents a storage entity, right-click an entity that appears in the content pane, and select **Add to new Storage Resource Group...** from the drop down menu.

To edit a storage resource group:

- In the navigation tree, expand **IBM Tivoli Storage Productivity Center**, click **Storage Resource Group Management**, then click the magnifying glass icon next to the storage resource group you want to edit.
- In the navigation tree, expand **IBM Tivoli Storage Productivity Center** > **Topology**, click **Storage Resource Groups**, right-click the storage resource group you want to edit in the content pane, and select **Launch Detail Panel** from the drop down menu.

Field Descriptions

Creator
Displays the user name of the creator.

Name Displays the name of the storage resource group or unnamed, if it is not yet named.

Description
Optional: Displays the user defined description for the storage resource group.

Selected Elements
Lists the elements selected to be members of this storage resource group.

Add Adds one or more selected elements to the list.

The Storage resource group element selection panel is displayed.

Remove
Removes one or more selected elements from the list.

Default Provisioning Profile
Lists the available provisioning profiles which can be associated with storage resource groups. The list also includes None. If this storage resource group is used as input to the SAN Planner, the settings defined in this profile will be used to pre-populate the planner inputs.

Create a New Profile
Launches the Provisioning Profile creation wizard. When you complete the wizard, the **Provisioning Profile** list is updated.

User defined property 1 (UDP1)
Specifies any user-defined properties that will be used by the topology viewer to provide custom groupings.

User defined property 2 (UDP2)
Specifies any user-defined properties that will be used by the topology viewer to provide custom groupings.

User defined property 3 (UDP3)
Specifies any user-defined properties that will be used by the topology viewer to provide custom groupings.

Storage Resource Group Element Selection

Use the Storage Resource Group Element Selection panel to select the storage entities that you want to include in a storage resource group. The storage entities that you can select are displayed in the topology viewer.

Available elements section

Displays the storage entities that you can include in a storage resource group. Use the topology viewer to select storage entities. See Overview of Topology Viewer for more information about the topology viewer.

Topology viewer
Use the topology viewer to select the storage entities that you want to include in a storage resource group. To do this, perform the following steps.

1. Open a detailed view of a storage entity type on the topology viewer's **Overview** page.
2. Expand the storage entity type to view the storage entities that are part of that type.
3. Select the storage entities you want to include in the storage resource group.
4. Click >>.

Selected elements section

>> Click >> to move the selected subsystems into the **Selected Elements** pane.

<< Click << to remove the selected subsystems from the **Selected Elements** pane.

Selected Elements
Displays the storage entities that are included in a storage resource group. You can include the following entities in a storage resource group:

- computer (vm, physical or hypervisor)
- computer disk (including vm disks)
- fabric

- switch
- other entities
- storage subsystem
- storage pool
- subsystem disk
- subsystem volume
- tape drive
- tape library
- tape media changer
- other storage resource groups

OK Click **OK** when you are satisfied with all your selections. The Create Storage Resource Group panel displays with your selections.

Cancel
Click **Cancel** if you do not want to add storage entities to a storage resource group. The Create Storage Resource Group panel displays. No data is saved.

Working with scripts

Use the triggered action facility to run scripts based upon conditions/alerts that are detected within your storage environment. Use the scheduled action facility to run scripts according to the schedule that you define, without having to associate the script with an alert.

Both the triggered action and scheduled action facilities are script-based, allowing you to use any third-party tools for actions, such as recovery and provisioning.

Running scripts with alerts

Monitoring, alerting, and policy management jobs enable you to use the triggered action facility to run scripts based on storage events that you define.

The triggered action facility allows you to define thresholds and constraints that result in notification or action. A script can be run when a job fails, an alert condition is detected in a storage resource, a constraint is violated (Data Manager only), or a quota is exceeded (Data Manager only). The triggered action facility is script-based, allowing you to use any third-party tools for actions such as recovery or provisioning.

Perform the following steps to run a script in response to a condition in a job:

1. Create or edit a monitoring, alerting, or policy management job. For example, the following jobs enable you to define a condition that triggers a script:
   - Data Manager:
     - **Monitoring → <all_monitoring_jobs>**
     - **Alerting → <all_alert_jobs>**
     - **Policy Management → Constraints**
     - **Policy Management → Quotas**
     - **Policy Management → Scheduled Actions**
   - Data Manager for Databases:
     - **Monitoring → <all_monitoring_jobs>**
     - **Alerting → <all_alert_jobs>**

- **Policy Management → Quotas**
- **Policy Management → Scheduled Actions**

• Disk Manager:
- **Monitoring → Subsystem Performance Monitors**
- **Alerting → Storage Subsystem Alerts**

• Fabric Manager:
- **Monitoring → Switch Performance Monitors**
- **Alerting → <all_alert_jobs>**

2. Check **Run Script** on the **Alert** page for the job.
3. Click **Define**. The **Specify Script** dialog box displays.
4. Enter the name of the script in the **Script Name** field. This script name is used to tell the server what script to send to the agent. If the script sent by the server is run by the agent, the name of the script will not be the name you specify in the **Script Name** field. Instead, the agent names the script as temporary file, runs it, then deletes it after the script has run. Because of this, you should not have code in your script that depends on this script name.
5. Specify where you want the script to run in the **Where to Run** field. You can select a specific agent, or the triggering computer. For example, if a file system alert is triggered on a on a specific computer, you would typically want the script to run on the computer that triggered the alert. (For some types of alerts such as *Job Failed*, you will not have the option to run the script on the triggering computer. Here you can run the script on any agent you choose.)
6. The **Script Parameters** section displays what parameters will be passed to the script when it is run. Each triggering condition passes different parameters. The parameters will be passed in the order specified on the dialog box.
7. Click **OK**.

Running scripts using scheduled actions

Use the scheduled action facility to run a script according to a schedule that you define. The scheduled action facility is available in Data Manager and Data Manager for Databases.

Use scheduled actions to run scripts against selected computers and computer groups. The scheduled actions facility is script-based, allowing you to use any third-party tools for actions, such as recovery, or provisioning. You can schedule scripts to run:

- Immediately
- Once at a specified time and date
- Repeatedly according to a schedule you specify

Scheduled actions should not be confused with triggered actions. Triggered actions enable you to run scripts based upon conditions and alerts that are detected within your storage environment. Scheduled actions enable you to run scripts according to the schedule that you define, without having to associate the script with an alert.

To create a scheduled action, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager → Policy Management → Scheduled Actions** or **Data Manager for Databases → Policy Management → Scheduled Actions**.
2. Right-click **Scripts**, and click **Create Script**. The Create Script window opens.

3. In the **Description** field, type a brief description of the job.
4. Select computers and computer groups against which to run the action:
   a. In the **Available** field, click the computers and computer groups against which you want to run the action. You can select individual computers as well as computer groups that you have previously defined using the group facility. Note that cluster resource groups in an HACMP or MSCS clustered environment are listed as computers.
   b. Click 

. The computers and computer groups are moved to the **Current Selections** field.
5. Click the **Script Options** tab; the Script Options page opens. Use this page to select the script you want to run against the selected computers and computer groups.
6. Select a script from **Script Name**. The scripts that appear in **Script Name** are stored in the \<*TPC_installation_directory*>\scripts directory on the Data server. The default installation directory where scripts are located is:
   - Windows: \program files\IBM\TPC\data\scripts
   - UNIX, Linux: /opt/IBM/TPC/Data/scripts or /usrs/IBM/TPC/Data/scripts

   Note: The scripts that appear in the Script Name list are stored on the computer where the Data server component is located. If you want to run a script that is stored on a Data agent, you must type the name of that script in the Script Name field.
7. Click the **When to Run** tab; the When to Run page opens. Use the When to Run page to specify when the scheduled action will run.
8. In the **How often to run** area, schedule how often you want to run the script associated with the action. You can choose from the following options:

   Run Now
   Run the script immediately.

   Run Once at
   Run the script once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

   Note: TheTivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

   Run Repeatedly
   Run the script repeatedly according to a schedule you specify. You can:
   - Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
   - Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
   - Indicate which days on which run the job. A check mark will appear next to the days you select.
9. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:

a. 9:00 AM of the time zone where the Data server resides
b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

Use the timezone that the server runs in
Select this option to use the time zone of the location where the Data server resides.

Use this timezone
Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field will activate.

10. Click the **Alert** tab; the Alert page opens. Use the Alert tab to define the alerting criteria for a scheduled action.
11. In the **Triggering Condition** area, the only condition is **Script Failed**. This indicates that an alert will automatically be triggered if the scheduled action fails to run.
12. In the **Triggered Actions** area, select the actions that will occur as a result of the triggering condition. For an explanation of triggered actions, see Appendix F, "Triggered actions for alerts," on page 801.
13. Save the scheduled action:
    a. Click **File** → **Save**. The Save As window opens.
    b. In the **Specify Script name** field, type a name for the job and click **OK**. The scheduled action is displayed in the Navigation Tree pane.

How scripts are run

Scripts can be run from the Data server or the Data agents.

During installation, IBM Tivoli Storage Productivity Center creates a directory on every server and agent computer called: \<*install_directory* >\scripts, where <*install_directory* > represents the directory where the agent or server is installed. The default installation directory is:

- (Windows Data agents) `\program files\IBM\TPC\ca\subagents\TPC\Data\scripts`
- (UNIX, Linux Data agents) `/opt/IBM/TPC/ca/subagents/TPC/Data/scripts` or `/usr/tivoli/ep/subagents/TPC/Data/scripts`
- (Windows Data server) `\program files\IBM\TPC\Data\scripts`
- (UNIX, Linux Data server) `\opt\IBM\TPC\Data\scripts` or `/usr/IBM/TPC/Data/scripts`

When a script is run:

1. The server looks in its local `\scripts` directory.
2. If the server can resolve that script name, it loads that script and sends it to the agent where you designated the script to be run.

   When running a script against a NAS filer, the script is run from the agent assigned to the file system where the triggering condition occurred.
3. The agent receives the script, names it as a temporary file, and runs it.
4. After the agent has finished running the script, the temporary file is deleted.

There are two main reasons to run a script:

- In response to an alert

- For batch reports

There are two reasons why a script may not run when sent by the server:

- The script already exists in an agent's \scripts directory. In this case, the agent ignores the script sent by the server and just run its local version.
- If you did not check the **Agent may run scripts sent by server** option in the installation parameters screen, the agent ignores the script sent by the server.

  If you selected the **Agent may run scripts sent by server** option in the installation parameters screen, the agent will run scripts sent by the server.

  The advantage of selecting this option is you need only store one copy of the script in a server's \scripts directory and do not have to keep a copy of the script on every agent computer. When a script needs to be run on a particular agent, the server will access the script from it's local \scripts directory and send it to the agent to be run.

  Note: You can change this option by editing the honorSentScripts parameter in the agent.config file. The agent.config file exists in every agent's installation directory. See *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for information about the agent.config file.

When the script is pushed out to the agent, it is named as temporary file. After the agent has finished running the script, the temporary file is deleted.

If the **Agent may run scripts sent by server** option is unchecked, you must make sure that the script is stored in every agent's \scripts directory. If the **Agent may run scripts sent by server** option is checked, the server will push out the script to the Data agent computer that has the low freespace condition. In this case, you only have to ensure that the script exists in a server's \scripts directory, and not on any of the agents \scripts directory.

Note: We recommend storing the scripts you want to run globally (across all your agents) on the server. Store the scripts you want to run locally on the agent computers where you want them to run.

How the server and agent resolve script names

An IBM Tivoli Storage Productivity Center server processes requests for agents running under UNIX, Linux, and Windows.

When associating and running scripts with triggered and scheduled actions, the Tivoli Storage Productivity Center server resolves the names of those scripts based on the following criteria:

- The OS under which an agent is running
- How the script name is specified on the Specify Script dialog box

Agent and Operating System

The operating system under which an agent runs determines the scripts that can be run by that agent:

- **UNIX, Linux**. An agent running under a UNIX or Linux operating system will not run or receive (from the server) scripts that have an extension. Agents running under UNIX or Linux only accept and run script files that do not contain an extension within their name. The first line in a script file determines what interpreter is needed to run the script.

- **Windows**. An agent running under a Windows operating system will only run or receive (from the server) scripts that have an extension. The extension of a script file determines what interpreter is needed to run the script.

Specifying a Script Name

Use the Specify Script dialog box when associating a script with a triggered or scheduled action. Enter the name of a script on the **Script Name** field to associate that script with the action you want to perform. The name of the script you enter might or might not contain an extension (including an extension is optional). Depending on the agent that receives the script, the following occurs after a triggered or scheduled action that contains a script is run:

- **Agent running under UNIX or Linux**. If the specified script name contains an extension, the agent ignores the extension and searches for a file of the same name (without an extension). For example, if you enter `backup.vbs` in the **Script Name** field, a UNIX or Linux agent will search for a file with the name `backup`. If the script exists, UNIX or Linux will read the first line of the script and use the appropriate interpreter to run the script. If the file does not exist, the triggered or scheduled action will not run the script.
- **Agent running under Windows**. If the specified script name contains an extension, the agent will search for a file of the same name/extension. For example, if you enter `backup.vbs` in the **Script Name** field, an agent will search for a file with the name `backup.vbs`. If the script does exist, the agent will determine what interpreter is needed to run the script based on its extension. If the file does not exist, the triggered or scheduled action will not run the script.

  If the specified script name does not contain an extension, an agent will look for the file name in the following order (based on its extension): `.com, .exe, .bat, .cmd, .vbs, .vbe, .js, .jse`

The process for script name resolving enables you to run scripts with the same name across multiple operating systems. To run a script on both a UNIX or Linux and Windows operating systems, you must have two versions of that file on either the server or the agent. For example, if you want to run a provisioning script against a UNIX or Linux and Windows agent, you must have two versions of that file (`provision.bat` and `provision`) stored in the `\scripts` directory on either the server or agent.

Using the topology viewer

This topic describes how to use the topology viewer to view detailed information about your storage environment and monitor entities and groups of entities.

You can view health status, performance, and other details about switches, computers, storage subsystem servers, fabrics, and other entities. To launch the topology viewer, expand **IBM Tivoli Storage Productivity Center Topology** in the Navigation Tree pane.

To display detailed information about an entity (for example, the host name of a computer, the used and unused capacity of a storage subsystem), right-click the entity and select **Launch Detail Panel**.

To select an entity in the graphical view, click the entity image. In the table view, click the name of the entity. To make multiple selections, press CTRL and click.

Note:

- The availability of certain action items depends on the entity class and view level.
- When you right-click an object in the graphical view, a context menu is displayed. The action items in this menu are identical to the action items displayed when you right-click the corresponding object in the table view and click the arrow in the **Action** field.
- If the discovery process is not completed and there are entities in the environment that have not been rendered into the current view, the view will build dynamically as entities are discovered.
- Some columns and fields displayed within the Topology Viewer might not be populated with data for the storage resources that you are monitoring. Columns and fields appear empty if you have not configured Tivoli Storage Productivity Center to discover or collect the information intended for those columns and fields. Use discovery, probe, ping, or scan jobs to ensure that you collect information for all the columns and fields that appear in the Topology Viewer.
- In environments where you are monitoring a large number of storage entities, Tivoli Storage Productivity Center might not be able to display all of those entities within the Topology Viewer. See "Viewing large environments in the Topology Viewer" on page 379 for more information on how to ensure that all your monitored entities are displayed. It is recommended that you perform this task in larger environments.

The topology viewer

This topic describes how the topology viewer provides a graphical representation of the physical and logical resources that have been discovered in your storage environment by IBM Tivoli Storage Productivity Center.

The information displayed by the topology viewer is collected by discovery, probing, performance data collection, and alerting. The topology viewer consists of two views, a graphical view and a table view, organized vertically with the graphical view on top of the table view. The table view shows the same information as the graphical view, but in a tabular format.

The topology viewer supports four levels of detail, from a high-level overview of the entire storage environment to detailed information about individual entities. This allows you to view as much or as little detail as you need.

Overview

A global, highly aggregated view of the entire storage environment. This is the default view when you expand **IBM Tivoli Storage Productivity Center → Topology**.

Groups level (L0)

This level focuses on groups of a particular type. The default groups are:

- **Computers**: All discovered computer groups in your storage environment
- **Fabrics**: All fabric groups in your storage environment. A *fabric* is a network of entities that are attached through one or more switches. The topology viewer also displays Virtual SANs (VSANs) and Logical SANs (LSANs).
- **Switches**: All switches in your environment.
- **Storage**: All storage subsystems and tape libraries in your environment.
- **Other**: All entities that are discovered in your environment but that do not fit in either the computer, fabric or storage classes.

The groups level is the default level upon launching the topology viewer by clicking either **Computers**, **Fabrics**, **Switches**, **Storage** or **Other** under the **Topology** node. The groups level shows one or more groups of entities that correspond to the topology class that you selected. Information in the graphical view at this level can include any individually pinned entities.

Group level (L1)
This level focuses on one selected group and its related entities (for example, a group of computers). At L1, individual entities can be shown by expanding the group box. You can also display lines representing connections between entities or between entities and groups.

Detail level (L2)
This level focuses on individual entities (for example, a single computer, switch, subsystem, or tape library) and the paths to associated logical and physical entities. This level shows details all the way to the port level.

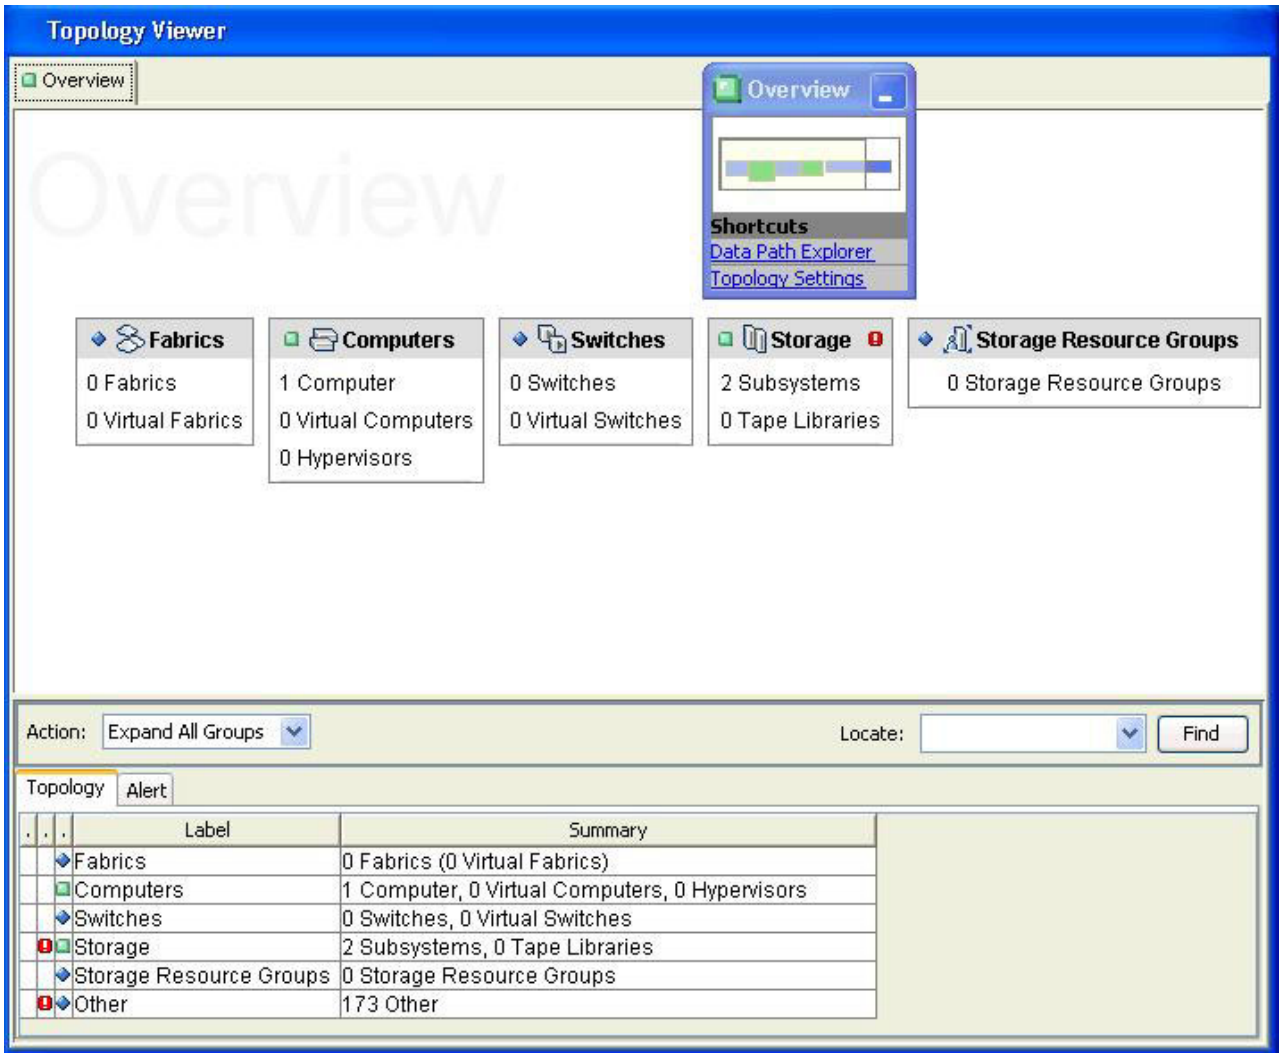

Figure 32. Topology overview

In addition, the topology viewer depicts the relationships among resources (for example, the disks comprising a particular storage subsystem). Detailed, tabular information (for example, attributes of a disk) is also provided. With all the

information that the topology viewer provides, you can more quickly and easily monitor and troubleshoot your storage environment and perform critical storage-management tasks.

Topology viewer interface controls

Progressive information disclosure reduces the visual complexity of a system. In the topology viewer, this means that you can identify and focus on entities in your storage environment that are of particular interest and hide other entities from view.

The tabs displayed in the tabular view represent classes of entities (a physical device or logical resource discovered byIBM Tivoli Storage Productivity Center). Individual entities (for example, a computer named *accounting1*, a tape library named *tapelib4*) are listed in their corresponding tabs. Entity attributes (for example, the operating system of a computer) are organized into table columns. The classes entities are assigned to are Computers, Fabrics, Storage Subsystems, and Other.

Note: The Other class contains entities discovered by Tivoli Storage Productivity Center but for which the type cannot be determined.

Changes to the graphical view or the tabular view are automatically applied to the other view, keeping both views synchronized. For example, if an entity is selected in the graphical view, it is also highlighted in the table view.

As more of your resources are discovered and displayed in the topology viewer, it is not possible to see all the discovered resources in a single topology window. To make it easier for you to find the resources you want to view, Tivoli Storage Productivity Center provides a tool called the *minimap*. The minimap allows you to shift the focus of the topology viewer to different areas within a particular view. The minimap is a small window that is initially located in the upper-right corner of any topology view.

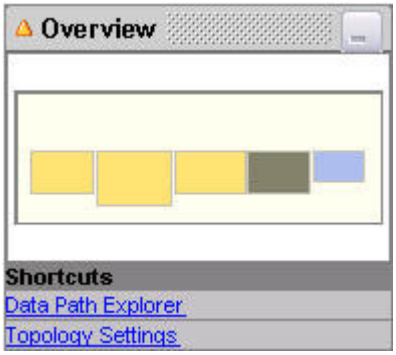

Figure 33. Minimap. You can click the shortcuts to launch Data Path Explorer or to view or change topology settings.

The minimap provides a miniature representation of the current topology view. Clicking on the minimap displays a multidirectional cursor that allows you to move the focus of the current topology view in any direction. As you drag the cursor around the minimap, the topology view will move accordingly. This makes navigating around complex topologies much easier.

The minimap uses the following colors to indicate the aggregated health status of groups:

- Green (normal): All entities in the group are operating normally.
- Yellow (warning): At least one entity is operating normally, and one or more entities have a health status of warning, critical, unknown, or missing.
- Red (critical): The status of all entities is either critical or warning.
- Gray (unknown): None of the entities in the group has an known health status.
- Black (missing): All of the entities of the group are missing.

In the Configuration History page, the minimap uses the following colors to indicate the aggregated change status of groups:

- Blue (changed): One or more entities in the group changed (an addition or deletion of an entity is considered a change).
- Gray (unchanged): All entities in the group are unchanged.

Included at the bottom of the minimap are shortcuts that you can click to launch the Data Path Explorer or to view or change the topology settings:

- To launch the Data Path Explorer, select one or more hosts, subsystems, disks, volumes, or Mdisks and click the Data Path Explorer shortcut.

  Note: The shortcut to Data Path Explorer is not available in the Configuration History topology view.
- To view or change the topology settings, click the Topology Settings shortcut. The Settings window displays for you to specify the type of overlays to display, the type of filtering to perform, whether connections between topology objects should display, and whether the Zone tab should display in the table view.

The graphical view:

This topic describes how the topology viewer groups entities that share certain characteristics. A collapsed group consisting of many entities or subgroups takes up much less space in the graphical view than the individual entities or subgroups would.

In the topology viewer, you can expand or collapse a group or subgroup by clicking the plus (**+**) character or minus (**-**) character, respectively, located in the upper-right corner of a group box.

Note: Pinned entities within a group remain visible, even when the group is collapsed.

By default, the topology viewer groups entities by class (computers, fabrics, switches, storage subsystems, and other entities). However, you can create additional groups (but not additional classes) using your own criteria. These custom groups can provide new perspectives on and important information about your storage environment. For example, you can create a special group belonging to the computer class that consists of computers that use the same operating system or storage subsystems of a particular model.

The following figure shows an expanded computer group representation in the graphical view:

Figure 34. Computer group

The title bar of the group box shows, from left to right, the following:

- An icon indicating the health status of the group (for example, normal).
- An icon indicating the types of entities contained in the group.
- The class of group (for example, computers).
- The grouping criterion (for example, Windows).
- The total number of entities in the group. If filtering is turned on, the number in the brackets is the ratio of displayed entities to total entities. For example, 5/10 indicates that 5 entities are shown out of 10 total entities.

The topology viewer provides the following default groups. These groups are accessible, if they exist in your topology, in the overview or by expanding the primary default groups (**Computers**, **Fabrics**, **Switches**, **Storage** and **Other**) that are displayed when you first launch the topology viewer.

Table 20. Group locations in the topology viewer

Type of group	Description	Path to group in topology viewer
Computers	All the computers that are discovered in your environment or a group of computers associated with a particular fabric or switch	**Computers**
Fabrics	All the fabrics and virtual fabrics that are discovered in your environment	**Fabrics**
Storage	All the storage-related devices that are discovered in your environment	**Storage**
Other	All the entities that are discovered in an environment but that do not fit in either the computer, fabric, or storage categories	**Other**
Fabric	All entities (for example, switches) that belong to or are related to a fabric	**Fabrics** → **Fabric**
Virtual Fabric	All entities that belong to a virtual fabric	**Fabrics** → **Fabric**
Switches	A set of switches	**Fabrics** → **Fabric**
Subsystems	A set of subsystems	**Storage** → **Subsystems**
Tape libraries	A set of tape libraries	**Storage** → **Tape Libraries**
Device	Devices associated with a particular entity (for example, a computer)	**Computers** → **Computers**

Table 20. Group locations in the topology viewer (continued)

Type of group	Description	Path to group in topology viewer
Volumes	A set of volumes that are associated with an entity	Different types of volumes groups exist at the L2 level of each of the four main categories: **Computers**, **Fabrics**, **Storage**, **Other**
HBA	A set of HBAs that are associated with an entity	**Computers** → **Computers** → **Computer**
Disks	A set of disks that are associated with an entity	**Storage** → **Subsystems** → **Subsystem**
Pools	A set of pools that are associated with an entity	**Storage** → **Subsystems** → **Subsystem**
FC Ports	A set of FC ports that belong to an entity	**Fabrics** → **Fabric** → **Switch**
Media	All media attached to an entity	**Storage** → **Tape Libraries** → **Tape Library**
Media Changers	All media changers attached to an entity	**Storage** → **Tape Libraries** → **Tape Library**
Tape Drives	All tape drives attached to an entity	**Storage** → **Tape Libraries** → **Tape Library**
Tape Drive	All ports on a tape drive attached to an entity	**Storage** → **Tape Libraries** → **Tape Library**
Switch	All ports on a switch. Due to the large number of ports on some switches, ports are shown as FC port groups within the switch	**Fabrics** → **Fabric** → **Switch**

As shown by the tabs at the top of the graphical panel (**L0:Computers**, **L1:Computers**, and **L2:Computers**), other views have been opened but are hidden by the overview. The topology tab in the table view displays information about the entities in the groups. The minimap, which you can use to shift the focus in the view, is positioned in the upper-right corner of the graphical view.

Right-click the object and select **Open Detail View** to view entities in a more detailed view. For example, at a high level an object representing a storage subsystem is depicted as an icon with a label. At lower levels, the same storage subsystem can be rendered as a box containing sub-entities representing disks, with the corresponding table view showing attributes of each individual disk. At yet another detail level, the ports of the storage subsystem might be shown as well as status information for each port.

GUI Topology Viewer has new L0 and L1 Switches views:

The IBM Tivoli Storage Productivity Center GUI has a Switches node under the Topology node. The Switches node opens the new L0 Switches view, from which you can also access the new L1 Switches view.

Overview
The Overview panel now shows a box with the number of switches present in your environment.

L0 Switches
The L0 Switches view shows all switches in your environment.

If you right-click on a Switches group, then click Group By, you will get a dialog box for group settings. You can choose to group by a single group or you can group by multiple categories:

- ByHealth
- ByUDP1
- ByUDP2
- ByUDP3
- ByFabric
- ByGroupSize

Click OK to close the dialog box and update the group settings.

L1 Switches
The L1 Switches view shows information on a group of switches in your environment that you selected to display. This also includes connectivity information: what entities this group is connected to.

L2 Switch
This view shows one switch, with all detailed port and connectivity information. You can access this view by double-clicking a single switch icon in either the L0 or L1 Switches view, or through the other Topology Viewer views.

The table view:

This topic describes how tabs in the table view provide detailed information about the entities (for example, computers, subsystems and so on) that are relevant in the current view.

The following is a list of all the tabs available in the topology viewer. Not all tabs are visible in all views.

Topology
Detailed information about all the entities in the topology grouped by class of the entity (computers, fabric, storage, and other).

Computer
Detailed information about any computer in the computer group or groups in the current topology view.

Fabric Detailed information about a fabric groups and virtual fabric groups in the current topology view.

Switch
Detailed information about switch groups and switch entities in the current topology view.

Subsystem
Detailed information about volume groups and volume entities in the current topology view.

Disk Detailed information about disk groups and disk entities in the current topology view.

Pool Detailed information about pool groups and pool entities in the current topology view.

Note: One of the columns in the tabular view refers to whether or not a storage pool is primordial. A primordial pool represents unallocated storage capacity on a storage device. Storage capacity can be allocated from

primordial pools to create storage pools. This means that primordial pools are disk/device sources for allocation of storage pools.

Volume
Detailed information about volume groups and volume entities in the current topology view.

Tape Library
Detailed information about tape-library groups and tape-library entities in the current topology view.

Media Changer
Detailed information about media-changer groups and media-changer entities in the current topology view.

Tape Drive
Detailed information about tape-drive groups and tape-drive entities in the current topology view.

Other/Unknown
Detailed information about objects in your environment other than computers, switches, subsystems, and tape libraries in the current topology view.

FC Port
Detailed information about port groups and port entities in the current topology view.

Zone Detailed information about any of the following groups or entities in a zone. By default, the **Zone** tab is disabled.

- Fabric groups
- Zone-set groups
- Zone groups
- Computer entities
- Switch entities
- Subsystem entities
- Tape library entities

Alert Detailed information about all alerts that are associated with entities in the current view of the topology.

Identifying storage resources using icons:

This topic describes how the topology viewer displays entities and groups of entities with an appropriate icon based on the entity class.

The storage resource icons can help you quickly and easily identify storage resources. For example, an entity of class "computer" will be displayed in the graphical view with an icon that resembles a computer.

Most icons exist in large and small versions, as shown in the following table:

Topology Object	Icon	Group Icon
Computer		

Topology Object	Icon	Group Icon
Computer HBA		
Fabric		
Switch		
Switch port		
Storage subsystem		
Storage subsystem port		
Storage virtualizer (for example, SAN Volume Controller, HDS Tagmastore)		
Tape library		
Tape library media changer		
Database volume		
File system volume		
Other volume		
Storage disk		
Storage pool		
Unknown		

Note: *Unknown* entities are entities that are discovered by IBM Tivoli Storage Productivity Center but that do not have an agent installed or the agent is not working.

The topology viewer and virtual SANs:

This topic describes how Cisco switches support virtual SANs (VSANs), which is the logical partitioning of a fabric into multiple fabrics. The overall network is referred to as the physical infrastructure, and the logical fabrics are the VSANs.

SNMP queries to a Cisco switch collect information about the Cisco physical infrastructure and about any VSANs in that infrastructure. In-band fabric agents can also collect information about the VSANs but are unable to collect information about the physical infrastructure. The status of Cisco physical infrastructures shown in the topology view is determined based on the status propagated from the VSANs in the infrastructure and the status of the infrastructure itself. If the status of the physical infrastructure is missing but the status of all or some of the VSANs in the infrastructure is normal, infrastructure information can no longer be collected. This typically indicates that the out-of-band agent is no longer working properly or that SNMP queries to the Cisco switches are no longer configured because the switch has been removed from the list of agents in the **Out of Band Fabric** node.

The consequences of this situation are illustrated in the following scenario: A fabric includes a Cisco switch with a VSAN on it. An out-of-band agent on the switch collects physical infrastructure and VSAN information, and an in-band Fabric agent collects information about the VSAN. If the out-of-band agent is removed or Fabric Manager is no longer able to communicate with the agent, IBM Tivoli Storage Productivity Center still detects the VSAN through the in-band agents, but physical infrastructure information can no longer be collected. As result, in the fabric topology view, the status of the physical infrastructure box is shown as missing, while the status of one or more VSANs in the infrastructure shows as normal. To correct this situation ensure that the out-of-band agent is configured for the Cisco switch.

The topology viewer and VMware:

This topic describes what the topology viewer displays for VMware support.

The Overview view of the topology viewer displays a count of hypervisors and virtual machines in the Computers summary.

The L0: Computers view displays the hypervisors and virtual machines in addition to the normal computers and includes a new grouping - By Virtualizer.

The L2: Computer view for a hypervisor includes a mapping of virtual machines and shows the virtual machine disk to physical disk relationship.

The L2: Computer view for a virtual machine will display the physical disk (what the virtual machine thinks is a physical disk) to virtual machine disk relationship and hypervisor connectivity.

Note: IBM Tivoli Storage Productivity Center supports the mapping of storage between the storage subsystem volume and ESX Server's corresponding disk. This is supported for the ESX Server 3.5 or VirtualCenter 2.5. Previous versions of ESX Server or VirtualCenter are not supported for the mapping of storage to the storage subsystem.

The topology viewer and switches:

Learn about important considerations when viewing and collecting information about McDATA and Brocade switches in the topology viewer.

McDATA Switches, CIMOMS, and their status as displayed through the topology viewer

McDATA CIMOMs communicate with all switches in a fabric as long as they are all managed by the Enterprise Fabric Connectivity Management (EFCM) Server and the CIMOM is pointed at the EFCM Server. Alternatively, if a McDATA CIMOM is configured in "Direct Mode", it communicates with all switches in a fabric as long as the CIMOM has been configured to manage all switches.

When a switch is disabled, the CIMOM can still communicate with the switch so it does not go missing. McDATA leaves the operational status and health indications the same for the switch and then changes the operational status for all its ports to dormant.

Therefore, only the port-to-port connections appear in the topology viewer as missing, as well as any nodes connected to those ports. The ports might have a yellow status indicator due to their dormant status. This occurs whether the fabric is set up in a flat or a loop/mesh topology.

Brocade Switches, CIMOMS, and their status as displayed through the topology viewer

Brocade CIMOMs do not communicate directly with all switches in a fabric even when those switches are all managed by the CIMOMs. Brocade fabrics use a "proxy switch" mechanism for communcation. This proxy switch is the only switch in the fabric to which the CIMOM communicates.

When the proxy switch is disabled, the CIMOM can still communicate with the switch so it does not go missing. Brocade leaves the operational status and health indications the same for the switch and all its ports, but then changes the **Enabled State** for all its ports to **disabled**. Therefore, the switch, its ports, and all its port-to-port connections appear in the topology viewer with the same **Operational Status** as they had before the switch was disabled.

However, because the proxy switch is disabled, the rest of the switches in the fabric go missing, as well as any nodes connected to switches in the fabric. This occurs whether the fabric has its switches in a flat or a loop/mesh topology.

When a "non-proxy switch" is disabled, the CIMOM can no longer communicate with that switch using the proxy switch. Therefore, the switch, its ports, and all its port-to-pot connections appear in the topology viewer as missing. For a Brocade fabric with a loop/mesh topology only, the disabled switch is affected because the proxy switch can reach the other switches through other ISL connections.

To workaround this issue, it is recommended you set up redundant agents (such as Out of Band Fabric Agents) for the Brocade switches in your fabrics.

Viewing large environments in the Topology Viewer

In environments where you are monitoring a large number of storage entities, IBM Tivoli Storage Productivity Center might not be able to display all of those entities within the Topology Viewer. To avoid this issue, there are some tasks you can perform to improve overall scalability and performance of the user interface.

If the Topology Viewer is unable to display all the entities in your environment, the user interface generates javacore and heapdump files in its installation directory. It is recommended that you check this directory to determine if this display problem is occurring in your environment. To do this, perform the following steps:

1. View the contents of the installation directory where Tivoli Storage Productivity Center's user interface is installed. The default installation directories are as follows:
   - Windows: C:\Program Files\IBM\TPC\gui\
   - UNIX/Linux: /<usr or opt>/IBM/TPC/gui/
2. Determine if javacore and heapdump files appear in the installation directory. You can identify these files by searching for files that begin with the word **javacore** or **heapdump** and have a **.txt** or **.phd** extension. For example: javacore.20080605.151906.8088.txt and heapdump.20080605.151903.8088.phd.

If these files appear in the installation directory, it might indicate that the display problem is occurring when you use the Topology Viewer.

There are two methods for working around this issue:

- Reducing the number of open topology views and disable unneeded overlays (such as the alert, performance, and zoning overlays). The number of topology views that are open and overlays that are enabled can affect how many storage entities are able to be displayed in the Topology Viewer.
- Increasing the memory allocation for IBM Tivoli Storage Productivity Center's user interface component. If reducing the number of open topology views and disabling unneeded overlays does not correct the problem, you can increase the amount of memory that is allocated to the user interface to help resolve the issue.

Reducing the number of open topology views and disabling unneeded overlays:

Perform these tasks to help improve the scalability of the Topology Viewer.

The number of topology views that are open can affect how many storage entities are able to be displayed within those views. The more views (tabs) that are open in the Topology Viewer, the more memory is used by the user interface. If the user interface requires more memory than is available to display all the entities within the open views of the Topology Viewer, some of the entities are not displayed. To workaround this problem, close any views that you no longer need to view and refresh the content of the views that you leave open.

The number of overlays that are enabled in the Topology Viewer can also affect how many storage entities are able to be displayed. The topology viewer uses color and icons to indicate the status of objects in the topology and tabular views. Overlays are colors and icons that used to indicate health, performance, change, and alerts. To disable overlays that you do not need, complete the following steps:

1. Right-click on any blank space within a topology view.
2. Select **Global Settings...** from the pop-up menu.
3. Uncheck the overlays you do not need to view. For example, clear the **Performance** and **Alert** check boxes.
4. Click **OK** to accept the changes. The content of the open views within the Topology Viewer are automatically refreshed according to your changes.

Increasing the memory allocation for the user interface:

Perform this task to allocate more system memory to the IBM Tivoli Storage Productivity Center user interface.

Increasing the amount of allocated memory might enable the user interface to correctly display a large number of storage entities in the Topology Viewer. This task is generally recommended and encouraged in larger environments.

To do this, perform the following steps:

1. Close any current sessions of the Tivoli Storage Productivity Center user interface.
2. Locate the batch or script file that is used to start the user interface. This file is located in the Tivoli Storage Productivity Center installation directory on the computer where the user interface is installed. The default file names and installation directories are as follows:
   - Windows: C:\Program Files\IBM\TPC\gui\tpc.bat
   - UNIX/Linux: /<usr or opt>/IBM/TPC/gui/TPCD.sh
3. Open the tpc.bat or TPCD.sh file in a text editor and locate both occurrences of the parameter that defines how much memory is allocated to the GUI component:
   - Windows parameter: **-mx256M**. The default setting for Windows is 256 megabytes.
   - UNIX/Linux parameter: **-xmx512M**. The default setting for UNIX and Linux is 512 megabytes.
4. Increase the value for all occurrences of this parameter. For example:
   - Windows: change **-mx256M** to **-mx1024M**
   - UNIX/LINUX: change **-xmx512M** to **-xmx1024M**

   Note: Make sure that the computer on which the user interface is installed has enough memory to support the setting that you define.
5. Save the batch or script file and restart the user interface.
6. Check the Topology Viewer to ensure that all the monitored storage entities are displayed as a result of increasing the amount of memory allocated to the user interface. If you still experience this display problem, repeat steps 1-5 and further increase the value for the parameter in the startup file.

Contact an IBM service representative if changing the memory allocation for the user interface does not help resolve the issue.

Arranging entities in a group

This topic describes how to arrange entities in an expanded group.

You can arrange entities within the group box by health status or by entity label if an expanded group contains many entities. With either type of arrangement, you can specify ascending or descending order. To arrange entities, follow this procedure:

1. Right-click the graphical view and select **Arrange By**.
2. Select one of the following:
   - **Label (ascending)**
   - **Label (descending)**
   - **Health (ascending)**

- **Health (descending)**

Checking the status of entities and groups

This topic describes how the topology viewer uses icons to indicate the health, performance, and operational status for entities.

To quickly check the health, performance, operational, change, and alert status of entities or groups of entities, look for the following icons associated with a particular entity or group in the topology. Position the cursor over these icons to see the health string in text format:

Table 21. Topology status indicators and their icons

Indicator	Icon
Health - Normal	
Health - Warning	
Health - Critical	
Health - Unknown	
Health - Missing	
Operational - Running	
Operational - Upgrading	
Operational - Missing	
Performance - Normal	
Performance - Unknown	
Performance - Warning	
Performance - Critical	
Performance - Missing	
Change - Objects did not exist at the two specified time periods	Objects in light gray with no icon
Change - Objects changed between the two specified time periods	
Change - Objects did not change between the two specified time periods	
Change - Objects were created between the two specified time periods	
Change - Objects were deleted between the two specified time periods	

Table 21. Topology status indicators and their icons (continued)

Indicator	Icon
Alert - Critical, Warning, or Informational	

The operational status icons (running, upgrading, missing) are displayed by default and cannot be hidden. However, you can specify whether to display health, performance, or alert icons. To display these icons, follow this procedure:

1. Right-click a blank space in the background of the graphical panel and select **Global Settings**.
2. In the Settings window, select **Health**, **Performance**, or **Alert**. Click **OK**.

Health status

This topic describes how the topology viewer displays the health status of entities.

The health of an entity falls into one of the following categories:

- Normal: The entity is operating normally.
- Warning: At least one part of the entity is not operating or has serious problems.
- Critical: The entity is either not operating or has serious problems.
- Unknown: The entity was discovered but is not recognized by IBM Tivoli Storage Productivity Center.
- Missing: The entity was previously discovered by Tivoli Storage Productivity Center, but was not detected in the most recent discovery.

The health of a group falls into one of the following four categories:

- Normal: All the entities in the group are operating normally.
- Warning: At least one entity in the group has normal status.
- Critical: All entities are either in critical or warning states.
- Unknown: None of the entities in the group have a known health status.

Health indicators for some entities, such as ports, are only normal and critical. In certain situations, the health of an entity might not be available. This happens, for example, when the entity is identified as missing. In the graphical view, missing entities are shown with a semitransparent light gray layer on top of an existing icon.

Level-indicator buttons at the top of the graphical view (for example, **L0:Computers**) display icons for health.

Multiple states from single entities are aggregated into group states or states of entities shown in upper layers, using the following rules:

- If all entities are in the same state, the aggregated state will represent this state.
- You will get an aggregated critical state (red) if:
  - One entity is in the missing state and at least one other entity shows a critical state.
  - One entity is in the critical state and at least one other entity shows a missing state.
- You will get an aggregated normal state if:
  - One entity is in the normal state and all others are in normal or undefined state.

  - One entity is in the undefined state and all others are in normal or undefined state.
- All other combinations will aggregate to warning state (yellow).
- States displayed within the tabs on top of the topology view will be aggregated according to the focus of the specific entity class.
  - For Computers and Storage, only the states of these entities are used for aggregation.
  - For Fabric, only the switch status is used for aggregation.
  - For level 1 and 0, the states of all displayed entities are used for aggregation.

At the Overview level, the health status indicated by the icon on **Overview** at the top of the graphical view indicates the aggregated status of all the groups inside the graphical view (**Computers**, **Fabrics**, **Switches**, **Storage** and **Other**). At the L0 level, the health status indicated by the icon on the **L0** tab at the top of the graphical view indicates the aggregated status of all the groups inside the that particular L0 graphical view (for example, the health status for the L0:Computers view would aggregate the health status of all discovered computers, including those that are missing, computers whose health status is critical, computers whose health status is warning). At the L1 level, the health status indicated by the icon on the **L1** tab at the top of the graphical view indicates the status of the particular group that is the focus of the view. (The status of the group is based on the aggregated status of the entities within the group.) The health status for the group is also indicated in the table view. At the L2 level, the health status indicated by the icon on the **L2** tab at the top of the graphical view is the same as the status of the particular entity that is the focus of the view. The health status for the entity is also indicated in the table view.

Note:

- At the L1:Fabric level the health status indicated by the icon on the **L1:Fabric** tab at the top of the graphical view might not be the same as the overall health of the group, based upon the aggregated status of all the entities within the group.
- At the L2:Subsystem level, the status of the volume group is the aggregated status of all the volumes that comprise the group and not the status of the host that owns those volumes.

Operational status

This topic describes how the topology viewer has three categories of operational status for entities.

The operational status of most entities falls into one of the following three categories:

Table 22. Categories of operational status for storage entities

Status	Description
Running	The entity is running.
Upgrading	The entity is upgrading.
Missing	The entity was not discovered during the last refresh cycle.

However, some entities display additional operational status. For example, switches can display a status of **OK**, **Error**, **Stopped**, and other types of status.

Note: When IBM Tivoli Storage Productivity Center no longer detects an entity that was previously displayed in the Topology viewer, the operational status of

that entity remains as the last known status before it went missing. This means that the operational status of an entity is not cleared or reset if that entity is not found in subsequent data collection jobs.

Performance status

This topic describes performance status details collected for entities.

Performance information is collected for entities in the L2:Switch and L2:Subsystem views. You can see the performance indicators for individual switches and subsystems in these two views by turning on the performance feature in the Global Settings window. This feature is not enabled by default.

The performance status of an entity falls into one of the following five categories:

- Normal: The entity is performing at the expected level.
- Warning: At least one part of the entity is not operating at the expected performance level.
- Critical: The entity is operating below the expected performance level.
- Unknown: The performance of the entity is not known.
- Missing: The entity was not found during the last refresh cycle and the performance status is not available.

The performance status does not get carried up to the L0 and L1 views. It is only available in the L2:Switch and L2:Subsystem views.

Performance status of an entity is not propagated upwards to its parent entity. For example, a SAN Volume Controller can cross several critical or warning stress thresholds (and show as yellow or red in the topology viewer) even though the underlying pools, disks or volumes do not have a degraded performance status.

After performance has been enabled, you will see a performance status indicator for any entity in these views that has performance data available. Position the cursor over a performance icon to see a list of performance metrics for that entity.

To enable the collection of performance information, perform the following steps:

1. Right-click a blank space in the background of the graphical panel and select **Global Settings**.
2. In the Settings window, select **Performance** and click **OK**.

Change status

This topic describes how the Configuration History topology view displays the status of entities over time.

The Configuration History feature uses the icons and colors of the change overlay to indicate historical changes in a SAN configuration's entities. The same icons and colors are used in the graphical view, table view, and snapshot selection panel. The following table lists the five possible states of an entity and shows the associated icons and colors.

State of an entity	Change overlay icon	Change overlay color indicator
Entity changed between the time that the snapshot was taken and the time that a later snapshot was taken.	Yellow pencil	Blue background

State of an entity	Change overlay icon	Change overlay color indicator
Entity did not change between the time that the snapshot was taken and the time that a later snapshot was taken.	Light gray circle	Dark gray background
Entity was created or added between the time that the snapshot was taken and the time that a later snapshot was taken.	Green cross	Green background
Entity was deleted or removed between the time that the snapshot was taken and the time that a later snapshot was taken.	Red minus sign	Red background
Entity did not exist at the time that the snapshot was taken or at the time that a later snapshot was taken.	Not applicable	Light gray background

Level-indicator buttons at the top of the graphical view (for example, **L0:Computers**) display the change icons.

For groups and other collections, the change overlay provides only aggregated status. That is, the change status for a group of entities is either "changed" or "unchanged." Details for the group as a whole are not provided. Aggregated status displays in the minimap. In the Configuration History page, the minimap uses the following colors to indicate the aggregated change status of groups:

- Blue (changed): One or more entities in the group changed (an addition or deletion of an entity is considered a change).
- Gray (unchanged): All entities in the group are unchanged.

An object that is moved from one group to another group appears only once in the new group in the graphical view. For example, if the health status of a computer has changed from Normal to Warning, the Configuration History page displays the computer as changed in the Warning health group (and no longer displays the computer in the Normal health group).

Important: If you manually deleted a missing entity from the Tivoli Storage Productivity Center , the configuration history function rediscovers and displays the entity in the configuration history viewer as both Missing and in a Normal state. For example, Entity A (with ID 12) is missing and you decide to manually delete it by using the Remove From Database option from the pop-up menu for Entity A in the topology viewer or configuration history viewer. Later, the Tivoli Storage Productivity Center rediscovers Entity A and assigns it a new ID (ID 22). The configuration history function treats the same Entity A (with ID 12 and ID 22) as two different entities and displays them in both Missing and Normal states in the configuration history viewer. This scenario can also occur if a missing entity is removed by the Tivoli Storage Productivity Center removed resource retention function and is rediscovered by theTivoli Storage Productivity Center again at a later time.

To display a summary of change for the entire row, hover over the changed row. The summary displays the title of the column. It also displays the old and new value for each cell or tab that changed. For example, if a group contains one new entity but everything else in the group is unchanged, the group as a whole displays as "changed." If you click the group, the one new entity displays as "created" and all other entities display as "unchanged."

The Change column in the table view also indicates changed entities.

Alert status

This topic describes how the topology viewer displays alert status for all entities in the current view.

The alert status of an entity falls into one of the following categories:

- Warning: At least one part of the entity is not operating or has serious problems.
- Critical: The entity is either not operating or has serious problems.
- Informational: Related data about the entity is supplied to aid you in decision making.

When enabled, in the topology viewer the alert overlay is located at the right of entity icons, below the pinning indicator. The icon for the alert overlay is a red exclamation mark  for the presence of alerts. In the title bar of groups and tabs, the alert icon is similarly located at the far right of the entity label. The alert overlay is disabled in the Configuration History page.

To find alerts that are associated with an entity in the table view, click the alert overlay indicator to go to the corresponding alert below the Alert tab. Or, right-click the entity to display a pop-up menu and click **Show Alerts in Tabular View** to go to the Alert tab and any corresponding alerts that are associated with that entity. The corresponding alerts are highlighted in purple. When you select an alert in the Alert tab, the entities that are associated with that alert are highlighted in purple in the graphical view. Note that more than one entity may become highlighted. An alert icon beside a collapsed group indicates that alerts exist within the collapsed group.

The alert overlay provides rollup, similar to pinning. If alerts are inside a collapsed group, the number of alerts within that group are aggregated and indicated in the collapsed group. All alerts are aggregated into one count, independent of type or entity involved. Hover over the alert indicator for a tab or group to determine the total number of alerts that are rolled up into that indicator and the time stamp of the newest alert.

In the table view, click the Alert tab to view all alerts associated with entities in the current view of the topology. The Alert tab contains alerts as a hierarchical list. For configuration analysis alerts, the top level provides information about which policy was violated during a particular analyzer run. The next level indicates the number of times that the policy was violated. This list contains one node for each policy violation. The Alert tab only lists the policy violations that contain at least one of the affected entities that are displayed in the graphical view. The remaining policy violations are filtered. The only exception is the Overview panel where the Alert tab lists all of the policy violations. Click each violation to view a list of affected entities. The parent entities for these affected entities are displayed in a separate column. You can resort the alert list by clicking the headers in the table. For instance, by clicking on the header above the entity the alert table is reorganized into a hierarchical list sorted by entity, then by alert type. If no alert was generated for an entity, the Alert column is blank.

To find an entity associated with an alert in the graphical view, click the alert below the Alert tab. The corresponding entity and any associated entities are highlighted (but not selected) in the graphical view. The highlighting is purple.

For more information, see “Viewing alerts” on page 259.

Creating custom groups

This topic describes how custom groups can be created based on criteria that is useful and relevant to you.

Custom groups are visible in both the graphical view and the table view. You can create custom groups using any of the following criteria. The grouping criterion is displayed in parentheses in the title bar of the group box.

For example, using the topology viewer, you can organize a collection of computers into groups based on the type of operating system (for example, Windows, Linux and so on). Using a UDP, you could also create subgroups for each operating system-based group of computers (for example, the Windows group could contain subgroups for Windows 2000 and Windows NT). Organization into groups can also be based on physical location of associated devices; for example, devices in Lab A could be organized in one group, devices in Lab B could be organized in another group, and so on. To create these types of groups, use UDPs (for example, a text string of ″Lab A″). To set UDPs use the pop-up menu item, Launch Detail Panel.

If you specify grouping by UDP, the groups are sorted in alphanumeric order, from left to right, in the graphical view. If the grouping is by health status, groups are sorted, from left to right, in the following order: missing, critical, warning, normal and unknown. You can also specify the arrangement of entities within a group box based on health status and label.

To group entities, follow this procedure:

1. Right-click a group and select **Group By**.
2. Select a group-by option. Options differ depending upon the class of the group.

   Single Group
   All entities are included in one group.

   Health
   Groups are based on health status (for example, normal or warning).

   Model The model of the storage subsystems server (for example, DS8000). This option is available for subsystem groups only.

   OS Type
   Groups are based on operating systems. This option is available for computer groups only.

   RAID level
   Groups are based on RAID level. This option is available only for pool groups at the **L2:Subsystems** level.

   Fabric This option is available for switch groups only.

   User-defined properties (UDPs)
   UDPs are characteristics that you define for an entity. For example, you could create a UDP for several entities using the text string ″new installation.″ All the entities with that UDP would be grouped together. The topology viewer provides three UDPs so that you can group resources in three different ways, depending on your needs. UDPs are not available for fabric groups.

Defining properties for grouping entities

This topic describes how User-defined properties (UDPs) are useful for describing specific properties of an entity or a set of entities.

You can also use UDPs to group entities with similar characteristics. The topology viewer lets you define three UDPs for an entity. UDPs are not available for fabric groups. To define one or more UDPs for an entity, follow this procedure:

1. Right-click an entity and select **Launch Detail Panel**.
2. Type a property in the **User Defined Property 1** field (for example, a specific operating system, location of computers, and so on).
3. Optional: Type additional properties in the **User Defined Property 2** and **User Defined Property 3** fields, and then click **Set**.
4. Type a description of the element manager and click **OK**.
5. Click **Launch**.

Displaying data paths from one endpoint in the topology to another

This topic describes how to view connectivity for selected entities in the topology using Data Path Explorer.

Data Path Explorer Overview

The Data Path Explorer allows you to follow connectivity between a host disk or computer and a storage subsystem. The Data Path Explorer also provides SAN Volume Controller support. Using Data Path Explorer you can follow the connectivity from a computer to SAN Volume Controller, and from an SAN Volume Controller to a storage subsystem.

Viewing a single data path allows you to monitor performance status and pinpoint weaknesses without navigating the many entities typically seen in Topology Viewer.

Data Path Explorer shows only those entities that are part of the selected data path, and only certain types of entities, such as host, switch, host bus adapter (HBA), port, storage subsystem, and storage volume.

Data Path Explorer functions like the rest of the Topology Viewer, with key variations.

Note: The shortcut to Data Path Explorer is not available in the Configuration History topology view.

Before Using Data Path Explorer

Different amounts of information can be shown by the Data Path Explorer, depending on which agents are deployed throughout the fabric.

To use Data Path Explorer for debugging performance problems, follow these steps to ensure that the information displayed in the Data Path Explorer view is current:

1. Enable Performance overlays. To see the overlays, go to topology **Global Settings > Active Overlays**: and select the **Performance** checkbox.
2. Ensure I/O is running on the host machine.
3. Create a threshold alert (see "Creating a performance threshold within an alert" on page 98).

4. Create and run performance monitors on the switches and storage subsystems of interest. See “Monitoring switch performance” on page 241 or “Monitoring storage-subsystem performance” on page 242 for detailed descriptions on creating and running performance monitors.
5. Run a discovery job for the CIM agents and Data agents.
6. Run a probe job for the subsystem, fabric, and Data agents.
7. Run a performance monitoring job for the subsystem and switches if you want to see performance information.

See the section about planning for the Data Path Explorer view in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for more information on configuring data path explorer.

Graphical View

When you open a data path view, three tiles display in the graphical view: Initiator Entities, Fabrics, and Target Entities. These contain all the entities in all the data paths related to the selected host or storage entity.

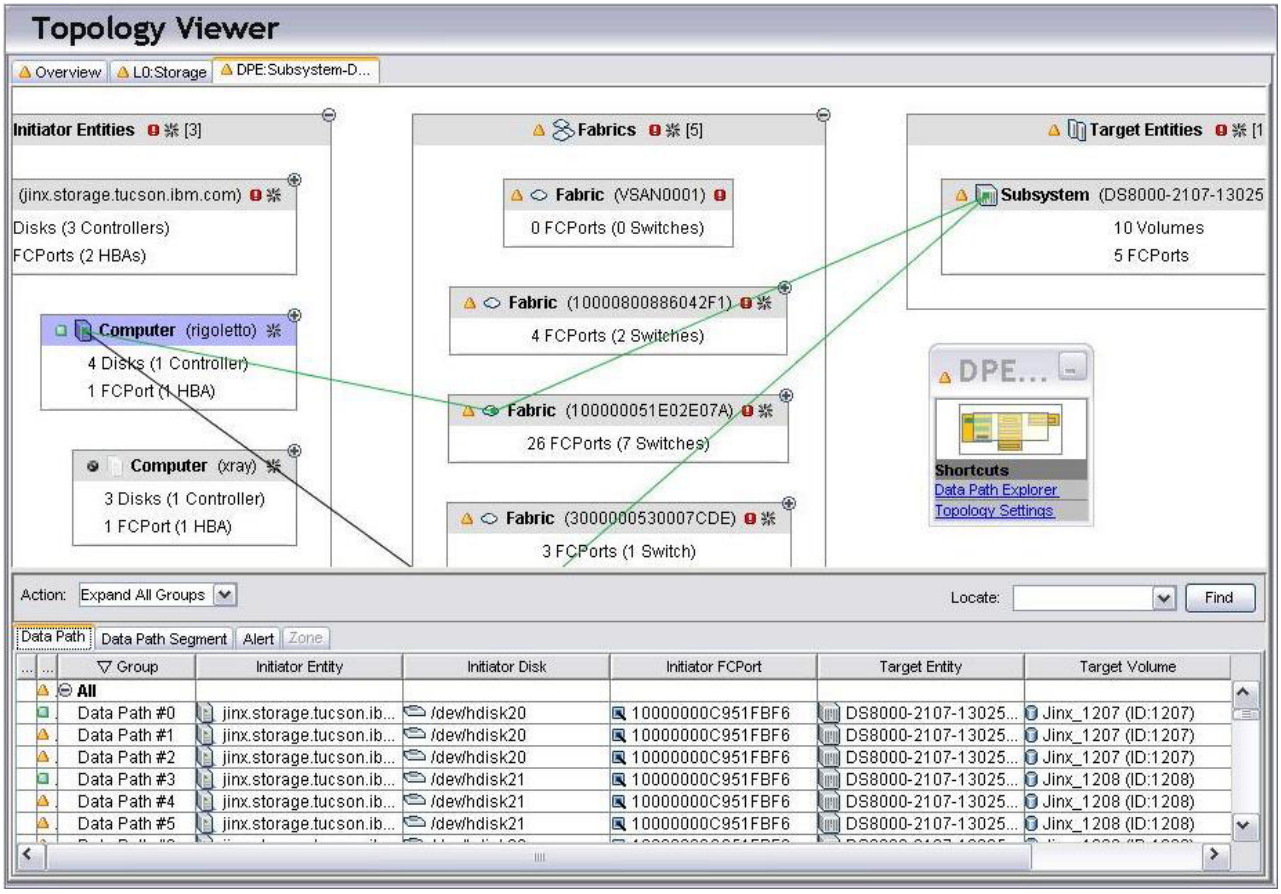

Figure 35. Data Path Explorer view

- Selecting an entity in any of the tile groups at any level displays connector lines showing the data paths the entity is involved in.

- Each tile behaves like an active topology viewer tab, allows grouping and regrouping, expanding and collapsing of groups and so forth. Each tile can be expanded to view the lower level details such as the disks, ports and volumes involved in a data path.

Initiator Entities

The Initiator Entities tile contains the computers and related volumes, disks, and HBA ports that have valid data paths to the storage system. Expanded computer tiles display each computer, its health status icons, and a count of the entities within the computer that are part of the valid data paths. Expanded computer tiles show the controllers and disks, HBA ports, and volumes on the computer that are part of valid data paths. The initiator entities tile can also contain SAN Volume Controllers. An SAN Volume Controller is a storage virtualization device and can act as an initiator entity. Expanded SAN Volume Controller tiles show the SAN Volume Controller managed disks and ports that are part of valid data paths.

Fabrics

The Fabrics tile contains the fabrics entities, mainly switches and ports, through which valid data paths for the selected endpoint traverse. A count of the switches and switch ports with health status icons and labels are shown in the group title. Expanded Fabric tiles display a view of the switches to the port level. Only switches and ports with valid data paths traversing through them are displayed.

Target Entities

The Target Entities tile contains the endpoint storage systems in data paths for the selected endpoint. Expanded Target tiles display each storage subsystem, its status icons, a count of entities, and the number of subsystem front-end ports and volumes involved. When these are expanded, the associated port and volume tiles display. These can also be expanded to show individual entities.

Table View

The data path table displays below the graphical view. Each row in the table view shows a valid data path for the selected endpoint.

Display a single data path in the graphical view by selecting a row in the data path table.

The data path table has two tabs: Data Path and Data Path Segment.

Data Path tab

This is the default view. This tab shows each data path associated with the entity for which the data path was opened.

- Select an entity in the graphical view to display individual data paths in the table.
- The health status shown for each data path is aggregated.
- The data path table can be sorted by any of the column headers.

See “Columns in Data Path tab” on page 393.

Data Path Segment tab

This tab displays the segments in a data path. One segment is listed per row.

There are two ways to populate the Data Path Segment table:

- Select a data path in the Data Path tab. The Segment tab then lists the segments in the selected data path.
- Click an entity in the graphical view. The Segment tab lists the segments in all of the data paths the entity is involved in. These are identified by data path number (0, 1, 4, and so on).

See "Columns in Data Path Segment tab" on page 394.

Related reference

"Opening a Data Path View"
Data Path Explorer provides an end-to-end topology view of the storage environment from a host disk to a storage volume. It displays the different data paths (SAN access paths or I/O paths) for each host disk.

"Displaying an entity's data paths" on page 393
You can view the data paths associated with any entity by selecting the entity in the graphical view. Lines connect the related entities in the data path or paths.

"Example use cases for Data Path Explorer" on page 395
The following examples show some of the use scenarios for Data Path Explorer:

Opening a Data Path View

Data Path Explorer provides an end-to-end topology view of the storage environment from a host disk to a storage volume. It displays the different data paths (SAN access paths or I/O paths) for each host disk.

You can open a data path view on any of the following endpoints:

host The host (computer) must have at least one disk that is assigned from a subsystem.

subsystem
The subsystem must have at least one volume assigned to a host.

volume
The volume must be assigned to a computer, and the fabric must be zoned so that the computer can see the volume.

disk The disk must be from a subsystem. Note that the disk can be host disk or SAN Volume Controller managed disk.

Open a data path explorer view in one of the following ways:

- Right click on an endpoint, and then select Open Data path view.
- Select an endpoint, and then select Open Data path view from the Action menu.
- Select an endpoint, and then select the Data Path Explorer shortcut in the minimap
- For multiple entities: Select one or more endpoints of the same type (for example, one or more storage volumes), and then select Open Data Path View. The data paths display for all of the endpoints selected. If you select an SAN Volume Controller entity, all other entities you select must be SAN Volume Controller entities.

Note:

1. Opening the Data Path Explorer view for a subsystem (including an SAN Volume Controller) that has several associated data paths can consume extensive system resources. Select volumes or disks instead, and then open the Data Path Explorer view.
2. If the endpoint is an SAN Volume Controller, two data path views are opened because the SAN Volume Controller can act as both a initiator and target entity.

The first view displays the data paths associated with the SAN Volume Controller acting as a storage subsystem or a target entity. The second view displays the data paths associated with the SAN Volume Controller acting as an initiator entity.

When a data path view is opened, three tiles display: the Initiator Entities, Fabric, and Target Entities tiles. These tiles contain all the associated entities in the data paths. Select an entity within one of the tiles to see the data paths associated with it. You may need to further expand the tiles.

If you select an entity or a group of entities that do not have a valid data path, a dialog box displays with the following message: ″No available DataPath for the selected entities.″

The Data Path Table lists each data path (one per row).

To return to a normal topology view, right-click an entity in the graphical or table view and select **Open Detail View** from the pop-up menu.

Displaying an entity's data paths

You can view the data paths associated with any entity by selecting the entity in the graphical view. Lines connect the related entities in the data path or paths.

- In the table view-Data Path tab, the selected entity's data paths are highlighted.
- The table view-Data Path Segment tab displays the data paths belonging to the selected entity and the individual segments that make up the data path.

Multiple paths

IBM Tivoli Storage Productivity Center shows the multiple paths from the disks to the storage subsystem volumes in the Data Path Explorer view. IBM Subsystem Device Driver (SDD) software allows the operating system to use multiple paths (that is, different initiator and target port pairs) through the SAN for storage access. On the L2 Computer View and the Data Path Explorer view, the disks with multiple paths managed by SDD are grouped in a controller representing the SDD driver.

For a Windows computer with multiple paths managed by SDD to subsystem volumes, a controller representing the SDD driver is displayed, and this controller contains the disks managed by the SDD.

For a non-Windows computer with multiple paths managed by SDD to subsystem volumes, Data Path Explorer displays a controller representing the SDD driver, and this controller contains the vpath disks. A separate disk (for example, hdisk on AIX) is created by the operating system on the non-Windows computer for each path to the storage volume. These disks are grouped by the controller that contains the initiator port present in the path.

Related reference

"Columns in Data Path tab"
When an entity in the Data Path Explorer graphical view is selected, its data paths display in the Data Path tab.

"Columns in Data Path Segment tab" on page 394
The Data Path Explorer Segment tab shows each segment in a data path.

Columns in Data Path tab

When an entity in the Data Path Explorer graphical view is selected, its data paths display in the Data Path tab.

The Data Path tab columns display the following information:

Located
This column is associated with the Locate function in the Topology Viewer. If search text is present in a row, the Locate column displays a magnifier icon.

Health (status)
The aggregated status of the data path.

Group The number of the data path.

Initiator Entity
The name of the computer or SAN Volume Controller in the data path.

Initiator Disk
The computer disk or SAN Volume Controller managed disk in the data path.

Initiator FCPort
The computer Fibre Channel port in the data path. Data in this column comes from the HBA port entity in the Computer tile. If an SAN Volume Controller is the initiator entity, this column contains the SAN Volume Controller fibre channel port in the data path.

Target Entity
The storage subsystem in the data path.

Target Volume
The storage volume in the data path. An asterisk indicates that this is a space-efficient volume or vdisk.

Target FCPort
The storage Fibre Channel port in the data path.

Columns in Data Path Segment tab

The Data Path Explorer Segment tab shows each segment in a data path.

Each segment has two endpoints: the beginning device and the ending device connected by a communication link. Each endpoint resides in a parent entity, which can be a computer, fabric, or subsystem.

The Data Path Segment tab columns display the following information:

Located
This column is associated with the Locate function in the Topology Viewer. If search text is present in a row, the Locate column displays a magnifier icon.

Segment Health
The health status of the segment is an aggregation of From Endpoint Health, To Endpoint Health, From Entity Health, To Entity Health, and Connection Health.

Label The number of the data path and its segments. The data path number corresponds to the numbering on the Data Path tab. Segment numbering advances from left to right, from computer tile through fabric tile to subsystem tile.

From Entity Health
The health status of the parent entity beginning the segment.

From Entity Performance
The performance status of the parent entity that begins the segment. This column is displayed only when the performance overlay is enabled.

From Entity
The name of the computer, fabric, or subsystem that is the parent of the beginning device (endpoint) in the segment.

From Endpoint Health
The health status of the endpoint device that begins the segment.

From Endpoint Performance
The performance status of the endpoint device that begins the segment. This column is displayed only when the performance overlay is enabled.

From Endpoint
The name of the endpoint device that begins the segment.

To Entity Health
The health status of the parent entity that ends a segment.

To Entity Performance
The performance status of the parent entity that ends the segment. This column is displayed only when the performance overlay is enabled.

To Entity
The name of the computer, fabric, or subsystem that is the parent of the ending device (endpoint) in the segment.

To Endpoint Health
The health status of the endpoint device that ends the segment.

To Endpoint Performance
The performance status of the endpoint device that ends the segment. This column is displayed only when the performance overlay is enabled.

To Endpoint
The name of the endpoint device that ends the segment.

Connection Health
The health status of the connection between the endpoint device that begins the segment (From Endpoint) and the endpoint device that ends the segment (To Endpoint). The Connection Health is displayed as N/A if the health status data for the connection is not available.

Example use cases for Data Path Explorer

The following examples show some of the use scenarios for Data Path Explorer:

- An application on a host is running slowly, which causes you to question the health of the I/O path for that application. Are there component level performance problems that might cause the slow application response?
- Disk utilization on the storage subsystem is high, triggering performance alerts. You decide to see which hosts or applications, residing on separate host disks, are causing the problem.
- You suspect that the I/O paths for two applications (on separate host disks) are in conflict with each other. They could be sharing a common component (such as a switch). After viewing the I/O paths for these two applications, you make the required zoning or connectivity change to alleviate the problem.
- I/O path information can be used do load balancing and storage planning. For example, a host adapter port on a host has I/O from three host disks passing

through it, but a second port has I/O from ten host disks (and has more I/O). This information can be used to provision a new host disk using the first host adapter port.

Note: Application to host disk mapping is not shown in the Topology Viewer. This mapping must be done outside IBM Tivoli Storage Productivity Center. However, Tivoli Storage Productivity Center can provide mapping from database applications to host disks if configured to do so.

Planning for Data Path Explorer view

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

The Data Path Explorer is a new type of view in the topology viewer. Data path explorer combines several of the usual 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 allows you to use the Data Path Explorer to check the health and performance statistics of the storage components along the I/O path.

To use Data Path Explorer for debugging performance problems, follow these steps to ensure that the information displayed in the Data Path Explorer view is current:

1. Enable Performance overlays. To see the overlays, go to topology **Global Settings > Active Overlays**: and select the **Performance** checkbox.
2. Ensure I/O is running on the host machine.
3. Create a threshold alert (see "Creating a performance threshold within an alert" on page 98).
4. Create and run performance monitors on the switches and storage subsystems of interest. See "Monitoring switch performance" on page 241 or "Monitoring storage-subsystem performance" on page 242 for detailed descriptions on creating and running performance monitors.
5. Run a discovery job for the CIM agents and Data agents.
6. Run a probe job for the subsystem, fabric, and Data agents.
7. Run a performance monitoring job for the subsystem and switches if you want to see performance information.

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

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

- CIM agent to monitor the storage subsystem
- Data agent to monitor the host
- In-band Fabric agent to monitor the fabric. Also a CIM agent for the fabric to collect fabric performance information.

To show the host disk assignment from a given storage subsystem, IBM Tivoli Storage Productivity Center Data Server needs access to the subsystem's CIM

agents (SMI-S agent). In the case of the subsystem, the data path for the subsystem is defined by the host masking or mapping information maintained by the CIM agent.

Analysis on data paths involving SAN Volume Controllers requires one or more CIM agents pointing to the SAN Volume Controller and the back-end system.

To show accurate information, it is important to have the most current information from the CIM agents, Fabric agents, and Data agents. You must probe all the components within the data path (or a set of agents strategically placed throughout the fabric) before using the Data Path Explorer.

To show physical fabric connectivity, you must have an in-band Fabric agent, out-of-band Fabric agent, or a CIM agent for the fabric.

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 **Topology** → **Global Settings** → **ActiveOverlays** and select the **Performance** check box. See "Collecting performance data" on page 241 for detailed descriptions on creating and running performance monitors.

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

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

These are the general setup requirements for using the Data Path Explorer:

1. Install and configure the SMI-S agents (CIM agents) for the subsystems (including SAN Volume Controllers) and fabric to collect subsystem and fabric data. Contact your vendor for information on installing and configuring the CIM agent for the subsystem or fabric.
2. Install and configure the in-band Fabric agent, or out-of-band Fabric agent to monitor a given fabric.
3. Install the Data agent on the hosts.

Displaying membership in zones and zone sets

This topic describes how the topology viewer allows you to examine zones and zone sets.

Zone information is shown in the **Zone** tab in the table view. The **Zone** tab is disabled by default. Selecting a zone or zone set in the **Zone** tab highlights zone and zone set members in the graphical view. Ports and entities that are members of the selected zone or zone set are indicated in the graphical view by highlighting. You can select highlighted entities to show details (for example, connections).

The entities listed in the **Zone** tab are only those entities that are in the current view. However, zone or zone-set selections apply to all currently open views in the topology viewer. Therefore, switching to another view might indicate other entities in the selected zone.

Note: You cannot use the **Zone** tab to change membership status of entities.

To activate the **Zone** tab, follow this procedure:

1. Right-click in the graphical view and select **Global Settings**.
2. In the Settings window, select **Show zone tab** and click **OK**.

Displaying detailed information about entities

This topic describes how to view detailed information about entities using the topology viewer.

There are two ways to display information about entities:

1. Right-click the entity and select **Launch Detail Panel** to display detailed information about an entity (for example, the operating system of a computer, the used and unused capacity of a storage subsystem)
2. Right-click the entity or entities and select **Reports**. From the dialog displayed, you can choose which report to create. The reports you can choose depend on which entities you selected. Asset reports are available only if you selected a single entity. Select a report and click **Create Report**

Displaying storage resource connections

This topic describes how physical and logical connections among entities in the environment can be displayed using the topology viewer.

Connections are shown in both the graphical and table views. In the graphical view, connections are represented by lines between entities or groups. In the table view, a **Connections** tab lists both ends of a connection.

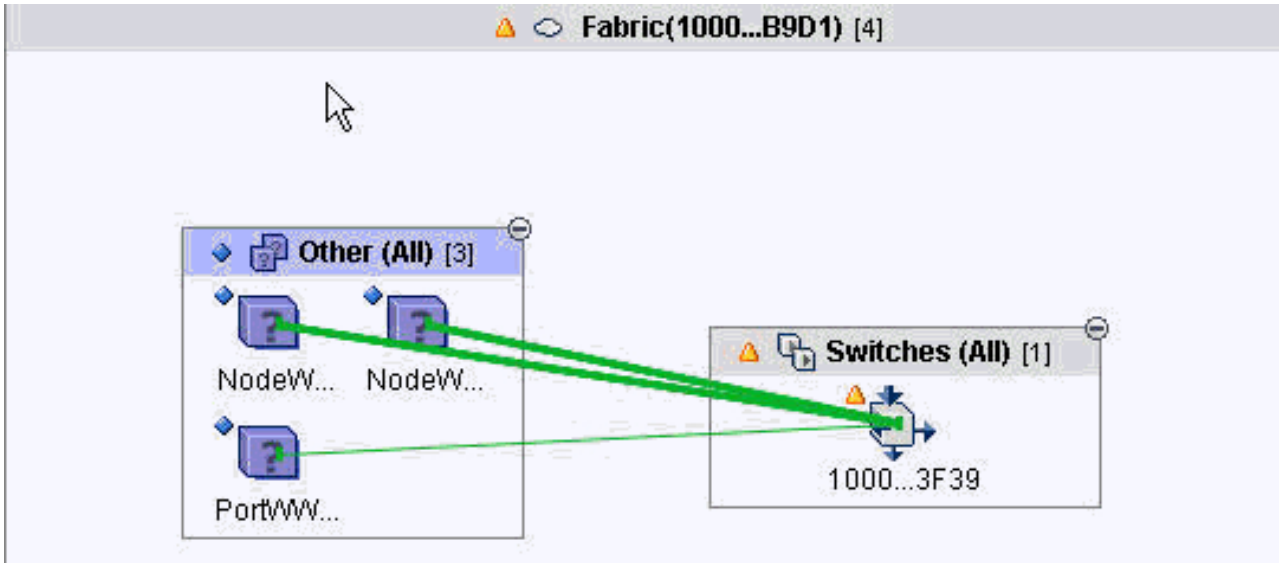

Figure 36. Connections

In the graphical view, the color of a connecting line between entities indicates the status of a connection. The color of the line between groups, individual entities, or between an individual entity and a group can be green (normal), yellow (warning), red (critical), black (missing), or grey (status unavailable).

Note: When you have a host with two HBAs or FC ports that are both in the same switch and you switch the ports between them, the connection will be shown with

two yellow (warning) lines. This is because you have not reset the history retention setting. Until the settings are reset, the connecting lines will remain yellow.

Connections between two individual entities are represented by a thin line. Connections between groups or between a group and an individual entity are represented by a thick line. A connection to a collapsed group can represent either a single connection or several connections.

To show connections, follow this procedure:

1. Right-click a blank space in background of the graphical view and select **Global Settings**.
2. In the Settings window, select **Show Connections** and click **OK**.

Filtering entities based on health status

This topic describes how to quickly narrow the health status of normal or non-normal entities in your environment.

In views with many entities and connections, the display can become very cluttered and can make it difficult to determine where problems are occurring. You can filter entities to display only those entities whose health status is non-normal (critical, warning, missing, or unknown).

When you apply filtering, you can also choose to display entities with a health status of normal but that are connected to entities whose health status is either critical or warning. To filter entities by health status:

1. Right-click in the background area of the graphical view and select **Global Settings**.
2. In the Settings window, select **Hide Normal Entities** and, optionally, **Unless connected to non-normal**, and click **OK**.

Labeling entities

This topic describes how you can specify or change label entities.

In both the graphical and table views, entities are identified by labels (for example, the name of a computer). You can specify or change these labels. To specify or change the existing label for an entity, follow this procedure:

1. Right-click the entity and select **Launch Detail Panel**.
2. Type a label for the entity in the **Label** field.
3. Click the save icon in the menu bar to save the label, and then close the window.

Starting element managers

This topic describes how to start element managers for the entities that are defined in IBM Tivoli Storage Productivity Center.

If element managers on a system, tape library, switch, or other entity have been defined, you can start the element managers to access vendor-provided management tools.

To start an element manager for an entity, right-click the entity. A menu that contains launch in context options is displayed. The options that are displayed depend on the element managers that you have configured on the Element Manager tab of the Configuration Utility. The possible options are:

- ***Launch descriptor file defined application.*** This is an individual menu item for each application that is defined by a launch descriptor file. For Tivoli Storage Productivity Center, launch descriptor files are provided for IBM System Storage DS8000 andIBM Tivoli Storage Productivity Center for Replication.

  If the launch descriptor file cannot be found, the **Launch file items error** menu item is displayed. Click this item to view a description of the problem.
- **External Tools.** This menu item shows any external tools that are defined in the Configuration Utility. Examples of external tools include Web-based user interfaces and local executable files. For example, you might add the address of a Web site that you frequently access or an application that you often use (a text editor, a spreadsheet application, and so on). External tools are not device specific. Any defined external tool is displayed in the launch in context menu for every device that is defined in Tivoli Storage Productivity Center.
- **Element Manager.** This menu item shows any element managers discovered for the device by a Storage Management Initiative Specification (SMI-S) query. To avoid menu item duplication, this item is not displayed when the associated element manager is defined by a launch descriptor file. In this situation, the element manager is displayed as an ***Launch descriptor file defined application*** menu item.
- **User Defined Tools.** This menu item shows any user-defined element managers that are defined in the Configuration Utility. These element managers differ from the element managers in the Element Manager menu because they are user-defined and are not discovered by a SMI-S query.

If the selected entity has no element manager, these options are not available in the menu.

Launching SAN Planner

This topic describes how to launch the SAN Planner for an entity from the topology viewer. SAN Planner assists end-to-end planning for SAN components.

Right-click a computer, fabric or storage subsystem and select **Launch Planner**. A new SAN Planner configuration page will be displayed, and the element you selected will be added to it in the selection pane.

Related concepts

"SAN Planner overview" on page 278
The SAN Planner assists the user in end-to-end planning involving fabrics, hosts, storage controllers, storage pools, volumes, paths, ports, zones, and zone sets . Once a plan is made, the user can select to have the plan implemented by the SAN Planner.

Starting Storage Optimizer

This topic describes how to start the Storage Optimizer from the topology viewer. The Storage Optimizer uses data in the Tivoli Storage Productivity Center database to analyze your storage subsystems to identify performance bottlenecks, and recommend changes to improve performance.

Right-click one or more storage subsystems and select **Launch Optimizer**. A new Storage Optimizer analysis definition is created, and the storage subsystems you selected are added to the **Select Storage** pane.

Related concepts

"Storage Optimizer overview" on page 291
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.

Locating storage resources in the topology

This topic describes how you can search for a particular entity using the entity name or a keyword (for example, a host name or operating system).

The open view (for example, L0:Other) will be searched for the string you enter. The results of a locate action are shown by highlighting.

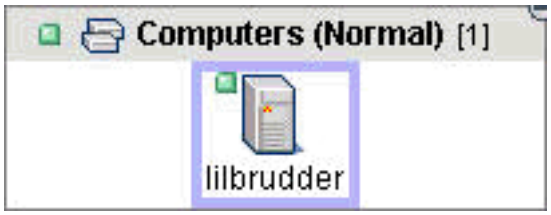

Figure 37. Highlighted entity

To locate an entity, type a search string in the **Locate** field in the right corner of the table view and click **Find**. The current view (for example, L0:Other) will be searched for the string you enter. The locate function does not search the entire topology, only the currently open view.

When the entity is found, it will be displayed in both the graphical and table views. If the located entity is inside a collapsed group in the topology, you will not be able to see that entity until you expand that group. However, the entity will be highlighted in the table view. If the entity is in a collapsed group in the table view, that group will automatically be expanded and the highlighted table entry will receive focus. You can then click on that table entry to see the located entity in the topology view.

If you specify a Locate criterion that applies to more than one entity in a view, all the applicable entities are displayed.

Managing Storage Subsystems

This topic describes how to display the Storage Subsystems page, where you can perform storage-subsystem management tasks.

Right-click a storage subsystem and select **Manage Storage Subsystems**.

Managing Zone Control

This topic describes how to display the Fabric page, where you can perform zone-control tasks.

Right-click a fabric and select **Manage Zone Control**.

Pinning entities

This topic describes how pinning entities is an easy way to identify selected entities for future reference or quick access.

Pinned entities display a small flag next to the entity. This flag is propagated to all views that contain that entity. This is useful for marking entities for various reasons, such as a reminder to look at this entity later and for upward propagation. For example, the user may pin a computer in an L0 or L1 Computers view and then quickly wish to see which fabric it belongs to by opening an L0 Fabrics view and looking at which fabric group contains the pinned entity - indicated by the small flag. Pinned objects remain visible as you change views in the topology viewer and when you log off and log back in. Pin lists are stored on a per user basis. This means that different users can have separate pin lists. Pinned objects are indicated by a pin icon in the **Pinned** column of a table or by a flag icon next to the entity in the graphical view.

Figure 38. Pinned entity

To pin an entity, right-click the entity in the graphical view and select **Pin**. To unpin an entity, right-click the entity in the graphical view and select **Unpin**. The same function is available in the **Action** list in the table view.

Pinned entities are visible in groups at all levels and in all views regardless whether the group is expanded or collapsed. You can only pin individual entities; you cannot pin groups (for example, a group of computers).

Printing information from the topology viewer

You can print information displayed in the topology viewer on a single page for reference.

Printing capability is available for the graphical view but is not available in the tabular view. To print the information displayed in the topology viewer, right-click the graphical view and select **Print** or **Print Preview**. The content is scaled to fit onto one page if necessary.

Refreshing views

This topic describes how you can refresh the topology viewer to view the most up-to-date details about your environment.

To provide you with the most up-to-date information, the topology viewer automatically refreshes all open views, including graphics and tables, every five minutes from the IBM Tivoli Storage Productivity Center database. You can also refresh views manually or change the default refresh rate. Opening a new view automatically shows the most current information. To manually refresh views or change the refresh rate, right-click in the graphical view or click the arrow in the **Action** field in the table view and choose one of the following options:

- **Refresh View**: Updates the view that is currently displayed.
- **Refresh All Views**: Updates all views that are currently open. You can also use this option to remove the alert indicators from the devices in the topology view.
- **Refresh Settings**: Opens the **Refresh Rate Settings** window, in which you can select a new automatic refresh rate (specified in minutes).

Removing entities from the database

Computers in the missing state were once in the topology viewer, but were not discovered by IBM Tivoli Storage Productivity Center during the latest refresh cycle. You might want to remove them from the view and the database in order to make it easier to work with other resources.

To remove a computer from the topology viewer (both graphical and tabular views) and the database, you first need to uninstall the agents on that system.

Note: This action is available only for entities that are in the missing state.

1. See *Uninstalling the agents* in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for instructions on uninstalling the agents.
2. In the topology viewer *graphical* view, right-click the computer entity and select **Remove From Database**..
3. In the topology viewer *tabular* view, right-click the computer entity and select **Remove From Database**.

Removing deleted entities

When a discovered entity (fabric, subysystem, switch, or volume) is deleted using another application, such as an element manager or other management application, the next discovery or probe of the topology viewer displays a missing entity. The status of missing entities is aggregated and propagated up through the device hierarchy.

It can be difficult to determine if there are other lower-level entities that have ″missing″ status, but have not been intentionally removed by a management application. The Removed Resource Retention setting in the topology viewer can be set to a low value, such as 0 or 1 day to remove the missing entity from the display. This removes the missing entities from the display the next time the Removed Resource Retention process runs.

You can also remove deleted entities from the topology viewer using the following procedure:

1. Highlight the intentionally deleted entities in the topology view and select **Remove Resource** in the **Action** list in the tabular view.

   Note: Not all entity classes can be removed from the topology using the **Action** list. For those entities, use the Removed Resource Retention setting.
2. A prompt asks you if you are sure you want to remove the resource from the database. Click **Yes**.
3. To refresh the view, click **Refresh View** in the **Action** list. The missing entities will no longer be displayed in the topology viewer.

Selecting entities in graphical and table views

This topic describes how to select entities in the graphical and table views of the topology viewer.

To select an entity in the graphical view, click the entity image. In the table view, click the name of the entity. To make multiple selections, press CTRL and click.

Showing Alerts

This topic describes how to display alert logs for entities in the topology viewer.

You can display alert logs for:

- computers
- fabrics
- storage subsystems
- tape libraries
- switches

All alerts for the entity-type you select will be displayed, even if no alerts exist for the entity.

Right-click an entity and select **Show <entity type> Alerts**.

For example:

Right-click a computer and select **Show Computer Alerts**.

Working with tape libraries in the topology viewer

Perform storage management functions on your tape library using the topology viewer.

The topology viewer displays information about the discovered tape libraries and tape library groups. The topology viewer has two views:

- Topology view
- Tabular view

Each view shows information about the discovered storage devices in your enterprise. Changes to one view are automatically displayed in the other view, so that both views are always synchronized.

To see tape library information in the topology viewer, click the **Storage** node in the **IBM Tivoli Storage Productivity Center → Topology** node in the console. The level zero (L0) view of the storage topology view is displayed. The view includes the tape libraries and tape library groups that have been discovered.

Viewing tape libraries in the topology view

This topic describes how you can see information about your tape libraries and tape library groups in the topology view.

Clicking a tape library icon in the topology view highlights (selects) that tape library in both the topology view and the tabular view. Right-click a tape library to display a pop-up menu that lists actions that you can perform on the library, as well as other general actions that you can perform. The actions specific to tape libraries are:

Open Detail View
Opens the next level of detail for the selected topology object. For tape libraries, the view hierarchy is

- L0:Storage- overall view of the discovered storage subsystems and tape libraries
- L1:Tape Libraries- tape library groups
- L2:Tape Library- view of an individual tape library and its related devices and connections

Launch Detail Panel
Displays the detail panel for the tape library. The detail panel shows detailed information about the tape library, such as the name of the tape library vendor, the serial number, status, owner/administrator contact information, and properties that can be defined by the user. It also has **Launch** button that will launch the element manager for the tape library.

Launch Element Manager
Launches the element manager application for the selected tape library. The element manager is an application or Web site provided by the tape library vendor for performing management tasks for the selected tape library. For example, the element manager for and IBM 3584 tape library is the IBM Tivoli Storage UltraScalable Tape Library Specialist.

Pin Pins the selected tape library. When a topology object is pinned, it remains visible even after its respective group or view level has been collapsed. A small blue flag icon to the right of the object's icon indicates that the object is pinned. Once an object has been pinned, this action will change to **Unpin** in the pop-up menu for the object.

As you add more storage devices to your enterprise, your topology view can become very crowded and increasingly difficult to navigate. You can use the **minimap** in the upper right corner of the topology view to quickly move around the view and locate the objects you need to work with.

Viewing tape libraries in the tabular view

This topic describes how you can see information about your tape libraries and tape library groups in a tabular view.

The tabular view presents the same information as the topology view. Changes to the tabular view are immediately reflected in the topology view, so that both views are always synchronized.

At the top of the tabular view are two fields:

Action
Lists the actions that can be performed on the selected object in the table, as well as the more general actions, such as Expand All Groups, that can be performed. These actions are the same ones that appear in the context menu for a selected object.

Locate Allows you to find objects in the tabular view. This is very useful when your tabular view has many entries.

The columns in the tabular view change as you select different objects in the topology. For example, the columns available for objects in the L1:Tape Libraries view are different than the columns in the L2:Tape Library view. These columns show information about the tape libraries, such as the operational status, vendor name, and number of drives in the library.

Managing disk storage

This section describes managing storage subsystems, virtual disks, volumes, and creating and managing alerts.

Working with workload profiles

This topic describes creating, modifying, and deleting workload profiles.

Only storage subsystems that have been discovered and had performance data successfully collected can be used with the Volume Planner function of the SAN Planner. SAN Planner and Volume Planner are only available for Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 systems.

For information on storage subsystem discovery, see "Discover storage subsystem, tape library, and fabric information" on page 197. For information on the Volume Planner, see "Using the SAN Planner" on page 283

Creating a workload profile

This topic describes how to create a workload profile. Workload profiles contain information about Enterprise Storage Server (Tivoli Storage Enterprise Storage Server) volumes and are used by the Volume Planner.

1. In the left pane, expand **Disk Manager → Profile Management → Workload Profiles**. In the right pane, the currently-defined workload profiles are displayed.
2. Click **Create**; the Workload Profile wizard starts.
3. On the Welcome to the Workload Profile Wizard page, click **Next**.
4. On the Creation Method page, select the basis for the new workload profile. Click one of the following choices:

   Create profile from volume performance data
   Base the new profile on the volume-performance data that you provide.

   Create from an existing profile
   Base the new profile on an existing workload profile. In the **Profile Name** field, select the workload profile.

   Click **Next**. If you clicked **Create from existing profile**, go to step 6. Otherwise, the Name and Description page opens.
5. On the Name and Description page, provide information about the new workload profile:
   a. In the **Name** field, type the name of the workload profile.
   b. Optional: In the **Description** field, type a description of the workload profile.
   c. Click **Next**.
6. On the Select Volumes page, specify the volumes whose performance data will be used as the basis for the workload profile.
   a. In the **Available volumes** field, select the volumes, and then move them to the **Selected volumes** field.
   b. Click **Next**.
7. On the Peak Activity periods page, specify the time period for the performance data:
   a. To use all available performance data, click **Use all available performance data**, and go to step 8.

Note: Analyzing all performance data might be a lengthy process.

b. To use performance data for a specified time period, specify the time period. In the **Start date** field, select the month, day, and year. In the **End date** field, select the month, day, and year. By default, the time period is set to the previous month. In the **Time reference** field, select one of the following options:

   Client Time

   Server Time

   GMT

   Device Time

c. To specify that performance data is collected only during specified hours, click **Use only a portion of each day**. In the **Start time** field, select the hour and minutes. In the **End time** field, select the hour and minutes.

d. Click **Next**.

8. On the Review Profile Results page, review the following performance data:

Data quality

Lists a data-quality rating. This rating is based on the percentage of volumes that had valid data:

- Good (100% - 67%)
- Fair (66% - 33%
- Poor (32% - 0%)

I/O Demand

I/O rate per GB per sec
The data-transfer rate in gigabytes per seconds (GBps).

Avg. transfer size (KB)
Default size for data transfer in kilobytes (KB). You can edit this value.

Read / Write

Sequential reads (%)
Percentage of read/write activity that was sequential reads.

Sequential writes (%)
Percentage of read/write activity that was sequential writes.

Random reads (%)
Percentage of read/write activity that was random reads.

Random writes (%)
Percentage of read/write activity that was random writes.

Total (%, must be 100)
Total of the preceding percentages. This value must equal 100%. If it does not, you cannot continue with the wizard.

Cache

Random read cache hits (%)
Percentage of cache activity that was random hits to read cache.

Random write destage (%)
Percentage of cache activity that consisted of moving data from cache to nonvolatile storage.

9. Click **Finish**.

Deleting a workload profile

This topic describes how to delete a workload profile.

Only user-defined workload profiles can be deleted.

1. In the left pane, expand **Disk Manager** → **Profile Management** → **Workload Profiles**. A list of the currently defined workload profiles is displayed in the right pane.
2. Select the workload profile you want to delete.
3. Click **Delete**.

Modifying a workload profile

This topic describes how to modify a workload profile.

Only user-defined workload profiles can be modified.

1. In the left pane, expand **Disk Manager** → **Profile Management** → **Workload Profiles**. A list of the currently defined workload profiles is displayed in the right pane.
2. Select a workload profile and click **Details**. The Workload Profile Properties window displays.
3. Click the **General** tab to view the name and description of the workload profile. You can modify only the description of the workload profile.
4. Click the **Performance** tab to view I/O demand, read/write, and cache performance data for the workload profile. You can modify any of the performance data for the workload profile.

Viewing a workload profile

This topic describes how to view a workload profile.

1. In the left pane, expand **Disk Manager** → **Profile Management** → **Workload Profiles**.
2. In the right pane, select a workload profile and click **Details**. The Workload Profile Properties window displays.
3. Click the **General** tab to view the name and description of the workload profile.
4. Click the **Performance** tab to view I/O demand, read/write, and cache performance data for the workload profile.
5. Click the **Activity** tab to view peak activity settings for the workload profile. To view all data collected during a period of time, select **Use all available performance data** and specify a time reference (options are Client Time, Server Time, GMT, and Device Time) and start and end dates. To view only the data collected during a portion of each day, select **Use only a portion of each day** and specify start and end times. Click **OK** to view the peak activity.

Working with provisioning profiles

This topic describes creating, modifying, and deleting provisioning profiles.

Only storage subsystems that have been discovered and had performance data successfully collected can be used with the Volume Planner function of the SAN Planner. SAN Planner and Volume Planner are only available for Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 systems.

For information on storage subsystem discovery, see “Discover storage subsystem, tape library, and fabric information” on page 197. For information on the Volume Planner, see “Using the SAN Planner” on page 283

Creating a provisioning profile

This topic describes how to create a provisioning profile. Provisioning profiles are used by SAN Planner, and in general, by features that require a model for volume, multipathing and zoning requirements.

1. In the left pane, expand **Disk Manager → Profile Management → Provisioning Profiles**. In the right pane, the currently-defined provisioning profiles display.
2. Click **Create**. The Provisioning Profile wizard starts.
3. On the Welcome to the Provisioning Profile Wizard page, select one of the following:

   Create profile from scratch.
   Base the new profile on the volume size and redundancy, multipathing and zoning requirements data that you provide.

   Create from an existing profile.
   Base the new profile on an existing provisioning profile. In the **Profile Name** list, select the provisioning profile.

   Click **Next**. If you clicked **Create from an existing profile**, go to step 5. Otherwise, the Name and Description page opens.
4. On the Name and Description page, provide information about the new provisioning profile.
   a. In the **Name** field, type the name of the provisioning profile.
   b. Optional: In the **Description** field, type a description of the provisioning profile.
5. In the Volume Settings section, specify the volumes whose performance data will be used as the basis for the provisioning profile.
   a. Select **Specify how storage will be allocated and its performance characteristics** to allocate the storage.
   b. Enter the **Total Capacity** in gigabytes.
   c. Select **Use existing unassigned volumes, if available** to use existing volumes.
   d. Select one of the two options listed under**Divide Capacity**.
      - **By Number of Volumes.** Select the number of volumes that you want to create. By default, this is set to 1. For an Tivoli Storage Enterprise Storage Server subsystem, the maximum is 255.
      - **By Volumes of Size.** Select the size of the volumes to be created. The maximum and minimum size depends on the type of storage subsystem.
      - **RAID Level.** Select a RAID level based on the type of storage subsystem. The choices are:
        – DS6000 and DS8000: RAID 5, RAID 10. RAID 5 is the default setting.
        – Tivoli Storage Enterprise Storage Server: RAID 5 or RAID 10. RAID 5 is the default setting.
        – FAStT or DS4xxxx: RAID 0, RAID 1, RAID 3, or RAID 5. RAID 0 is the default setting.
   e. In the **Volume Name Prefix** field, type in this name. This string will be used as the prefix to all volumes created. This string cannot be longer than 15 characters.

f. The **Workload Profiler** list box is updated to include the newly created workload profile.

6. In the Path Settings section, specify the time period for the performance data.
   a. To set up the multipath options, click **Setup multipath options (if supported by the host drivers)**.
      1) **Multipath Option** specifies how the driver uses the paths between the host and the storage subsystem. The options are:
         - **Load Balancing** sends Input/Output on all paths.
         - **Round Robin** sends Input/Output on one path until a time interval expires (set in an SDD setting at the host) or stops to use another path.
         - **Fail-Over** sends Input/Output on one path until a failure occurs and fails over (switches) to another path.
      2) **Specify number of paths** specifies the number of paths between each host and the storage subsystem.
      3) When selected, **Use fully redundant paths (requires 2 fabrics)** causes the Path Planner to check for redundant fabrics between each host and storage subsystem and to create paths in each fabric. This requires at least two fabrics.
7. In the Security Settings section:
   a. Click **Automatically change zoning to ensure hosts can see the new storage**.
   b. Select **Automatically Create Zone** to indicate where zoning will be done. The predefined options are:

      <auto-zone>
      This is the default. The plan is generated creating the maximum number of zones without grouping the data paths based on host, host ports, or host bus adapters (HBAs).

      ...for each host
      Creates a zone for each host.

      ...for each HBA
      Creates a zone for each HBA.

      ...for each host port
      Creates a zone for each host port.

   c. In **Maximum Number of Zones** enter the maximum number of zones in the box.
   d. In **Maximum zone members per zone** enter the maximum number of zone members (per zone) in the box.
   e. If you want to set a prefix name for each zone, enter a zone name prefix in **Append zone name prefix.**
   f. If you do not want two controllers with different types to be used in the same zone, select **No two controllers with different types should be used in the same zone**.
   g. If you want any zone set that is available selected, select **Use active zone set**.
8. Click **Finish**.

Deleting a provisioning profile

This topic describes how to delete a provisioning profile.

Only user-defined provisioning profiles can be deleted.

1. In the left pane, expand **Disk Manager** → **Profile Management** → **Provisioning Profiles**. A list of the currently defined provisioning profiles is displayed in the right pane.
2. Select the provisioning profile you want to delete.
3. Click **Delete**.

Modifying a provisioning profile

This topic describes how to modify a provisioning profile.

Only user-defined provisioning profiles can be modified.

1. In the left pane, expand **Disk Manager** → **Profile Management** → **Provisioning Profiles**. A list of the currently defined provisioning profiles is displayed in the right pane.
2. Select a provisioning profile and click **Details**. The Provisioning Profile Properties window displays.
3. Click the **General** tab to view the name and description of the provisioning profile. You can modify only the description of the provisioning profile.
4. Click the **Performance** tab to view I/O demand, read/write, and cache performance data for the provisioning profile. You can modify any of the performance data for the provisioning profile.

Viewing a provisioning profile

This topic describes how to view a provisioning profile.

1. In the left pane, expand **Disk Manager** → **Profile Management** → **Provisioning Profiles**.
2. In the right pane, select a provisioning profile and click **Details**. The Provisioning Profile Properties window displays.
3. Click the **General** tab to view the name and description of the provisioning profile.
4. Click the **Performance** tab to view I/O demand, read/write, and cache performance data for the provisioning profile.
5. Click the **Activity** tab to view peak activity settings for the provisioning profile. Click **OK** to view the peak activity.

Managing storage subsystems

This topic contains information about storage-subsystem management tasks.

Grouping storage subsystems

A storage-subsystem group is a user-defined set of storage subsystems. You might want to group storage subsystems in order to make it easier to run monitoring or reporting jobs.

Note: A storage subsystem can belong to only one storage-subsystem monitoring group, but can belong to multiple storage subsystem reporting groups.

Creating a storage subsystem group:

This topic describes how to create a storage subsystem group.

1. In the left pane, expand **Disk Manager** → **Monitoring** → **Groups**.
2. Right-click **Storage Subsystem** and then click **Create Storage Subsystem Group**. In the right pane, the Create Storage Subsystem Group window opens.
3. In the **Description** field, type a short description of the group.
4. In the **Available** field, click the storage subsystems that you want to add to the group, and then click  . The storage subsystems are added to the **Current Selections** field.
5. Click **File** → **Save**. A window opens requesting a name for the storage-subsystem group.
6. Type the name of the storage-subsystem group and click **OK**.

Deleting a storage-subsystem group:

This topic describes how to delete a storage-subsystem group.

1. In the left pane, expand **Disk Manager** → **Monitoring** → **Groups** → **Storage Subsystems**. All existing storage-subsystem groups are listed.
2. Right-click the storage-subsystem group that you want to delete.
3. Click **Delete**. A window opens, asking if you are sure that you want to delete this group.
4. Click **Yes**.

Editing a storage-subsystem group:

This topic describes how to edit a storage-subsystem group.

1. In the left pane, expand **Disk Manager** → **Monitoring** → **Groups** → **Storage Subsystem**. All existing storage-subsystem groups are listed.
2. Click the storage-subsystem group that you want to edit. The **Edit Storage Subsystem Group** window opens.
3. Make any necessary changes.
4. Click **File** → **Save**.

Defining element managers

Element managers are programs that you can use to configure and manage 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.

Element managers are managed in Tivoli Storage Productivity Center using the Configuration Utility. You can access the Configuration Utility from a storage system that is visible in the Storage Subsystems page of Tivoli Storage Productivity Center as shown in the following steps:

1. Expand **Disk Manager** → **Storage Subsystems** in the navigation tree. The Storage Subsystems page is displayed.
2. In the **Storage Subsystems** table, double-click a storage subsystem. The Details dialog page is displayed.
3. Click **Element Management**. The **Element Manager** tab in the Configuration Utility is displayed.

Related concepts

"Managing DS8000 storage systems using element managers" on page 167
Use the Configuration Utility in IBM Tivoli Storage Productivity Center to add and manage element managers for IBM System Storage DS8000. *Element managers* are applications that are used to configure and manage a storage device. You can add and manage connections to the element manager that is provided with DS8000, DS8000 Storage Manager, or you can specify other element managers for a device.

"Managing other storage systems using element managers" on page 179
Use the Configuration Utility in IBM Tivoli Storage Productivity Center to add and manage element managers for storage devices other than IBM System Storage DS8000. *Element managers* are applications that are used to configure and manage a storage device. You manage connections to the element managers that are provided with the storage device or you can specify other element managers for a device.

Starting element managers

You can start element managers from a storage system that is visible on the Storage Subsystems page of Tivoli Storage Productivity Center. You can also start element managers from other points in Tivoli Storage Productivity Center as described in other areas of this documentation.

To start an element manager from the Storage Subsystem page, complete the following steps:

1. Expand **Disk Manager** → **Storage Subsystems** in the navigation tree. The Storage Subsystems page is displayed.
2. In the **Storage Subsystems** table, select a storage system.
3. Use one of the following methods to start the element manager:
   - To start the default element manager for the storage system, select **Launch Element Manager**.
   - To start another element manager for the storage system, right-click the system and click **Launch Element Manager**. Select the element manager that you want to start from the menu.

If the connection information for an element manager is valid, the element manager is started. If unsuccessful, an error message is displayed. The error messages that you receive describes one of the following problems:

- that the element manager at the specified IP address could not be started
- that the required connection information for the element manager has not been provided.
- that the element manager is not installed and configured correctly (for local element manager installations such as XIV Storage Manager)

If you are prompted to enter additional information such as an ID and password for the element manager, complete the following steps:

1. Click **OK** in the message dialog box. The Modify Element Manager dialog box is displayed.
2. Enter the required information for the element manager and click **Save**.

Removing a storage subsystem

This topic describes how to remove a storage subsystem.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystem** table, click a storage subsystem, and then click **Remove**. The Delete Storage Subsystem window opens.

3. Click **Yes**.

Setting user-defined properties

This topic describes how to set properties for storage subsystems. You can define a name and up to three properties for each storage subsystem.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystem** table, double-click a storage subsystem. The Storage-Subsystem Details page opens.

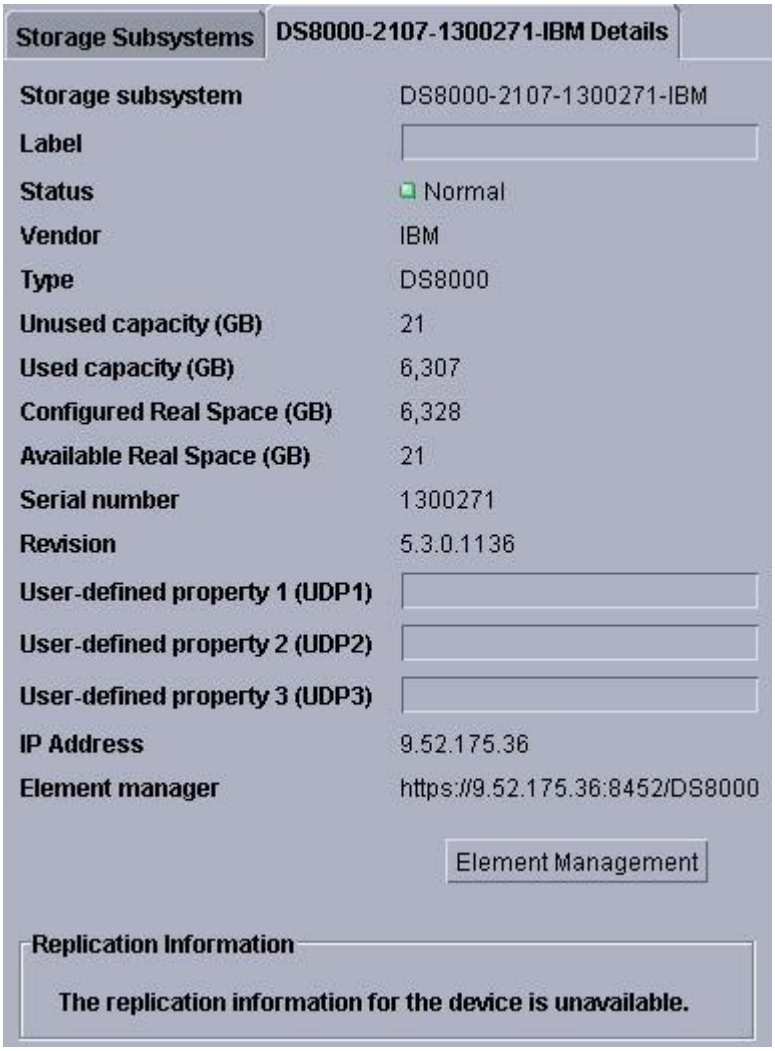

Figure 39. Storage-Subsystem Details page

3. Set the user-defined properties:
   a. In the **Label** field, type a name for the storage subsystem.
   b. In the **User-defined properties** fields, type information that you want to associate with the storage subsystem.
4. Click **File** → **Save**.

Viewing storage-subsystem information

This topic describes how to view information about the storage subsystems in your environment.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.

   The **Storage Subsystems** table contains the following information about each storage subsystem:

   Subsystem
   A system ID which is in the following format:

   Type-Model-ID-Vendor

   - *Type* is the type of the storage system, for example, Tivoli Storage Enterprise Storage Server or SAN Volume Controller.
   - *Model* is the model number, for example, 2105 or 2145.
   - *ID* is the serial number or cluster name. Cluster names only are used for SAN Volume Controllers.
   - *Vendor* is the manufacturer of the storage system, for example, IBM.

   Label A user-defined name or the system ID.

   Type The type of system:
   - DS6000
   - DS8000
   - Tivoli Storage Enterprise Storage Server
   - FAStT
   - Hitachi
   - IBM XIV Storage System
   - Other
   - SAN Volume Controller

   Status Status of the system:
   - Normal
   - Critical
   - Warning
   - Unreachable

   Available Space (GB)
   The available space on the computer.

   Consumed Space (GB)
   The amount of space already consumed on the computer.

   Configured Real Space (GB)
   Total amount of real available and formatted storage space in the storage pools that are associated with the storage system

   Available Real Space (GB)
   Amount of real storage space available to allocate to volumes in the storage pools that are associated with the storage system.

2. To view detailed information about a specific storage subsystem, click the storage subsystem, and then click  . The Storage-Subsystem Details page opens.

   The Storage-Subsystem Details page contains the following information:

   Storage subsystem
   Storage subsystem ID.

   Label User-defined name of the subsystem. You can edit this field.

Status Status of the subsystem:

- Normal
- Critical
- Warning
- Unreachable

Vendor
Manufacturer of the subsystem.

Type Type of subsystem:

- DS6000
- DS8000
- IBM XIV Storage System
- Tivoli Storage Enterprise Storage Server
- FAStT
- Hitachi
- SAN Volume Controller
- Other

Available Space
Total storage capacity that is not formatted in volumes or virtual disks and assigned to hosts.

Tip: For space-efficient volumes, the Available Space is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage space reported for a storage subsystem using space-efficient volumes. This also applies to other space calculations, such as the calculations for the Storage Subsystem's Consumable Volume Space and FlashCopy Target Volume Space.

Consumed Space
Total storage capacity that is formatted in volumes or virtual disks and assigned to hosts.

Configured Real Space (GB)
Total amount of real available and formatted storage space in the storage pools that are associated with the storage subsystem.

Available Real Space (GB)
Amount of real storage space available to allocate to volumes in the storage pools that are associated with the storage subsystem.

Serial number
Serial number.

Revision
Revision.

User-defined property 1 (UDP1)
A user-defined property. You can edit this field.

User-defined property 2 (UDP2)
A user-defined property. You can edit this field.

User-defined property 3 (UDP3)
A user-defined property. You can edit this field.

IP Address
IP address of the element manager.

Element manager
Identifies the element manager. This either can be a URL, the fully qualified name of the executable file that launches the element manager, or the short name (for IBM XIV Storage System).

Element Management
Opens the Configuration Utility, which enables you to modify the element manager information.

Managing virtual disks

This topic contains information about managing virtual disks.

Adding a managed disk to a managed-disk group

This topic describes how to add a managed disk to a managed-disk group.

Attention: If you add a managed disk that contains existing data to a managed disk group, you will lose the data that it contains. The *image mode* is the only mode that will preserve this data.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystems** table, click a SAN Volume Controller.
3. Click **Create Virtual Disk**. The Create Virtual Disk wizard starts and the ″Define the virtual disks″ page opens.
4. In the **Managed-disk group** field, select the managed-disk group to which you want to add a managed disk.
5. Click **Add Managed Disk**. The ″Add managed disk″ window opens.
6. Select the managed disk, and then click **OK**.

Assigning host ports

This topic describes how to assign host ports to a virtual disk. You also can make host-port assignments when you create a virtual disk.

Note: When you assign a host port to a virtual disk, all host ports defined on the SAN Volume Controller model will be assigned to the virtual disk. If you want to assign only one host port to a virtual disk, you must define the host with only one host port on the SAN Volume Controller.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems window is displayed.
2. In the **Storage Subsystems** table, click the SAN Volume Controller with which you want to work.
3. Click **Virtual Disks**. The Virtual Disks page is displayed.
4. Click either **Filter by MDisk group** or **Filter by host port**, and then select the specific filtering condition.
5. Click **Go**. A list of virtual disks associated with the selected SAN Volume Controller is displayed.
6. In the **Virtual Disks** table, select the virtual disk, and then click **Assign Host Ports**. The Assign Host Ports wizard starts.
7. On the ″Assign the virtual disk to host ports″ page, assign or unassign host ports:

a. Optional: In the **Available ports** field, select the ports that you want to assign to the virtual disks, and then click >>.
b. Optional: In the **Assigned ports** field, select any ports that you want to unassign from the virtual disks, and then click <<.
c. Click **Next**.

8. (This page only is displayed if Fabric Manager is installed, and you have Fabric administrator authority.) On the "Zone actions" page, perform one of the following zone actions:

   Create a new zone
   Click **Create new Zone**. In the Zone name field, type the name of the zone. The ports listed in the field below are added to the new zone.

   Update an existing zone
   Click **Update existing zone**, and then select a zone. The ports listed in the field below are added to the selected zone.

   Leave the zone configuration as it is
   Click **Do neither** to avoid performing any fabric-control tasks. The ports listed in the field below will not be able to communicate with the subsystem on this fabric.

9. Click **Next**. The "Review settings" page opens.
10. Review the settings that you made earlier in the wizard, and then click **Finish**.

Creating virtual disks

This topic describes how to create virtual disks.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystem** table, click a SAN Volume Controller, and then click **Create Virtual Disk**. The Create Virtual Disk wizard starts and the "Define the virtual disks" page opens.

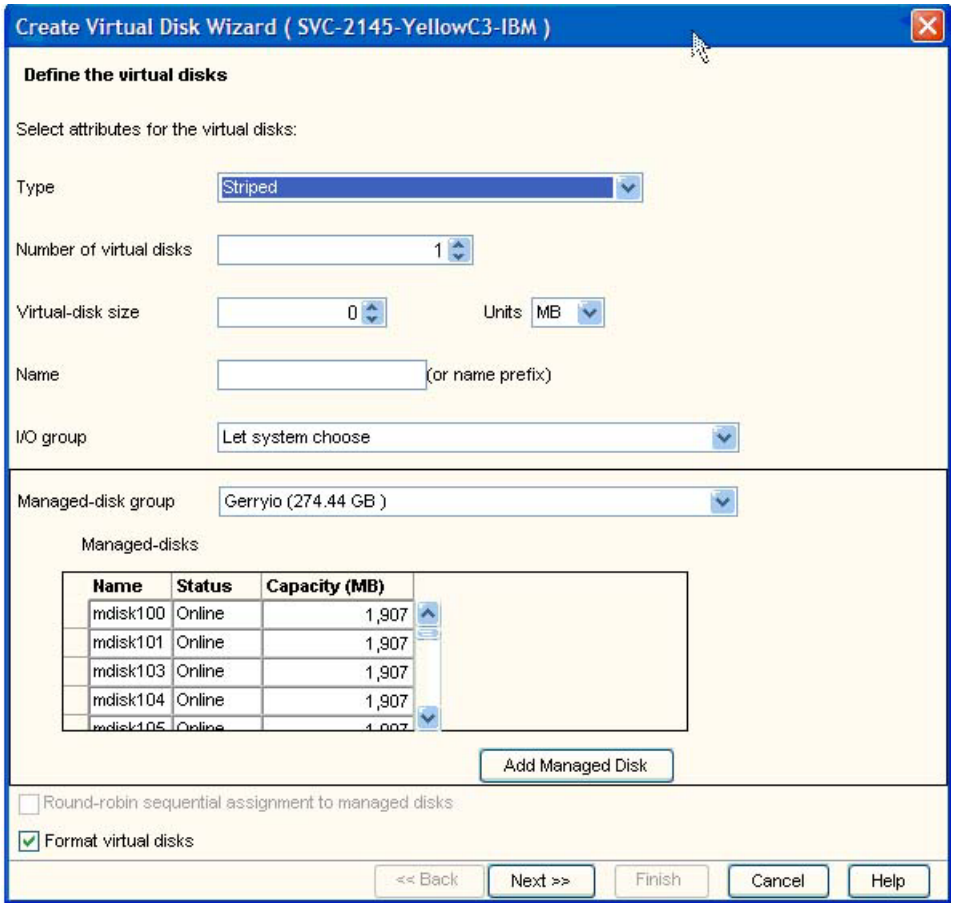

Figure 40. Create Virtual Disk wizard: "Define the virtual disks" page

3. On the "Define the virtual disks" page, specify the attributes for the virtual disks:
   a. In the **Type** field, select one of the following values:
      - Striped
      - Sequential
   b. In the **Number of virtual disks** field, select the number of virtual disks that you want to create. Each available I/O group can contain 1024 virtual disks.
   c. In the **Virtual-disk size** field, select the size of the virtual disks to be created.
   d. In the **Units** field, select one of the following units of measurement for the size of the virtual disks:
      - B (byte)
      - GB (gigabyte)
      - KB (kilobyte)
      - MB (megabyte)
      - PB (petabyte)
      - TB (terrabyte)
   e. In the **Name** field, type a string to identify the virtual disk. This string must adhere to the following conventions:
      - The maximum length of the name is 15 characters.

- Permissible characters include uppercase letters (A-Z), lowercase letters (a-z), numerals (0-9), hyphens (-), and underscores (_).
- The first character cannot be a numeral.
- The name cannot begin with an abbreviation commonly used to specify the object type, for example, the name cannot begin with VDisk, VDISK, or vdisk.

If you leave this field blank, a name is generated automatically. If you create multiple virtual disks, the name is used as a prefix.

f. In the **I/O group** field, select the I/O group. By default, **Let the system choose** is selected. Other options are the actual I/O groups on the cluster.

g. In the **Managed-disk group** field, select the managed-disk group. The managed-disk groups are listed using the following convention:

```
MDisk-group name (N)
```

where *N* is the available capacity. After you select the managed-disk group, the managed disks in the managed-disk group are displayed.

To add a managed disk to the managed-disk group, click **Add Managed Disk**.

h. (if you are creating more than one sequential virtual disk) Select the **Round robin sequential assignment to managed disks** check box to ensure that the virtual disks are assigned to the managed disks in a round-robin manner. If selected, the number of virtual disks must equal the number of managed disks selected in the managed-disk table.

Clear the **Round robin sequential assignment to managed disks** check box to ensure that all virtual disks are created on a single managed disk. If the check box is cleared, there must be *only* one managed disk selected in the managed-disk table.

i. Select the **Format virtual disks** check box to ensure that the virtual disks are formatted (written with zeros) after they are created.

j. Click **Next**.

4. On the "Assign the virtual disk to host ports" page, assign host ports:
   a. Optional: In the **Available ports** field, select the ports that you want to assign to the virtual disks, and then click  .
   b. Click **Next**.
5. (This page only is displayed if Fabric Manager is installed, and you have Fabric administrator authority.) On the "Zone actions" page, perform one of the following zone actions:

   Create a new zone
   Click **Create new Zone**. In the Zone name field, type the name of the zone. The ports listed in the field below are added to the new zone.

   Update an existing zone
   Click **Update existing zone**, and then select a zone. The ports listed in the field below are added to the selected zone.

   Leave the zone configuration as it is
   Click **Do neither** to avoid performing any fabric-control tasks. The ports listed in the field below will not be able to communicate with the subsystem on this fabric.
6. Click **Next**.

7. On the Summary page, review the selections that you made earlier in the wizard.
8. Click **Finish** to create the virtual disks.

Deleting a virtual disk

This topic describes how to delete a virtual disk.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems window is displayed.
2. In the **Storage Subsystems** table, click a SAN Volume Controller, and then click **Virtual Disks**. The Virtual Disks page opens.
3. Click either **Filter by MDisk group** or **Filter by host ports**, and then select the specific filtering condition.
4. Click **Go**. A list of virtual disks associated with the selected SAN Volume Controller is displayed.

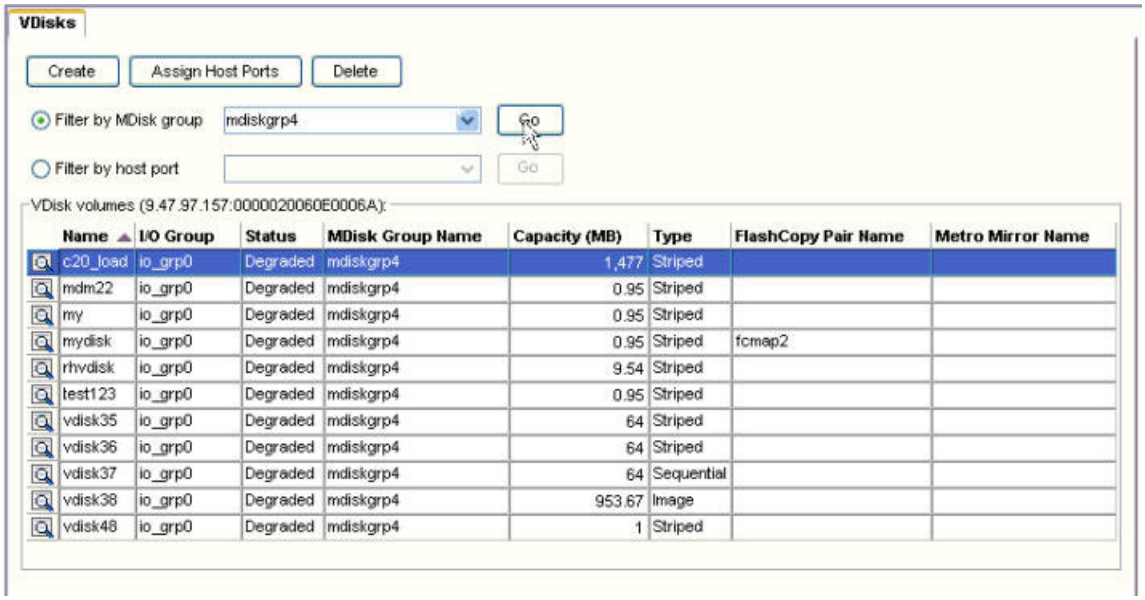

Figure 41. Virtual Disks page

5. In the **VDisk** table, click the virtual disk that you want to delete, and then click **Delete**.
6. When prompted to confirm that you want to delete the virtual disk, click **Yes**.

Viewing virtual-disk information

This topic describes how to view information about the virtual disks in your environment.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystems** table, click a SAN Volume Controller, and then click **Virtual Disks**. The Virtual Disks page opens.
3. Click either **Filter by MDisk group** or **Filter by host ports**, and then select the specific filtering condition.
4. Click **Go**. A list of virtual disks associated with the selected SAN Volume Controllers is displayed. The **VDisk** table contains the following information about each virtual disk:

   Name Name of the virtual disk or mirror copy.

I/O Group
Name of the I/O group that is associated with the virtual disk or mirror copy.

Status Status of the virtual disk or mirror copy:

- Offline
- Online
- Degraded

Copy Type
Indicates the type of the virtual disk. Valid values for this column include the following:

- **Standard:** Indicates standard VDisks (from pre-SAN Volume Controller 4.3 systems)
- **Primary:** Indicates VDisks from SAN Volume Controller 4.3 or newer systems.
- **Copy:** Indicates a mirror copy (not the primary copy) of a VDisk.

Note: You can assign virtual disks to hosts if those disks have a **Copy Type** of **Standard** or **Primary**.

Copy ID
Indicates the ID of a mirror copy for a virtual disk. Depending on the configuration history, a single mirror copy might have an ID of either 0 or 1.

Note: The value **N/A** appears in this column for virtual disks.

MDisk Group Name
Name of the managed-disk groups to which the virtual disk or mirror copy belongs.

Capacity (MB)
Capacity of the virtual disk or mirror copy in MB. This value is rounded to 0.1 MB.

Type Type of virtual disk or mirror copy:

- Image
- Sequential
- Striped

FlashCopy Name
Name of the FlashCopy mapping to which the virtual disk belongs.

MM Name
Name of the Metro Mirror relationship to which the virtual disk belongs.

5. To view detailed information about a specific virtual disk, click the virtual disk, and then click  . The Virtual-Disk Details page opens.

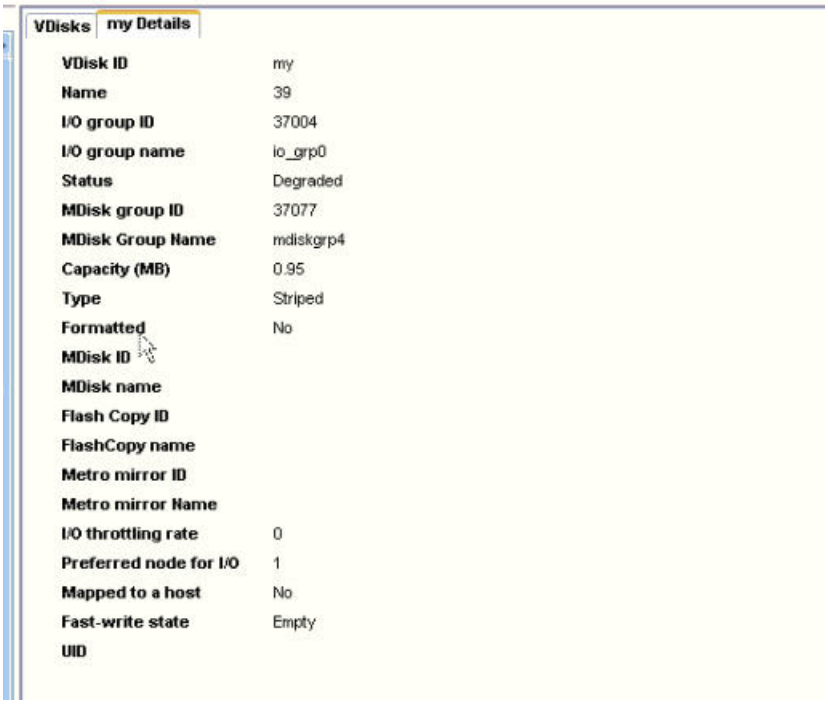

Figure 42. Virtual-Disk Details page

The Virtual-Disk Details page contains the following information:

VDisk ID
ID of the virtual disk.

Name Name of the virtual disk.

I/O group ID
ID of the I/O group.

I/O group name
Name of the I/O group.

Status Status of a virtual dis:

- Offline
- Online
- Degraded

MDisk group name
Name of the managed-disk group.

Type Type of a virtual disk:

- Image
- Sequential
- Striped

Capacity (MB)
Capacity of a virtual disk in MB. This value is rounded to 0.1 MB.

Formatted
Indicates if a virtual disk is formatted.

MDisk ID
IDs of the MDisks that a virtual disk is using.

MDisk Name
Names of the MDisks that a virtual disk is using.

Flash Copy
Indicates if the virtual disk is in a FlashCopy (FC) relationship and whether it is a FC source or FC target. Virtual disks that are not in a FC relationship will be displayed with a None value whether or not the virtual disk is on a subsystem that supports FC.

Note:

- This value is available for virtual disks of the following subsystems only: Tivoli Storage Enterprise Storage Server, DS6000, DS8000, and SAN Volume Controller.
- Tivoli Storage Enterprise Storage Server subsystems must have at least the following microcode levels: ess800mincodelevel = 2.4.3.56, essf20mincodelevel = 2.3.3.89. IBM Tivoli Storage Productivity Center will not report FC information if the Tivoli Storage Enterprise Storage Server subsystems do not meet this requirement.

FlashCopy ID
ID of the FlashCopy mapping to which the virtual disk belongs.

FlashCopy name
Name of the FlashCopy mapping to which the virtual disk belongs.

Metro-mirror ID
ID of the Metro Mirror relationship to which the virtual disk belongs.

Metro-mirror name
Name of the Metro Mirror relationship to which the virtual disk belongs.

I/O throttling rate
Maximum rate at which an I/O transaction is accepted for a virtual disk.

Preferred node for I/O
IP address of the node to be used for I/O operations.

Mapped to a host
Yes or No.

Fast-write state
Yes or No. This indicates whether the virtual disk has data in the cache.

UID Unique ID.

Managing volumes

This topic contains information about managing volumes.

Assigning host ports

This topic describes how to assign host ports to volumes. You also can make host-port assignments when you create volumes.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems window is displayed.
2. In the **Storage Subsystems** table, click the storage subsystem with which you want to work.
3. Click **Volumes**. The Volumes page is displayed.

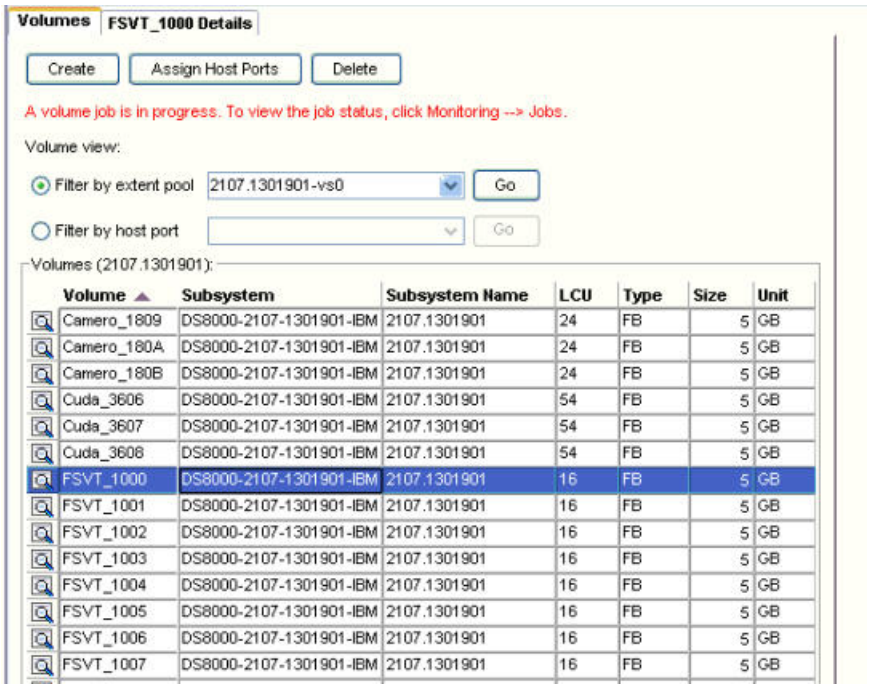

Figure 43. Volumes page

4. Select a filtering condition, and then click **Go**. A list of volumes is displayed.
5. In the **Volumes** table, select the volume, and then click **Assign Host Ports**. The Assign Host Ports wizard starts.
6. On the "Assign the volume to host ports" page, assign or unassign host ports:
   a. Optional: In the **Available ports** field, select the ports that you want to assign to the virtual disks, and then click >>.
   b. Optional: In the **Assigned ports** field, select any ports that you want to unassign from the virtual disks, and then click <<. To avoid lost or inconsistent data, make sure that the volume is unmounted before unassigning.
   c. Click **Next**.
7. On the "Assign the volume to subsystem ports" page, assign the volumes to specific subsystem ports.
   a. Optional: In the **Available ports** list, select subsystem ports and click >> to assign them to the volumes. The ports are identified using the following convention: `<Worldwide port name> (<host name>)`.
   b. Click **Next**.
8. (This page only is displayed if Fabric Manager is installed, and you have Fabric administrator authority.) On the "Zone actions" page, perform one of the following zone actions:

   Create a new zone
   Click **Create new Zone**. In the Zone name field, type the name of the zone. The ports listed in the field below are added to the new zone.

   Update an existing zone
   Click **Update existing zone**, and then select a zone. The ports listed in the field below are added to the selected zone.

Leave the zone configuration as it is

Click **Do neither** to avoid performing any fabric-control tasks. The ports listed in the field below will not be able to communicate with the subsystem on this fabric.

9. Click **Next**. The ″Review settings″ page opens.
10. Review the selections that you made earlier in the wizard, and then click **Finish**.

Creating volumes

This task describes how to create volumes.

The performance of a volume creation job and the general performance of a CIMOM might decrease if you create a large number of volumes on a subsystem at the same time a performance monitor is collecting data about that subsystem. We recommend creating volumes on a subsystem when a performance monitoring job is not running or scheduled to run on the target subsystem.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystem** table, click a storage subsystem, and then click **Create Volume**. The Create Volume wizard starts.
3. On the ″Define the volumes″ page, specify the attributes for the volumes:
   a. In the **RAID level** field, select one of the following RAID levels, based on the type of storage subsystem:
      - DS6000 and DS8000: RAID 5, RAID 10, or Mixed. RAID 5 is the default setting.
      - Tivoli Storage Enterprise Storage Server: RAID 5 or RAID 10. RAID 5 is the default setting.
      - FAStT or DS4xxxx: RAID 0, RAID 1, RAID 3, or RAID 5. RAID 0 is the default setting.

   b. In the **Available arrays** field, select an array from the list of arrays with the specified RAID level. These arrays are identified using the following convention: `<Device-specific array name> (<free space> GB)`. For DS6000 and DS8000 subsystems, only arrays with fixed-block (FB) tracks are listed.
   c. In the **Number of volumes** field, select the number of volumes that you want to create. By default, this is set to 1. For an Tivoli Storage Enterprise Storage Server subsystem, the maximum is 255.
   d. In the **Volume size** field, select the size of the volumes to be created. The maximum and minimum size depends on the type of storage subsystem.
   e. In the **Units** field, select one of the following units of measurement for the size of the volumes:
      - KB
      - MB
      - GB

   f. In the **Name** field, type a name for the volume. If you are creating multiple volumes, this string is used as a prefix. The string cannot be longer than 15 characters.
   g. Click **Next**.
4. On the ″Assign the volume to host ports″ page, assign the volumes to specific host ports.

a. Optional: In the **Available ports** list, select host ports and click  to assign them to the volumes. The ports are identified using the following convention: `<Worldwide port name> (<host name>)`.

b. Click **Next**.

5. On the "Assign the volume to subsystem ports" page, assign the volumes to specific subsystem ports.

a. Optional: In the **Available ports** list, select subsystem ports and click  to assign them to the volumes. The ports are identified using the following convention: `<Worldwide port name> (<host name>)`.

b. Click **Next**.

6. (This page only is displayed if Fabric Manager is installed, and you have Fabric administrator authority.) On the "Zone actions" page, perform one of the following zone actions:

Create a new zone
Click **Create new Zone**. In the Zone name field, type the name of the zone. The ports listed in the field below are added to the new zone.

Update an existing zone
Click **Update existing zone**, and then select a zone. The ports listed in the field below are added to the selected zone.

Leave the zone configuration as it is
Click **Do neither** to avoid performing any fabric-control tasks. The ports listed in the field below will not be able to communicate with the subsystem on this fabric.

7. Click **Next**.
8. On the Summary page, review the selections that you made earlier in the wizard.
9. Click **Finish**.

Note: If you are creating a large number of volumes, for example, more than 50, an hourglass icon might be displayed. This indicates that the graphical user interface (GUI) is waiting for the volume-creation jobs to be completed. Click  . This enables you to continue using the GUI while the volumes are created.

Creating volumes and assigning them to a host

1. Expand **Disk Manager** and click **Storage Subsystems**.
2. In the **Storage Subsystem** table, select a storage subsystem, and click **Create Volume**. The Create Volume wizard starts.
3. In the "Define the volumes" page, specify the attributes for the volumes:

a. In the **RAID level** field, select one of the following RAID levels, based on the type of storage subsystem:

- DS6000 and DS8000: RAID 5, RAID 10, or Mixed. RAID 5 is the default setting.
- DS6000: RAID 5 or RAID 10. RAID 5 is the default setting.
- FAStT or DS4xxxx: RAID 0, RAID 1, RAID 3, or RAID 5. RAID 0 is the default setting.

b. In the **Available arrays** field, select an array from the list of arrays with the specified RAID level. These arrays are identified using the following

convention: <*Device-specific array name*> (<*free space*> GB). For DS6000 and DS8000 subsystems, only arrays with fixed-block (FB) tracks are listed.

c. In the **Number of volumes** field, select the number of volumes that you want to create. By default, this is set to 1. For an Tivoli Storage Enterprise Storage Server subsystem, the maximum is 255.

d. In the **Volume size** field, select the size of the volumes to be created. The maximum and minimum size depends on the type of storage subsystem.

e. In the **Units** field, select one of the following units of measurement for the size of the volumes:

- KB
- MB
- GB

f. In the **Name** field, type a name for the volume. If you are creating multiple volumes, this string is used as a prefix. The string cannot be longer that 15 characters.

g. Click **Next**.

4. On the "Assign the volume to host ports" page, assign the volumes to specific host ports.

a. Optional: In the **Available ports** list, select host ports and click >> to assign them to the volumes. The ports are identified using the following convention: <*Worldwide port name*> (<*host name*>).

b. Click **Next**.

5. On the "Assign the volume to subsystem ports" page, assign the volumes to specific subsystem ports.

a. Optional: In the **Available ports** list, select subsystem ports and click >> to assign them to the volumes. The ports are identified using the following convention: <*Worldwide port name*> (<*host name*>).

b. Click **Next**.

6. On the Summary page, review the selections that you made earlier in the wizard.
7. Click **Finish**.

Deleting volumes

This topic describes how to delete volumes.

Note that volumes that are part of a replication session cannot be deleted.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystems** table, click a storage subsystem, and then click **Volumes**. The Volumes page is displayed.
3. On the Volumes page, select the specific filtering condition, and then click **Go**. A list of volumes on the selected storage subsystem is displayed.
4. In the **Volumes** table, select the volume that you want to delete, and then click **Delete**. You can select one volume at a time only. You cannot delete Tivoli Storage Enterprise Storage Server volumes from within IBM Tivoli Storage Productivity Center.
5. When prompted to confirm that you want to delete the volume, click **Yes**. The volume is deleted automatically when you click Yes. You can reuse the volume name immediately after the volume is deleted.

Note: For SAN Volume Controller (SVC), if you choose to delete a volume using the SVC console rather than the IBM Tivoli Storage Productivity Center user interface, that volume is retained for 14 days before being deleted. You cannot reuse the volume name during that time.

Viewing volume information

This topic describes how to view information about the volumes in your environment.

Note: If a probe is run against a CIMOM that is not available, the volumes defined on the storage subsystem are not displayed in the IBM Tivoli Storage Productivity Center GUI. Check that probes run successfully and consider configuring an alert to ensure that you are notified when a scheduled probe fails.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystems** table, click the storage subsystem, and then click **Volumes**. The Volumes page opens.
3. Select the specific filtering condition, and then click **Go**. A list of volumes associated with the selected storage subsystem is displayed. The **Volume** table contains the following information about each volume:

   Volume
   Volume ID.

   Subsystem
   A subsystem ID in the following format:

   Type-Model-ID-Vendor

   - *Type* is the type of the storage subsystem, for example, Tivoli Storage Enterprise Storage Server or SAN Volume Controller.
   - *Model* is the model number, for example, 2105 or 2145.
   - *ID* is the serial number or cluster name. Cluster names only are used for SAN Volume Controllers.
   - *Vendor* is the manufacturer of the storage system, for example, IBM.

   Subsystem Label
   A user-defined name or the subsystem ID.

   LCU Logical control unit (LCU), which can have one of the following values:
   - DS6000, DS8000, and Tivoli Storage Enterprise Storage Server -- logical subsystem (LSS) ID
   - SAN Volume Controller --1

   Type Type of volume, either fixed-block (FB) or count-key data (CKD).

   Volume Space
   Size of the volume.

   Unit GB for fixed-block (FB) volumes and cyl for count-key data (CKD) volumes.
4. To view detailed information about a specific volume, click the volume, and then click  . The Volume Details page opens. The Volume Details page contains the following information:

   Volume
   Volume ID.

Volume size
Volume size.

Volume real space
Physical allocated space of the volume. For normal volumes this is equal to the volume space. For space efficient/thin provisioned volumes, this is equal to the real space allocated when data is written to the volume. If Tivoli Storage Productivity Center detects that a volume is thin provisioned but cannot retrieve the Volume real space value, ″N/A″ is displayed.

Array Array on which the volume is created.

RAID level
RAID level of the volume.

FlashCopy
Indicates if the volume is in a FlashCopy (FC) relationship and whether it is a FC source or FC target. Volumes that are not in a FC relationship will be displayed with a None value whether or not the volume is on a system that supports FC.

Note:

- This value is available for volumes of the following systems only: IBM Tivoli Storage Enterprise Storage Server (Tivoli Storage Enterprise Storage Server), IBM System Storage DS6000, IBM System Storage DS8000, and IBM System Storage SAN Volume Controller.
- Tivoli Storage Enterprise Storage Server systems must have at least the following microcode levels: ess800mincodelevel = 2.4.3.56, essf20mincodelevel = 2.3.3.89. IBM Tivoli Storage Productivity Center will not report FC information if the Tivoli Storage Enterprise Storage Server systems do not meet this requirement.

Subsystem
A system ID in the following format:

Type-Model-ID-Vendor

where:

- *Type* is the type of the storage system, for example, Tivoli Storage Enterprise Storage Server or SAN Volume Controller.
- *Model* is the model number, for example, 2105 or 2145.
- *ID* is the serial number or cluster name. Cluster names only are used in for SAN Volume Controllers.
- *Vendor* is the manufacturer of the storage system, for example, IBM.

Subsystem Label
This either can be a user-defined name or the system ID.

Assigned host ports
Host ports to which the volume is assigned.

Subsystem ports
System ports to which the volume is assigned.

Managing fabrics

This section describes operations available for directly manipulating the zone configuration in a fabric.

Zone control functions are also integrated in the Storage Provisioning function (see "Provisioning planner" on page 63) and the SAN Planner functions (see "SAN Planner overview" on page 278).

Note: After you click **Change** to begin your zone configuration changes and before you click **Update and Activate** or **Update Only** to conclude your changes, it is possible that other users might have made zoning changes to the same fabric. This may have happened through a switch element manager or switch CLI, as some switch vendors and models lack zone control session locking across the fabric. This could allow overwrite conditions of the changes to the active zoneset for the fabric and to the set of inactive zonesets.

Working with zones

This topic describes how to create, change, duplicate, activate or deactivate, and delete zones.

Creating zones

When you add zones, you assign assets to the zone and specify the zone set or sets to which the zone belongs.

To create a zone:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zones** area, click **Add**.
5. In the Zone Configuration window enter the name of the zone and, optionally, a description. The zone name must be unique in the fabric. Click **Next**.
6. In the next Zone Configuration window, select the assets (switches, devices, aliases) that you want to include in the zone, move them from the **Available assets** column to the **Selected assets** column, and click **Next**.
7. In the next Zone Configuration window, select one or more zone sets, move them from the **Available Zone Sets** column to the **Selected Zone Sets** column, and click **Finish**.

Adding storage to a server and setting up a zone

To add storage to a server and set up a zone, do the following:

1. **Add the CIMOM**
   a. Expand **Administrative Services → Data Sources → CIMOM Agents** .
   b. Click **Add CIMOM** in the right pane.
   c. In the **Add CIMOM** window, enter the required values and click **Save**.
2. **Run a CIMOM discovery**
   a. Expand **Administrative Services → Discovery**.
   b. Right-click **CIMOM** and select **Run Now**.
   c. Right-click **CIMOM** and select **Refresh Job List** for the submitted job to appear in the list of discovery jobs. The date and time of the job is displayed in the left pane. To see the latest status of the job, right-click **CIMOM** and select **Update Job Status**.
3. **Run a probe**

   When the discovery has completed successfully, perform the following steps:

   a. Expand **IBM Tivoli Storage Productivity Center → Monitoring** .

b. Right-click **Probes** and select **Create Probe**.
c. Make your selections and click **Save** to submit the job.
d. Right-click **Probes** and select **Refresh Job List** to display the job in the list of probes. The date and time of the job will be displayed in the left pane. To see the latest status of the job, right-click **Probes** and select **Update Job Status**.

4. In the left pane, expand **Data Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems window opens.
5. In the **Storage Subsystems** table, click the storage subsystem with which you want to work.
6. Click **Volumes**.
7. In the Volumes page select a filtering condition, and click **Go**.
8. In the **Volumes** table, select the volume, and then click **Assign Host Ports**. The Assign Host Ports wizard starts.
9. On the ″Assign the volume to host ports″ page, you can assign and unassign host ports:
   a. In the **Available ports** field, select the ports that you want to assign to the virtual disks, and then click >> .
   b. In the **Assigned ports** field, select any ports that you want to unassign from the virtual disks, and then click << .
   c. Click **Next**.
10. On the ″Assign the volume to subsystem ports″ page, assign the volumes to specific subsystem ports.
    a. Optional: In the **Available ports** list, select subsystem ports and click >> to assign them to the volumes. The ports are identified using the following convention: <*Worldwide port name*> (<*host name*>).
    b. Click **Next**.
11. ″Zone actions″ page is displayed only if Fabric Manager is installed, and you have Fabric administrator authority. Perform one of the following zone actions:

    Create a new zone
    Click **Create new Zone**. In the Zone name field, type the name of the zone. The ports listed in the field below are added to the new zone.

    Update an existing zone
    Click **Update existing zone**, and then select a zone. The ports listed in the field below are added to the selected zone.

    Leave the zone configuration as it is
    Click **Do neither** to avoid performing any fabric-control tasks. The ports listed in the field below will not be able to communicate with the subsystem on this fabric.

    Click **Next**.
12. In the ″Review settings″ page, review your selections and click **Finish**.

Changing zone properties

You can change a zone name, a description, the assets assigned to a zone, and the zone set or sets to which a zone belongs.

To change zone properties, perform the following steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zones** area, click **Change**.
5. In the Zone Configuration window enter the name of the zone and, optionally, a description. The zone name must be unique in the fabric. Click **Next**.
6. In the next Zone Set Configuration window, you can change the assets to be included in the zone:
   - In the **Available Assets** column, select any zones to be added and click >> to move them to the **Selected Assets** column.
   - In the **Selected Assets** column, select any zones to be removed from the zone set and click << to move them to the **Available Assets** column.
   - Click **Next**.
7. In the next Zone Set Configuration window, you can change the zone sets to which the zone is assigned:
   - In the **Available Zone Sets** column, select any zone sets to which the zone is to be assigned and click >> to move them to the **Selected Zone Sets** column.
   - In the **Selected Zone Sets** column, select any zone sets to which the zone is not to be assigned and click << to move them to the **Available Zone Sets** column.
8. Click **Finish**.

Duplicating zones

You can use the Zones dialog to duplicate a zone.

To duplicate zones, perform the following steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zones** area, click **Duplicate**.
5. In the Zone Configuration window enter the name of the zone and, optionally, a description. The zone name must be unique in the fabric. Click **Next**.
6. In the next Zone Configuration window, you can change the assets to be included in the zone:
   - In the **Available Assets** column, select any zones to be added and click >> to move them to the **Selected Assets** column.
   - Click **Finish**.
7. In the next Zone Configuration window, you can change the zone sets to which the zone is assigned:
   - In the **Available Zone Sets** column, select any zone sets to which the zone is to be assigned and click >> to move them to the **Selected Zone Sets** column.
   - In the **Selected Zone Sets** column, select any zone sets to which the zone is not be assigned and click << to move them to the **Available Zone Sets** column.

8. Click **Finish**.

Removing zones

You can use the Zones dialog to remove one or more zones from the list of available zones.

To delete zones, perform the following steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zones** area, select one or more zones and click **Remove**. Hold down the Control or Shift key to select multiple zones.
5. In the Zone Deletion window, click **Yes** to delete the selected zones or **No** to cancel the deletion.

Working with zone sets

This topic describes how to create, change, duplicate, activate or deactivate, and delete zone sets.

Creating zone sets

This topic provides information about creating new zone sets.

To create a zone set, follow this procedure:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zone Sets** area, click **Add**.
5. Type a name and description for the zone set, and click **Next**. The zone set name must be unique within the fabric.
6. In the Zone Set Configuration window, do one of the following:
   - Select one or more zones from the **Available Zones** column, move them to the **Selected Zones** column, and click **Finish**.
   - If there are no zones yet defined or if you are not adding any of the zones available at this time, click **Finish**. A Zone Set Configuration Error window opens stating: ″This zone set does not contain any zones. This could result in an error when the zone configuration is applied later. Do you want to continue?″ Click **Yes**.

Changing zone set properties

You can change the zone set definition, which includes the name, description, and the zones that belong to the zone set. When you change and save a zone set definition, the properties are checked by the switches when that definition is stored in the fabric. This can help identify possible problems with that definition before the definition is activated, giving you a chance to correct any problems found before they affect the operation of the fabric.

Note: Changing and saving the *definition* for the active zone set does not change the active zone set itself. You must perform an additional step to apply the new definition to the active zone set.

Note: After you click **Change** to begin your zoning changes and before you click **Update and Activate** or **Update Only** to conclude your zoning changes, it is

possible that other users might have made zoning changes to the same fabric. This may have happened through a switch element manager or switch CLI, as some switch vendors and models lack zone control session locking across the fabric. This could allow overwrite conditions of the changes to the active zoneset for the fabric and to the set of inactive zonesets.

To change the zone set properties, perform the following steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zone Sets** area, click the zone set to be changed and click **Change**.
5. In the Zone Set Configuration window, type a name and description for the zone set, and click **Next**. The zone set name must be unique within the fabric.
6. In the next Zone Set Configuration window, you can change the zones to be included in the zone set:
   - In the **Available Zones** column, select any zones to be added and click  to move them to the **Selected Zones** column.
   - In the **Selected Zones** column, select any zones to be removed from the zone set and click  to move them to the **Available Zones** column.
7. Click **Finish**.

Activating and deactivating a zone set

Only one zone set can be active at a time. If you deactivate the active zone set, no zone sets are active until you activate another zone set.

To activate or deactivate a zone set, follow these steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zone Sets** area, click **Activate** or **Deactivate**.
5. In the **Zone Configuration** area, click one of the following actions:
   - **Update and Activate** to commit the changes you have made to the zone set configuration and immediately activate the zone set.
   - **Update Only** to commit the changes you have made to the zone set configuration without making them active.

Selecting the inactive zone sets data source

Use the procedure below to change the data source for the inactive zone sets of a selected fabric.

Note: This procedure only applies to mixed Brocade and McDATA fabrics that are managed through CIMOMs.

To change the data source for an inactive zone set, follow these steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Select Inactive Zone Sets Data Source**. Use the **Select Inactive Zone Sets Data Source** dialog box to select a data source to be used for gathering and setting the inactive zone sets for the selected fabric. This dialog box has the following fields:

- **Service URL** displays the CIMOM service URL that can be selected and used for zone control of inactive zone sets for the selected fabric.
- **Display Name** lists the user-defined display name for the CIMOM.
- **Connection Status** lists the connection status to the CIMOM.

3. Select a data source and click **OK** to set the data source to be used for probing and controlling inactive zone sets for the fabric. After you select a different data source, a zone probe job is started in the background. A ″A probe job for fabric <*fabric ID*> has been submitted″ message box appears. Click **OK** to close the message box.

Duplicating zone sets

This section provides information about duplicating zone sets.

To create a duplicate zone set, perform the following tasks.

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zone Sets** area, click the zone set to be duplicated and click **Duplicate**.
5. In the Zone Set Configuration window, type a name and description for the zone set, and click **Next**. The zone set name must be unique within the fabric.
6. In the Zone Set Configuration window, you can change the zones to be included in the new zone set:
   - In the **Available Zones** column, select any zones to be added and click >> to move them to the **Selected Zones** column.
   - In the **Selected Zones** column, select any zones to be removed from the zone set and click << to move them to the **Available Zones** column.
7. Click **Finish**.

Working with zone aliases

This topic describes how to add, change, duplicate, and remove zone aliases. Using zone aliases is a convenient way to work with a named collection of zone members that you want to manage together.

Creating zone aliases

Creating a zone alias is a convenient way to assign a name to a collection of zone members that you want to manage together.

To create a zone alias:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zones Aliases** area, click **Add**.
5. Enter a unique name for the zone alias that will have zone members. Use the following guidelines for naming a zone alias:
   a. The maximum length for the zone alias name is 54 characters.
   b. Allowed characters are: A-Z and a-z; _ (underscores)
   c. The name must not begin with a number.
   d. Spaces are not allowed.

6. Enter an optional description of the zone alias. The description can be any combination of alphanumeric characters.
7. Click **Next**.
8. In the next Zone Alias Configuration window, select the assets (switches, devices, aliases) that you want to include in the zone alias, and move them from the **Available assets** column to the **Selected assets** column, and click **Next**.
9. In the final Zone Alias Configuration window, select one or more zones, move them from the **Available Zones** column to the **Selected Zones** column, and click **Finish**.

Changing zone alias properties

You can change a zone alias name, description, the assets assigned to it, and the zone members that are specified in the zone alias.

To change zone alias properties, perform the following steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zone Aliases** area, select the zone alias you want to change and click **Change**.
5. In the Zone Alias Configuration window change the name of the zone alias and, optionally, enter or change the description. The zone alias name must be unique in the fabric. Use the following naming guidelines:
   a. The maximum length for the zone alias name is 54 characters.
   b. Allowed characters are: A-Z and a-z; _ (underscores)
   c. The name must not begin with a number.
   d. Spaces are not allowed.
6. Click **Next**.
7. In the next Zone Alias Configuration window, you can change the assets to be included in the zone alias:
   - In the **Available Assets** column, select one or more assets to be added and click the  to move them to the **Selected Assets** column.
   - In the **Selected Assets** column, select one or more assets to be removed from the zone alias and click the  to move them to the **Available Assets** column.
   - Click **Next**.
8. In the final Zone Alias Configuration window, you can change the zone members as follows:
   - In the **Available Zones** column, select one or more zone members and click the  to move them to the **Selected Zones** column.
   - In the **Selected Zones** column, select one or more zone members and click the  to move them to the **Available Zones** column.
9. Click **Finish**.

Duplicating zone aliases

To create a new zone alias that's based on an existing one, you can duplicate a zone alias and then make changes to it.

To duplicate an existing zone alias and make changes to it, perform the following steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zone Aliases** area, select the zone alias you want to duplicate and click **Duplicate**.
5. In the Zone Alias Configuration window change the name of the zone alias and, optionally, enter or change the description. The zone alias name must be unique in the fabric. Use the following naming guidelines:
   a. The maximum length for the zone alias name is 54 characters.
   b. Allowed characters are: A-Z and a-z; _ (underscores)
   c. The name must not begin with a number.
   d. Spaces are not allowed.
6. Click **Next**.
7. In the next Zone Alias Configuration window, you can change the assets to be included in the new zone alias:
   - In the **Available Assets** column, select any assets to be added and click the  to move them to the **Selected Assets** column.
   - In the **Selected Assets** column, select any assets to be removed from the zone alias and click the  to move them to the **Available Assets** column.
   - Click **Next**.
8. In the final Zone Alias Configuration window, you can change the zone alias as follows:
   - In the **Available Zones** column, select one or more zones and click the  to move them to the **Selected Zones** column.
   - In the **Selected Zones** column, select one or more zones and click the 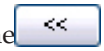 to move them to the **Available Zones** column.
9. Click **Finish**.

Removing zone aliases

The following procedure removes a selected zone alias from all zones in the currently selected fabric, and removes the zone alias from the fabric's zoning configuration definition. To remove a zone alias from an individual zone or a specific set of zones, use the Zone Configuration wizard.

To remove a zone alias, perform the following steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zone Aliases** area, select the zone alias you want to remove and click **Remove**. Hold down the Control or Shift key to select multiple zone aliases.
5. In the Zone Alias Deletion window, click **Yes** to delete the selected zone alias.

Managing Tape Libraries

IBM Tivoli Storage Productivity Center provides the capability to manage your tape libraries. For this release, tape library management capabilities are limited to support for the IBM 3584 tape library (full support) and the IBM 3494 tape library (partial support).

You must have at least Tape Operator authority in order to access the functions in the Tape Manager node. To view your authority level, open the ″Role-to-Group Mappings″ window in the **Administrative Services** → **Configuration** node in the navigation tree.

1. Discover IBM tape libraries attached to your SAN fabrics.
2. View tape library information.
3. Monitor the performance of groups of tape libraries by probing for data and viewing alerts. .
4. Create and view asset and capacity reports using data collected from your tape libraries.

   Asset reports can include such information as the library identifier and the status of the library. Capacity reports include such information as the number of drives in a library, the number of tape slots, and the number of occupied tape slots.
5. Use the topology viewer to view information about tape libraries, as described in “Working with tape libraries in the topology viewer” on page 404.

Discovering tape libraries

Discovering tape libraries uses the same general process for discovering CIMOMs as discovering storage subsystems does.

- Tape Manager automatically discovers available devices by contacting the CIM agent that is running on those devices.
- You can add Tape CIMOMs by automatic SLP or manually.
- You can create discovery jobs to discover other Tape Library CIM agents, as described in “Discover storage subsystem, tape library, and fabric information” on page 197.
- You can set up your discoveries to run on a regular schedule, for example weekly.

However, you view your tape libraries from a different location in the navigation tree.

- To confirm the tape libraries that were discovered, complete the following steps:
  1. In the navigation tree, expand **Tape Manager**.
  2. Click **Tape Libraries**. The tape libraries that have been discovered will be listed in the right pane.

IBM Tivoli Storage Productivity Center can detect a tape library that is connected to a fabric through queries to the switches in the fabric and through CIMOM queries. However, only basic information is available at this point, and the information from the two discoveries is not correlated. As a consequence, the tape library is displayed in the topology L0:Storage view as two separate icons:

- After the in-band discovery, as a tape library of unknown status (a blue diamond ) identified only by its world-wide name (WWN)

- After the CIMOM discovery, a tape library is identified as, for example, a 3494 or a TS3500

In order to correlate the information collected through the in-band and CIMOM discoveries, you must run a tape library probe. Then the tape library will be displayed by only one icon with its type identified.

Viewing tape library information

View collected tape library data to stay aware of discovered tape libraries and usage information.

- Click **Tape Manager**. Left-click **Tape Libraries**. In the right pane, the Tape Libraries page is displayed, which lists the discovered tape libraries.

  Select a library from the table and then click one of the buttons to see further details about that library. For example, select a library and click **Drives** to open a page that lists all the drives for that library.

  You can also right-click a library in the **Tape Libraries** table and select **View Details**.
- Use the **View Details** feature in the topology viewer, as described in "Working with tape libraries in the topology viewer" on page 404.

Monitoring tape libraries

Monitor the performance of tape libraries to be aware of events that occur, and to respond to alerts.

1. Create one or more tape library groups, as described in "Creating tape library groups" on page 237.
2. Create a probe job to collect data about your tape libraries and library groups, as described in "Creating tape library probes" on page 205.
3. View alerts that are generated by your tape libraries, as described in "Viewing alert logs" on page 262.

Launch in context

The Launch in context feature is available in two modes: inbound and outbound. With the inbound launch in context feature, you can start IBM Tivoli Storage Productivity Center at a specified point in the graphical user interface (GUI). With the outbound launch in context feature, you can use various points in the Tivoli Storage Productivity Center GUI to start an application such as an element manager at a specified point in the manager.

Inbound launch in context feature

With the inbound launch in context feature you can start IBM Tivoli Storage Productivity Center at a specified point in the graphical user interface (GUI).

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

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

- Create a volume or VDisk
- Open the storage area network SAN Planner
- Create a capacity report

- Create a wasted space report

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

- **Use Java Web Start** from a Web browser. See "Launch with Java Web Start."
- **Use the command line on Windows, AIX, Linux, or UNIX** to enter the command and arguments on the command line. See "Launch from the command line" on page 443.
- **Use the Windows Start menu**. See "Launch from Start menu" on page 444.
- **Use the Tivoli Enterprise Portal GUI** to define, edit, and launch commands. See "Launch from Tivoli Enterprise Portal" on page 445.

The launch in context feature has a number of parameters. See "Launch in context parameters" on page 785, "Launch in context task parameters" on page 786, and "Available windows" on page 788.

Launch with Java Web Start

Using Java Web Start, you can use the launch in context feature from a Uniform Resource Locator (URL).

The launch in context feature requires Java Web Start 1.5 on the computer that is running the Web browser. If the correct release of Java Web Start is not present, you receive an error message when you enter the URL. The message states that you must download the Java Runtime Environment, which includes Java Web Start. Follow the link provided in the message to install the required software.

If you are using Internet Protocol Version 6 (IPv6) and are using Java Web Start to start IBM Tivoli Storage Productivity Center on a Windows system, use a system host name such as **tpc332.storage.newyork.xyz.com:9549** in the Web browser address field rather than an IPv6 address such as **[2001:DB8:0000:1234:0000:0000:5678:ABCD]:9549**.

To use the launch in context feature to perform a task or display a window on a remote server from your browser, complete the following steps:

1. Open your browser and type the launch in context command.

   URLs cannot include spaces. In the command, a parameter and its value are separated by an equal sign (=). Parameter=value pairs are separated by an ampersand (&).

   For example, the format of a command to create a volume is:

   ```
   https://device_server:device_port/ITSRM/jsp/jnlp/
   tpcgui.jsp?consoleClass=com.ibm.usmi.console.lic.tpc.LICConsole&-
   user=userid&-passwd=password&-server=server:port&-
   function=disk.create_volume&subsystem_id=ESS-2105-17844-IBM&-name=myvol
   &-number=2&-size=10&-unit=MB
   ```

   For example, the format of a command to view a window on a remote server is:

   ```
   http://device_server:device_port/ITSRM/jsp/jnlp/
   tpcgui.jsp?consoleClass=com.ibm.usmi.console.lic.tpc.LICConsole&-
   user=userid&-passwd=password&-server=server:port&-function=window_name
   ```

   Tip: The Data server port value is 9549. The port value for the Device server is 9551 for https and 9550 for http. If the wrong the Device server port value is used, the return page cannot be interpreted by the browser. If the wrong Data server port value is used, the following error message is returned.

```
Unable to connect to the server
```

2. The first time a window or task is started, you receive a security warning informing you that the downloaded jar files have been digitally signed using a certificate from a trusted company. To prevent this message from being displayed in the future, select **Always trust content from this publisher** and click **Run**. If the Tivoli Storage Productivity Center graphical user interface (GUI) has not been downloaded on the local computer using Java Web Start, or the local computer does not have the latest version of the Tivoli Storage Productivity Center GUI, the GUI is downloaded.
3. If you included the user ID and password in the command, the sign-in dialog is bypassed, otherwise log on to the Tivoli Storage Productivity Center GUI. The Tivoli Storage Productivity Center GUI navigates to the tree node corresponding to the window name that you specified. You are connected using the user name and password that you entered. Depending on the value you specified for the `-function` parameter, the following actions might occur:
   - The specified window opens
   - For disk.create_volume, depending on the parameters, the Create Volume wizard or Create VDisk wizard opens.
   - For disk.san_planner, the SAN Planner panel opens with the given SAN plan ID.
   - For data.filesystem_capacity_by_computer, the Capacity Report opens.
   - For TPC.reports_data_wasted_space, the Wasted Space Report opens.

Restrictions:

- When you open a plan in the SAN Planner using launch in context, you must enter the exact name of the plan. The value for the **-plan_id** parameter is case-sensitive and must match the name of the plan when it was created. For example, if you created a plan named **UserPlan1**, you must specify **-plan_id UserPlan1** in the launch in context command.
- There is a known limitation when you launch the Filesystem Capacity > By Computer report using launch in context, you must specify the ID for the computer upon which you want to report. For example:

  ```
  http://tpcserver14:9550/ITSRM/jsp/jnlp/
  tpcgui.jsp?consoleClass=com.ibm.usmi.console.lic.tpc.LICConsole&-
  user=root&-passwd=passw0rd&-server=tpcserver14.storage.xyz.com:9549&-
  function=data.filesystem_capacity_by_computer&-computer_id=tpcserver14
  ```

  The value you enter for the **computer_id** parameter might not match the ID for that computer that is stored in the Tivoli Storage Productivity Center database repository. You receive a warning message if the ID for the computer that you enter in the URL does not match the ID that is stored in the data repository. To ensure that this report is displayed properly when you run from the launch in context facility, enter an ID for the computer that matches the ID that is displayed for it through the Topology Viewer or appropriate report in the product's user interface.

Use the information in the links below to enter parameters. If any of the parameters are spelled incorrectly, an error message indicates the incorrect parameter.

"Launch in context parameters" on page 785
You can use the launch in context parameters, definitions, and examples.

"Launch in context task parameters" on page 786
Use the launch in context feature task parameters to perform tasks.

"Available windows" on page 788
You can use windows with the inbound launch in context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the `-function window name` argument.

Launch from the command line

This topic describes how to start windows and tasks from the Windows AIX, Linux, or UNIX command line.

The bat file `tpc.bat` is used for Windows operating systems. The .sh file `TPCD.sh` is used for AIX, Linux, or UNIX systems. The parameters follow the file name. Parameters are separated by spaces. Parameters are optional and can appear in any order. See "Launch in context parameters" on page 785, "Launch in context task parameters" on page 786, and "Available windows" on page 788 for the parameters.

Frequently used commands can be saved as unique command files and rerun as needed.

1. In the Command Prompt window, type the command for the action that you want to take, and press Enter.

   The action occurs on the remote data server. The graphical user interface (GUI) window for that action opens on the local computer.
2. Enter parameters if you previously did not enter any parameters. If you do not provide a user name, password, or server name, the IBM Tivoli Storage Productivity Center GUI opens the connection dialog box. If you do not provide these values, the GUI presents an error message box and requests the values. If you specify incorrect parameters, you receive an error message.

The following example launches a task from the Windows command line:

```
tpc.bat -user admin -passwd password1 -server
tivolii122.storage.usca.ibm.com:9549 -function
data.filesystem_capacity_by_computer -computer_id <computerID>
```

The following example uses the launch in context feature to view a window on the remote machine:

```
tpc.bat -user admin -passwd password1 -server
tivolii122.storage.usca.ibm.com:9549 -function disk.storage_subsystem
```

Restrictions:

- When you open a plan in the SAN Planner using the command line or a Web browser for launch in context, you must enter the exact name of the plan. The value for **-plan_id** in the command line is case-sensitive and must match the name of the plan when it was created. For example, if you created a plan named **UserPlan1**, you might enter the following command on the command line:

  ```
  ./TPCD.sh -user administrator -passwd passw0rd -server
  tpcserver14.storage.xyz.com:9549 -function disk.san_Planner -plan_id
  UserPlan1
  ```
- There is a known limitation when you start the Filesystem Capacity > By Computer report using launch in context. To start this report, you must specify the ID for the computer upon which you want to report. For example:

```
./TPCD.sh -user root -passwd "passw0rd" -server
tpcserver14.storage.xyz.com:9549 -function
data.filesystem_capacity_by_computer -computer_id tpcserver14
```

The value you enter for "computer_id" might not match the ID for that computer that is stored in the Tivoli Storage Productivity Center database repository. You will receive a warning message if the ID for the computer you enter in the command line does not match the ID stored in the data repository. To ensure that this report is displayed properly when run from the launch in context command line or through Java Web Start, make sure to enter an ID for the computer that matches the ID that is displayed for it through the Topology Viewer or appropriate report in the product's user interface.

Launch from Start menu

This topic describes how to start windows and actions from the Windows Start menu.

The bat file, `tpc.bat`, is used for Windows operating systems. The parameters follow the bat file name. Parameters are optional and can appear in any order. To use the launch in context feature from the Windows Start menu, complete the following steps:

1. Click **Start > Programs > IBM Tivoli Storage Productivity Center**, right-click **Productivity Center** and select **Properties**. The Productivity Center Properties window opens.
2. Enter the command-line arguments in the **Target** text field and click **Enter**.

   If the Data server is remote, the action occurs on the Data server and the graphical user interface (GUI) window the action opens is on the local computer.
3. Enter the parameters, if you did not previously enter the necessary parameters. If you do not provide the user name, password, or server name, the Tivoli Storage Productivity Center GUI opens the connection dialog box. If you do not provide these values, the GUI presents an error message box and requests the values.

The following example launches a task from the Windows Start menu:

```
C:\Program Files\IBM\TPC\gui\tpc.bat -user admin -passwd password1 -server
tivolii122.storage.usca.ibm.com:9549 -function
data.filesystem_capacity_by_computer -computer_id computerID
```

The following example uses the Windows Start menu to view a window on the remote server is:

```
C:\Program Files\IBM\TPC\gui\tpc.bat -user admin -passwd password1 -server
tivolii122.storage.usca.ibm.com:9549 -function disk.storage_subsystem
```

For more information see:

"Launch in context parameters" on page 785
You can use the launch in context parameters, definitions, and examples.

"Launch in context task parameters" on page 786
Use the launch in context feature task parameters to perform tasks.

"Available windows" on page 788
You can use windows with the inbound launch in context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the `-function window name` argument.

Launch from Tivoli Enterprise Portal

This topic shows you how to start an existing launch definition from the Tivoli Enterprise Portal graphical user interface (GUI).

You must configure Tivoli Enterprise Portal before you can use the launch in context feature. For information on prerequisites and details on how to configure the launch in context feature with Tivoli Enterprise Portal, see "Configuring Tivoli Enterprise Portal." After Tivoli Enterprise Portal has been configured, you can start the application from the configuration window or use **Launch** from the Tivoli Enterprise Portal navigation tree.

If you use the Tivoli Enterprise Portal navigation tree, **Launch** will open the Tivoli Enterprise Portal graphical user interface with your predefined launch in context commands. To use one of these commands, complete the following steps:

1. In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems *your_server_name* > Windows OS**. Right-click **Disk** and select **Launch**. Disk is one of several values you can select. Values include Disk, Network, Enterprise Services, Memory, Printer, Process, Processor, and System.
2. The Create or Edit Launch Definitions window opens showing a list of defined launch definitions in the **Existing Launch Definitions** area.
3. Select the launch definition you want to use from the list and click **Launch**. The window or task that is defined by that launch definition opens.

For information on editing the launch in context command with the Tivoli Enterprise Portal GUI, see:

Related tasks

"Editing launch definitions" on page 447
You can edit an existing launch definition with the Tivoli Enterprise Portal graphical user interface (GUI).

Configuring Tivoli Enterprise Portal:

You can configure Tivoli Enterprise Portal so that you can use the launch in context feature.

To configure Tivoli Enterprise Portal complete the following steps:

1. The initial steps depend on whether you are going to open a window or perform a task.
   a. If you are opening a window, in the Tivoli Enterprise Portal navigation window on the local computer, select where you want the launch definition to originate from. This opens the corresponding view on the target computer. Right-click one of the following locations.
      - The navigator item
      - A row in a table or situation event console view
      - A slice of a pie chart
      - A bar of a bar chart

      If the launch definition originates from a table, chart, or situation event console view, open the workspace containing the view. If you think you might edit the argument later, create the launch definition from a table or chart view.
   b. If you are going to perform a task, create a launch definition for one of the following tasks:

- Create a volume or VDisk
- Open the SAN Planner
- Create a capacity report
- Create a wasted space report

2. In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems > *your_server_name* > Windows OS**. Right-click **Disk** and select **Launch**. Disk is one of several values you can select. Values include Disk, Network, Enterprise Services, Memory, Printer, Process, Processor, and System.
3. The Create or Edit Launch Definitions window opens. Click **Create New**.
4. In the **Name** text field, type a descriptive name for this launch definition. This name is used in the definition list.
5. In the **Target** text box click **Browse** to locate the `tpc.bat` file. The Select files for directories for Launch Definitions window opens. A typical location for Windows systems is `C:\Program Files\IBM\TPC\gui\tpc.bat`. Select the correct directory and click **OK**.

   When you start a launch definition from the command line, if any part of the path contains spaces, enclose the entire entry in quotation marks.
6. In the **Arguments** text field, enter any parameters that will be added when the application is started. Arguments are entered as *`-parameter value`*. In the Create or Edit Launch Definitions window, click **Browse** next to **Arguments** to locate the parameters that you want to use. The Select Browser Options window opens. Click **Substitutable Items** and then click **OK**. Select parameters from the **Selected Context** and **Extended Context** lists.

   The **help** argument lists the available **function** values. You must type the help parameter by itself into the **Arguments** text field or into the form `tpc.bat -help functions` on the command line for Windows or `TPCD.sh -help functions` for UNIX. The **-help functions** parameter provides a window that lists the nodes in the navigation tree that are available to start. If you type **-help** without the **functions** parameter, it results in an error message.

   When all parameters have been selected, click **OK** to return to the Create or Edit Launch Definitions window.
7. If the parameters can be changed at launch time (for example, you want to change the system ID), check the **Edit arguments before execution** box.
8. Specify a value for the **Start in** location. Click **Browse** to locate the IBM Tivoli Storage Productivity Center installation and select the folder **gui** where tpc.bat is located.
9. Click **Evaluate** to see how the command resolves when the application starts. Correct any of the parameters that are incorrect and reevaluate the command.
10. When you are satisfied with the launch definition, click **Apply** to save your changes. This keeps the window open so that you can start the current launch definition or create a new one. Click **OK** to save the launch definition and close the window.
11. To start the current definition, select **Launch**. The Tivoli Storage Productivity Center graphical user interface (GUI) with the function parameter that you have provided is opened.

Related reference

"Launch in context parameters" on page 785
You can use the launch in context parameters, definitions, and examples.

"Launch in context task parameters" on page 786
Use the launch in context feature task parameters to perform tasks.

"Available windows" on page 788
You can use windows with the inbound launch in context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the `-function` *`window name`* argument.

Editing launch definitions:

You can edit an existing launch definition with the Tivoli Enterprise Portal graphical user interface (GUI).

If the **Edit arguments before execution** box was checked when a launch definition was configured, you can edit the launch definition by changing the attribute values. To edit a launch definition, complete the following steps:

1. In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems > *your_server_name* > Windows OS**. Right-click **Disk** and select **Launch**. Disk is one of several nodes you can select. Nodes include Disk, Network, Enterprise Services, Memory, Printer, Process, Processor, and System.
2. The Create or Edit Launch Definitions window opens showing a list of launch definitions in the **Existing Launch Definitions** area.
3. Select the launch definition that you want to edit. The values are displayed.
4. Make the necessary changes in the text fields.
5. Click **Evaluate** to see how the arguments will resolve when the application starts. Correct any incorrect arguments and reevaluate the command.
6. Click **Apply** to save the changes and keep the window open. Click **OK** to save the launch definition and close the window.

   If you do not want to overwrite the existing launch definition, save the new launch definition with a unique name.

Outbound launch in context feature

With the outbound launch in context feature you can start applications such as element managers from various points in the IBM Tivoli Storage Productivity Center graphical user interface (GUI).

Note that the launch in context feature is not available for the IBM XIV Storage System element manager, XIV Storage Manager.

Launch in context menu

When you right-click a storage device in the Topology Viewer graphic view, for example, you see a launch in context menu for that device. The content of the launch in context menu varies as described in "Launch in context menu items" on page 448. The launch in context menu is displayed when you right-click various objects in the following areas of the Tivoli Storage Productivity Center GUI:

- **Alert Log Viewer**. Each row of the alert log is enabled for a launch in context menu. The context is determined by the type of the alert. Alert types include:
  - Replication
  - Storage Subsystem

  - Computer
  - Fabric
  - Switch
  - Port
  - Tape Library
- **Disk Manager Storage Subsystems table**
- **Fabric Manager Fabrics table**
- **Tape Manager Tape Libraries table**
- **Element Manager tab** of the Configuration Utility
- **Topology Viewer objects**, as listed in the table below

Topology Object	Icon
Fabric	
Switch	
Switch port	
Tape library	
Storage subsystem	
Storage subsystem port	
File system volume	
Storage disk	
Storage pool	

Launch in context menu items

The number of top-level menu items in the launch in context menu depend on the context of the entity that you select.

- **Launch in context defined application**. This menu item displays for each application that is defined by a launch descriptor file. Tivoli Storage Productivity Center provides launch descriptor files for DS8000 Storage Manager and IBM Tivoli Storage Productivity Center for Replication. If the corresponding launch descriptor files are found, the following menu items are displayed:

DS8000 Storage Manager
Click this menu item to start DS8000 Storage Manager.

Replication Storage Systems Overview
Click this menu item to start Tivoli Storage Productivity Center for Replication. The panel that is opened in the Tivoli Storage Productivity Center for Replication graphical user interface (GUI) depends on the context for the entity that you selected. For example, if you right-click a storage device on the **Element Manager** tab of the Configuration Utility and click **Replication Storage Systems Overview**, the Tivoli Storage Productivity Center for Replication Storage Systems window opens.

If the launch descriptor file cannot be found for DS8000 Storage Manager or Tivoli Storage Productivity Center for Replication, the **Launch file items error** menu item is displayed for that entity. Click this item to view a description of the problem. The **Launch file items error** menu item is also displayed in situations such as the IBM Tivoli Integrated Portal server is not installed or is not accessible, the Tivoli Storage Productivity Center Data server is not accessible, or another system error occurred.

- **External Tools.** This menu item shows any external tools that are defined in the Configuration Utility. Examples of external tools include Web-based user interfaces and local executable files. For example, you might add the address of a Web site that you frequently access or an application that you often use (a text editor, a spreadsheet application, and so on). External tools are not device specific. Any defined external tool is displayed in the launch in context menu for every device that is defined in Tivoli Storage Productivity Center.
- **Element Manager.** This menu item shows any element managers that were discovered for the device by a Storage Management Initiative Specification (SMI-S) query. To avoid menu item duplication, this item is not displayed when the associated element manager is defined by a launch descriptor file. In this situation, the element manager is displayed as a **Launch in context defined application** menu item.
- **User Defined Tools.** This menu item shows any user-defined element managers that are defined in the Configuration Utility. These element managers differ from the element managers in the Element Manager menu because they are user-defined and are not discovered by a SMI-S query.

Example launch in context menu

The following figure shows a launch in context menu for a storage system in the **L0:Storage** topology view. In this example, there is no **Element Manager** menu item because the element manager that was discovered by the SMI-S query is a duplicate of the element manager defined by the launch descriptor file (**DS8000 Storage Manager**).

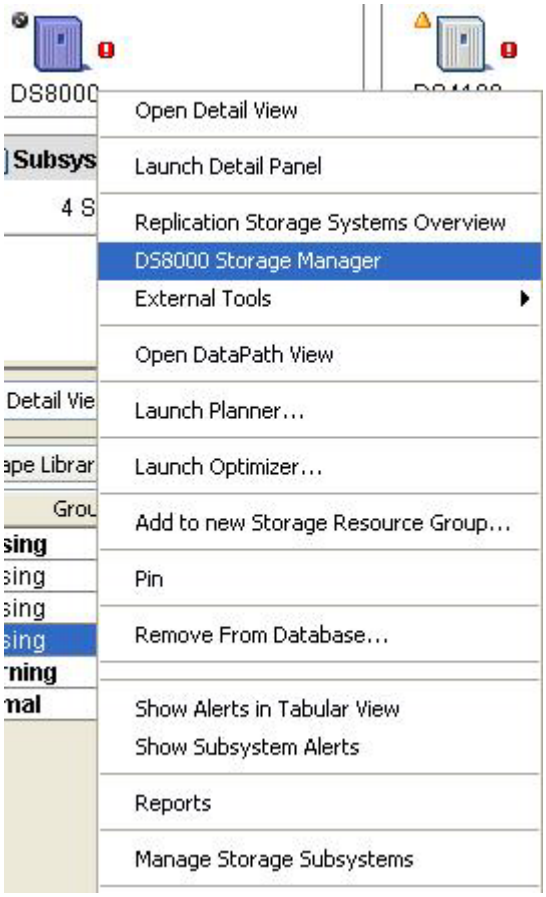

Using the single sign-on feature with launch in context

Single sign-on is an authentication process in which you can access more than one system or application by entering a single user ID and password. Single sign-on enables you to start applications from Tivoli Storage Productivity Center without providing logon credentials for the application.

For Tivoli Storage Productivity Center for Replication Storage Manager, single sign-on is enabled by default. For DS8000, single sign-on must be enabled as described in the Tivoli Storage Productivity Center information center. Single sign-on is not available for the IBM XIV Storage System element manager, XIV Storage Manager.

Related concepts

"Single sign-on" on page 54
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 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 in context feature to enable you to move smoothly from one application to a functionally-related location in a second application.

External tools

External tools include Web-based user interfaces and local executable files, scripts, and batch files that you want to access from IBM Tivoli Storage Productivity Center. For example, you might define the address of a Web site that you frequently access or an application that you often use (a text editor, a spreadsheet application, and so on) as an external tool. External tools such as executable files, scripts, and batch files must exist on the same system as the Tivoli Storage Productivity Center graphical user interface (GUI).

Both element managers and external tools are configured on the **Element Manager** tab of the Configuration Utility (**IBM Tivoli Storage Productivity Center → Configuration**. The following tasks can be performed for external tools in the **External Tools** section of the **Element Manager** tab. The tasks that you can perform depend on your Tivoli Storage Productivity Center authorization

- Add external tools (Administrator)
- Start external tools (all authorization levels)
- View external tools (all authorization levels)
- Edit external tools (Administrator)
- Remove external tools (Administrator)

Adding external tools

Use the Configuration Utility to add external tools to IBM Tivoli Storage Productivity Center.

Possibilities for external tools include:

- Executable files, shell scripts, or batch files. These files must exist on the same system as the Tivoli Storage Productivity Center graphical user interface (GUI).
- Uniform Resource Locators (URLs) for search engines or Web encyclopedias

To add an external tool in Tivoli Storage Productivity Center, follow these steps:

1. Expand the **IBM Tivoli Storage Productivity Center** node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **External Tools** section, select **Add Tool** from the **Select Action** list. The Add External Tool dialog box is displayed.
3. Define the information that is required to connect to the element manager in the following fields. Required fields are specified in bold in the dialog box. In the following field descriptions, required fields are marked by an asterisk.
4. The **Add External Tool** window opens. Enter definitions in the following fields:

   ***URL/Command**
   Enter the URL or shell command representing the external tool. You can

enter an IPv4 or IPv6 address depending on what is supported on the machine where the external tool is located:

- Enter an IPv4 address in the following format: ###.###.###.###. For example: 127.0.0.1
- Enter an IPv6 address. 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:DB8:0000:1234:0000:0000:5678: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.

  You must include brackets [] around IPv6 addresses to separate those addresses from their port numbers. For example: **[2001:DB8::1234:0000:0000:5678:ABCD]:9550**.

Note: If the computer on which the element manager is located is IPv6-only, the Tivoli Storage Productivity Center servers must be installed on an IPv6-only or dual stack (IPv4 and IPv6 addresses are configured) computer.

Label Enter the name of the tool. The label name is shown in the **Name** field in the external tools table. If this field is blank, the **Name** field is blank.

Description
Enter a short description for the external tool.

Save Click this button to save the external tool.

Cancel
Click this button to close the dialog box without saving the external tool.

Modifying external tools

Use the Configuration Utility to modify information for external tools in IBM Tivoli Storage Productivity Center.

To modify an external tool in Tivoli Storage Productivity Center, follow these steps:

1. Expand the **IBM Tivoli Storage Productivity Center** node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **External Tools** section, select the tool that you want to modify from the external tools table.
3. Select **Modify Tool** from the **Select Action** list. The Modify External Tool dialog box is displayed.
4. Modify any of the following fields:

   URL/Command
   Enter the URL or shell command representing the external tool. You can enter an IPv4 or IPv6 address depending on what is supported on the machine where the external tool is located:

   - Enter an IPv4 address in the following format: ###.###.###.###. For example: 127.0.0.1
   - Enter an IPv6 address. 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:DB8:0000:1234:0000:0000:5678: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.

You must include brackets [] around IPv6 addresses to separate those addresses from their port numbers. For example: **[2001:DB8::1234:0000:0000:5678:ABCD]:9550**.

Note: If the computer on which the element manager is located is IPv6-only, the Tivoli Storage Productivity Center servers must be installed on an IPv6-only or dual stack (IPv4 and IPv6 addresses are configured) computer.

Label Enter the name of the tool. The label name is shown in the **Name** field in the external tools table. If this field is blank, the **Name** field is blank.

Description
Enter a short description for the external tool.

Save Click this button to save the external tool.

Cancel
Click this button to close the dialog box without saving the external tool.

Starting external tools

Use the Configuration Utility to start external tools from IBM Tivoli Storage Productivity Center.

To start an external tool from Tivoli Storage Productivity Center, follow these steps:

1. Expand the **IBM Tivoli Storage Productivity Center** node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **External Tools** section, select the tool that you want to start from the external tools table.
3. Select **Launch Tool** from the **Select Action** list. The URL or command defined for the tool is run immediately.

Removing external tools

Use the Configuration Utility to remove external tools from IBM Tivoli Storage Productivity Center.

To remove an external tool from Tivoli Storage Productivity Center, follow these steps:

1. Expand the **IBM Tivoli Storage Productivity Center** node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **External Tools** section, select the tool that you want to remove from the external tools table.
3. Select **Remove Tool** from the **Select Action** list. The Remove External Tool dialog box is displayed.
4. Click **OK**. The element manager is deleted from the external tools table.

   To exit the window without removing the external tool, click **Cancel**.

   Note: If the computer on which the element manager is located is IPv6-only, the Tivoli Storage Productivity Center servers must be installed on an IPv6-only or dual stack (IPv4 and IPv6 addresses are configured) computer.

Working with invoices

With IBM Tivoli Storage Productivity Center you can develop invoice definitions and create invoices associated with data or resource usage.

Using Tivoli Storage Productivity Center you can:

- Charge users for the amount of *storage they are consuming* across a network and within relational database management system (RDBS) instances
- Charge users for the amount of *space they own* within the database and tablespace of RDBMS instances.

With Tivoli Storage Productivity Center you can collect usage information on a departmental, group, or user level. You can allocate costs by storage usage by user, disk capacity by computer, tablespace, or file system or physical device. You can create cost centers by creating user, computer, or tablespace groups, allowing organizations to chargeback individuals or business units for their storage usage. By understanding the costs associated with existing and future storage usage, you can improve the use of that storage and reduce the costs associated with its maintenances and upgrades.

Tivoli Storage Productivity Center uses the storage usage information gathered by probes and scans to generate invoices that chargeback for storage usage. Chargeback is the process of breaking down costs associated with data or resource usage and then distributing those costs back to the organizational consumers.

To collect the information needed for chargeback, you need to run probes and scans:

Probes
Define probes when you want to charge by Disk Capacity. Probes collect statistics on the storage assets in your enterprise, such as computers, disk controllers, hard disks, file systems, storage subsystems, instances, logs, RDBMS objects, and so on. The results of probes are stored in the database repository and are used to supply the data necessary for Tivoli Storage Productivity Center's reporting function.

Scans Define scans when you want to charge by OS User groups. Scans collect statistics on the usage and trending of your actual storage consumption. The results of scans are stored in the database repository and are used to supply the data necessary for Tivoli Storage Productivity Center's Capacity, Usage, Usage Violations, and Backup Reporting facilities.

In addition to providing basic storage usage invoicing, Tivoli Storage Productivity Center integrates with the chargeback systems already implemented in your environment. It provides you with a higher level, application-specific CIMS output format. CIMS is IBM Tivoli Usage and Accounting Manager (ITUAM) which provides IT financial knowledge by measuring, analyzing, reporting, and billing based on usage and costs of shared computing resources. With Tivoli Storage Productivity Center, you can export chargeback data directly into ITUAM, giving you the ability to integrate your storage usage data with other enterprise chargeback information and processes.

Planning for chargeback

Before generating data that will be used for chargeback, it is a good idea to plan ahead and gather basic information about your system and organization.

Understand your environment

Understanding the users, data, and resources within your organization will help you determine the most reasonable and enforceable chargeback policies. In the planning stages of developing a chargeback policy, it will help you to understand such things as:

- The types of data stored within your environment
- Patterns of data usage
- Different types of users and user groups
- The IT resources that exist within your organization

Determine the chargeback policy you want to implement

It is important to learn and understand the advantages and disadvantages of the charge back policy you want to implement. For example, with IBM Tivoli Storage Productivity Center, you can:

- Divide the cost of storage usage by all users
- Divide the cost of storage usage by individual users or groups of users
- Charge by resource
- Charge by individual usage of resources

Define monitoring groups

Organize the IT resources within your environment into functional groups. Defining groups enables you to classify your data and resources into billable entities, which in turn provides the information on who is using what storage. This enables you to bill specific groups according to their usage of storage space. See "Creating a group" on page 237.

Computer groups

Computer groups collect operating system and RDBMS storage statistics on computers that are related to each other or belong to a common group such as a department, a location, a territory, and so forth. Examples of computer groups are Sales, Accounting, Development, IT.

A computer can belong to one computer group only.

User groups

User groups are users that are related to each other or belong to a common group such as a department, a location, a territory, and so forth. Examples of user groups are Administrators, Developers, Marketing staff, HR staff.

A user can belong to one user group only.

OS User groups

OS user groups are users that are grouped together within the operating system.

Database-tablespace groups

Database-tablespace groups monitor entire databases or tablespaces across multiple instances or databases. Examples of tablespace groups are System Tablespaces, System Databases, databases-tablespaces related to a specific application or applications, and so forth.

A database-tablespace can belong to one database-tablespace group only.

Determine what information you want to gather

Determining what information you want to gather about the storage

resources within your enterprise is critical to helping you implement a chargeback strategy. Probes and scans are the monitoring jobs you can define to collect information about operating system and RDBMS storage resources. The information gathered by these jobs is reflected in the data included within a Tivoli Storage Productivity Center invoice.

As you become more familiar with how a specific charge back model works for your organization, you can continue to review and implement to more advanced charge back models.

Setting up invoices

This topic takes you through the task of setting up an invoice.

To set up an invoice, complete the following steps:

1. In the Navigation Tree pane, expand **Data Manager for Chargeback** → **Parameter Definition**. Data Manager for Chargeback - Configuration window opens.
2. In the **Output Format** area, specify the format for the output. You can click either of the following options:

   Invoice
   Creates a formatted text file suitable for opening in or importing into word processing applications.

   CIMS Creates a text file with comma-separated fields suitable for importing into a ITUAM application.
3. In the **Type of Charges** area, specify the information that you want the invoice to contain. You can select one or more of the following check boxes:

   OS Storage Usage By User
   The resulting invoice includes the following information:
   - Total space consumed by user groups and the cost of that usage
   - Total space consumed by each user in the environment and the cost charged to each user

   OS Disk Capacity By Computer
   The resulting invoice includes the following information:
   - Total space consumed by computer groups and the cost of that usage
   - Total space consumed on each computer in the environment and the cost of that usage

   Storage Usage By Database User
   The resulting invoice includes the following information:
   - Total space consumed by DB user groups and the cost of that usage
   - Total space consumed by each DB user in the environment and the cost charged to each user

   Total Size By Database-Tablespace
   The resulting invoice includes the following information:
   - Total space consumed by database-tablespace groups and the cost of that usage
   - Total space consumed on each database or tablespace in the environment and the cost of that usage
4. (Invoice output format only) For each type of charge that you selected in step 3, select the unit of measurement (kilobyte, megabyte, or gigabyte) and type a

numerical value. For example, if you want to charge $1.25 per megabyte of storage usage, select **Megabyte** and type 1.25 in the **Cost per** field.

(CIMS output format only) If you select CIMS as the output type, the **Cost Per** field will be deselected. This is because you will set the rates and costing within the ITUAM application.

5. If you selected **OS Storage Usage by User** or **Storage Usage by Database User**, select the profile that is used in the scan job. By default, this is either TPCUser.Summary By Owner for **OS Storage by User** or TPCUser.Db User Space for **Storage Usage by Database User**.
6. Click  to save the invoice definition.

Invoice definitions are not saved in actual files, but the last invoice definition saved is the one that will be used to create the invoice.

Generating an invoice

Generate an invoice and save it as a file.

To generate an invoice, complete the following steps:

1. In the Navigation Tree pane, expand the **Data Manager for Chargeback** node.
2. Click **Online Execution**. The Online Execution page opens.
3. Click the **Browse** button. The Open window opens, listing previously generated invoice files in the default directory. The default directory, or the directory where you have chosen to save the invoice files, is not part of the IBM Tivoli Storage Productivity Center application but on the computer where Tivoli Storage Productivity Center is installed.
   a. Click on an existing file or type a new file name, including extension, in the **File name** text field. A text file (.txt) is the default file type for both output types.

      Note: Unless you save the invoice as a separate file name, the default file (StorageChargeback.txt) will be overwritten each time you run an invoice. It is recommended that you save your invoices as a separate file name each time you run them.
   b. Click the **Open** button. You will be returned to the Online Execution page.

   Alternatively, in the Online Execution page, without clicking the **Browse** button, type the full path name of the invoice file to be generated in the **Output File** text field.
4. Click the **Generate Output File** button to generate the invoice.

The invoice is saved to the file name and path you specified.

Viewing and printing invoices

This topic discusses viewing or printing a generated invoice in both Invoice and CIMS output formats.

If you have generated the invoice in the Invoice output format, you can view and print that invoice with any text editor or word processing program. If you chose CIMS as the output format, use your ITUAM application to work with the invoice data.

To view and print an invoice generated in the Invoice output format, complete the following steps:

1. Start the text editor or word processing program.
2. Follow the steps in the word processing program to open the invoice. Open the invoice from the location you specified in the **Output File** field in the Online Execution page.
3. View the file and make any changes. If you make any changes to the file, be sure and save it under a different name.
4. Print the file from the word processing program.

To view and print an invoice generated in the CIMS output format, integrate IBM Tivoli Storage Productivity Center accounting data into the Tivoli Usage and Accounting Manager, by following the instructions in the Tivoli Usage and Accounting Manager Information Center. Alternatively, you can view the original invoice data in a text editor or word processing program. Remember that the data will be in the form of comma-separated fields.

Sample Invoices

This section gives examples of sample invoices.

OS Storage Usage By User

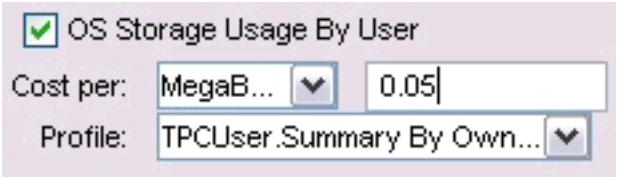

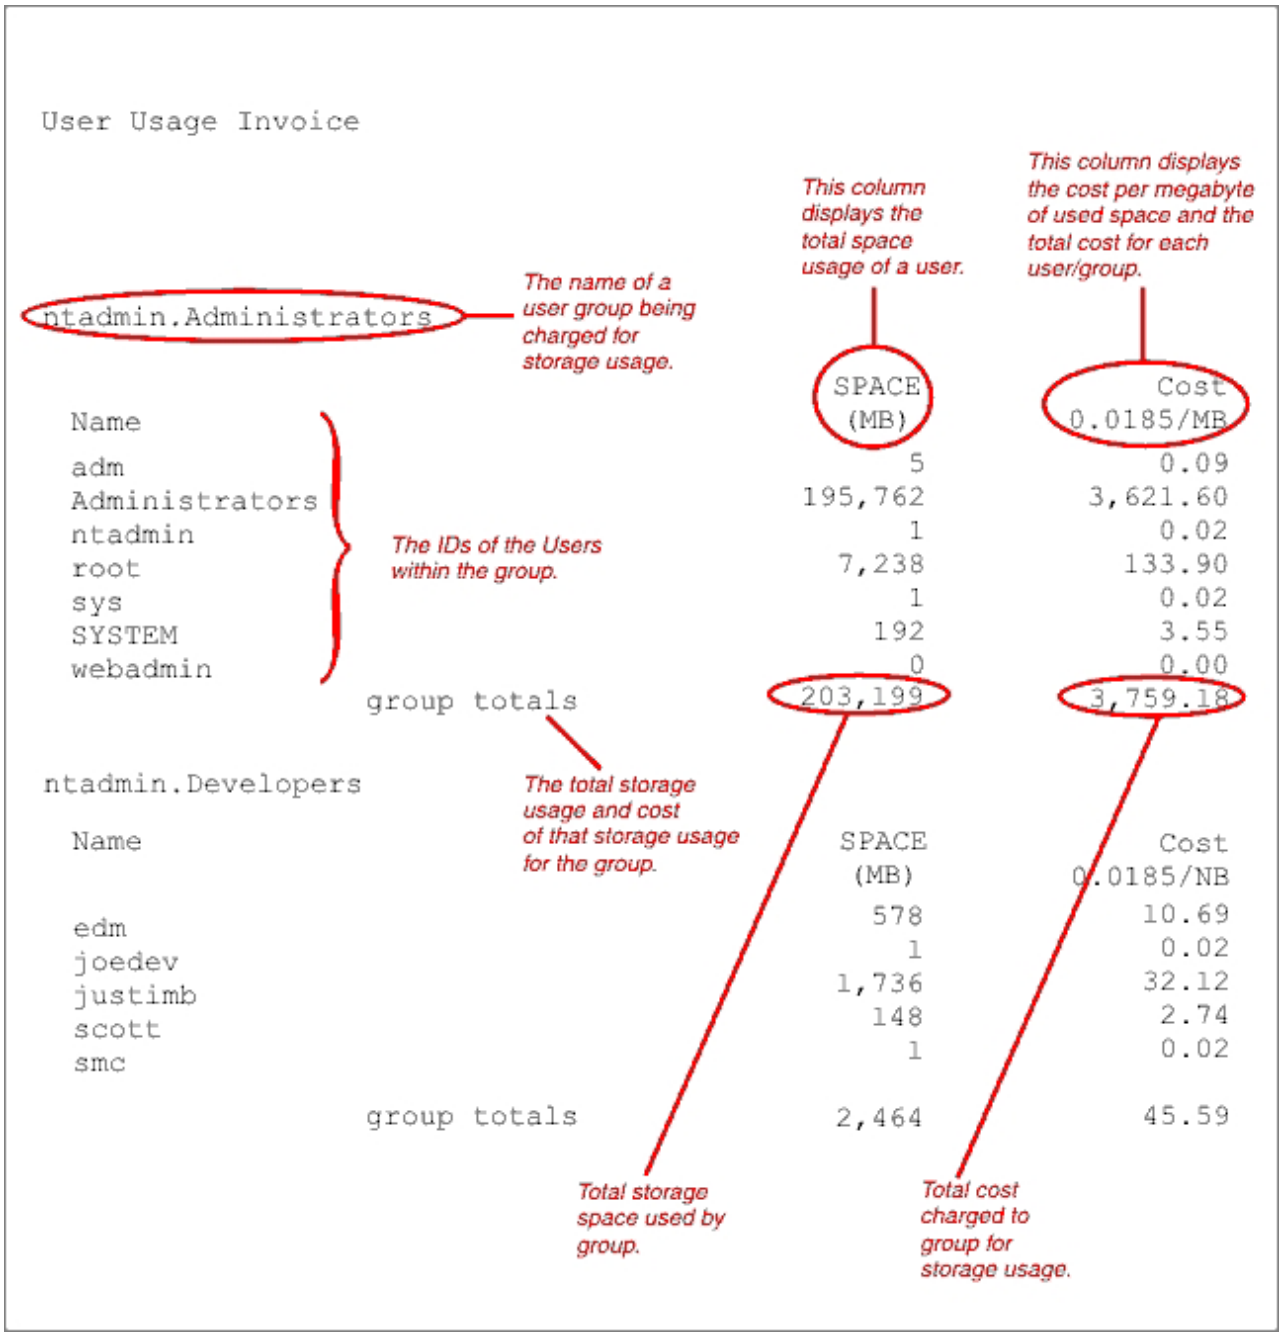

Figure 44. Sample Invoice — OS Storage Usage By User

OS Disk Capacity By Computer

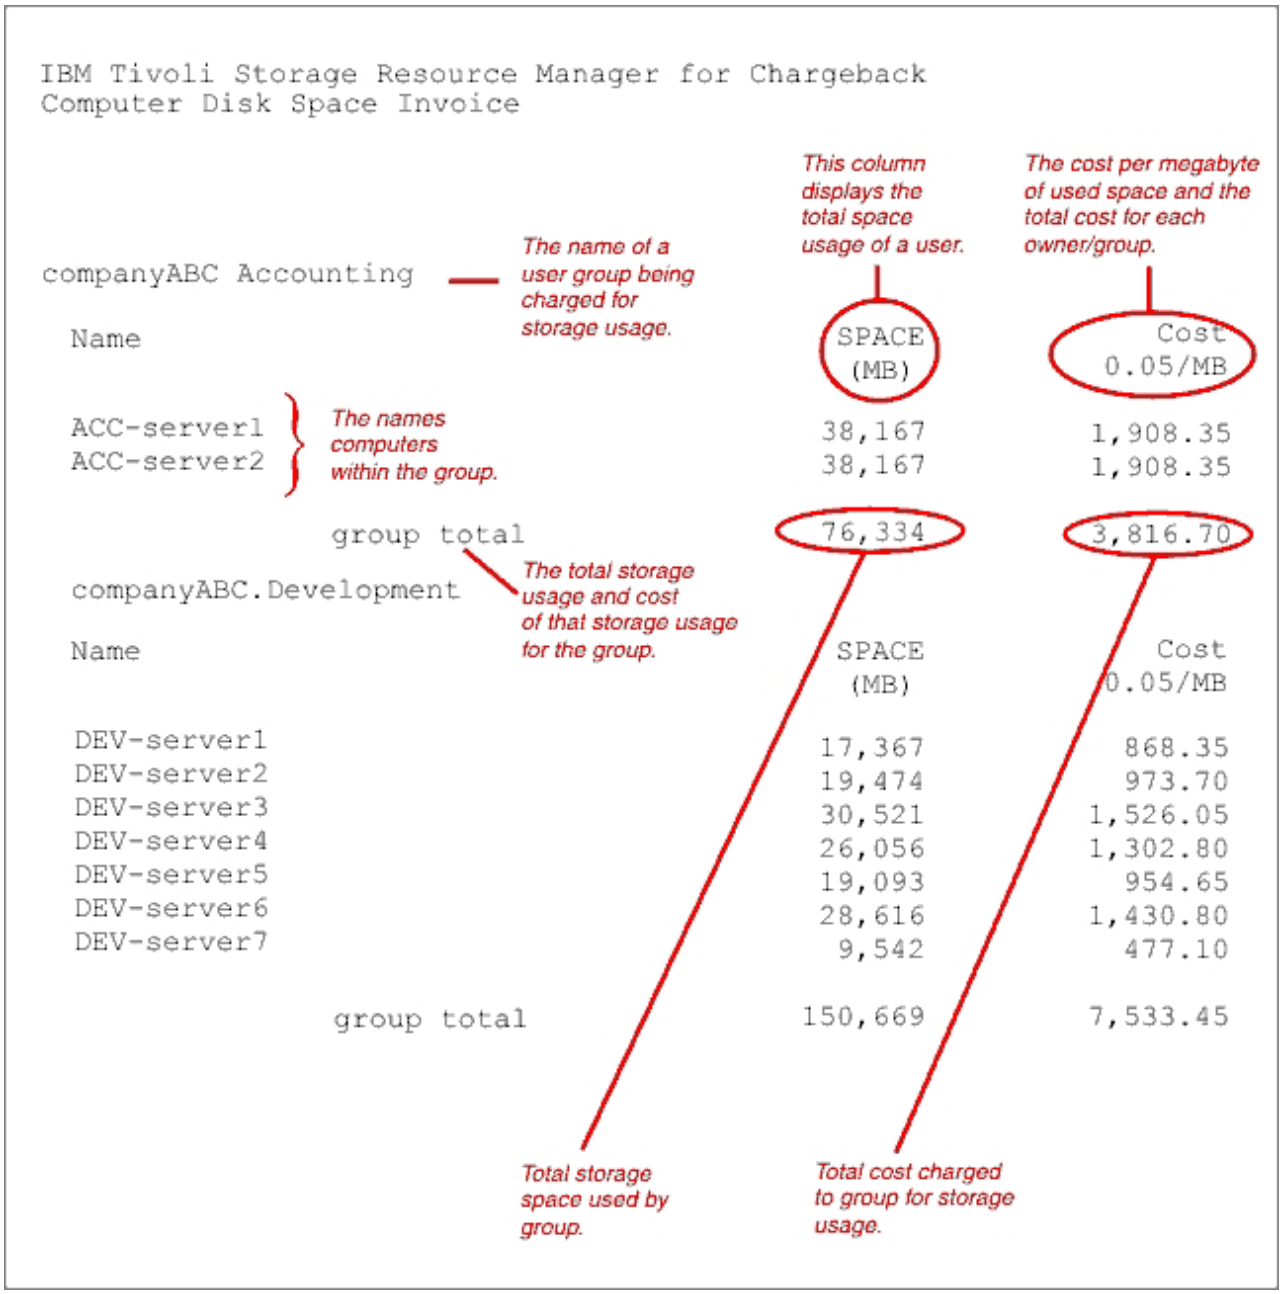

Figure 45. Sample Invoice — OS Disk Capacity by Computer

Storage Usage By Database User

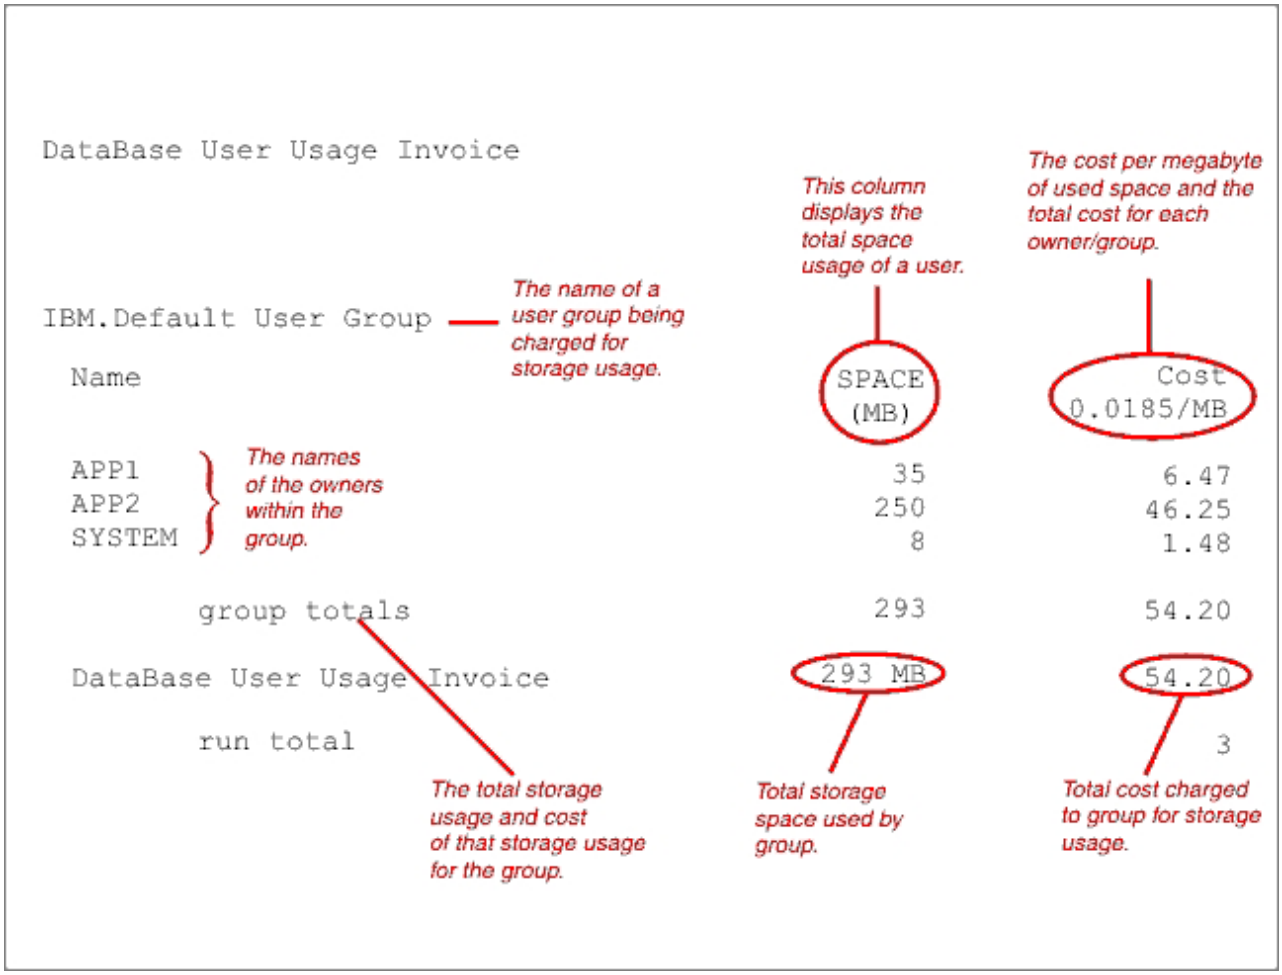

Figure 46. Sample Invoice — Storage Usage by Database User

Total Size By Database -Tablespace

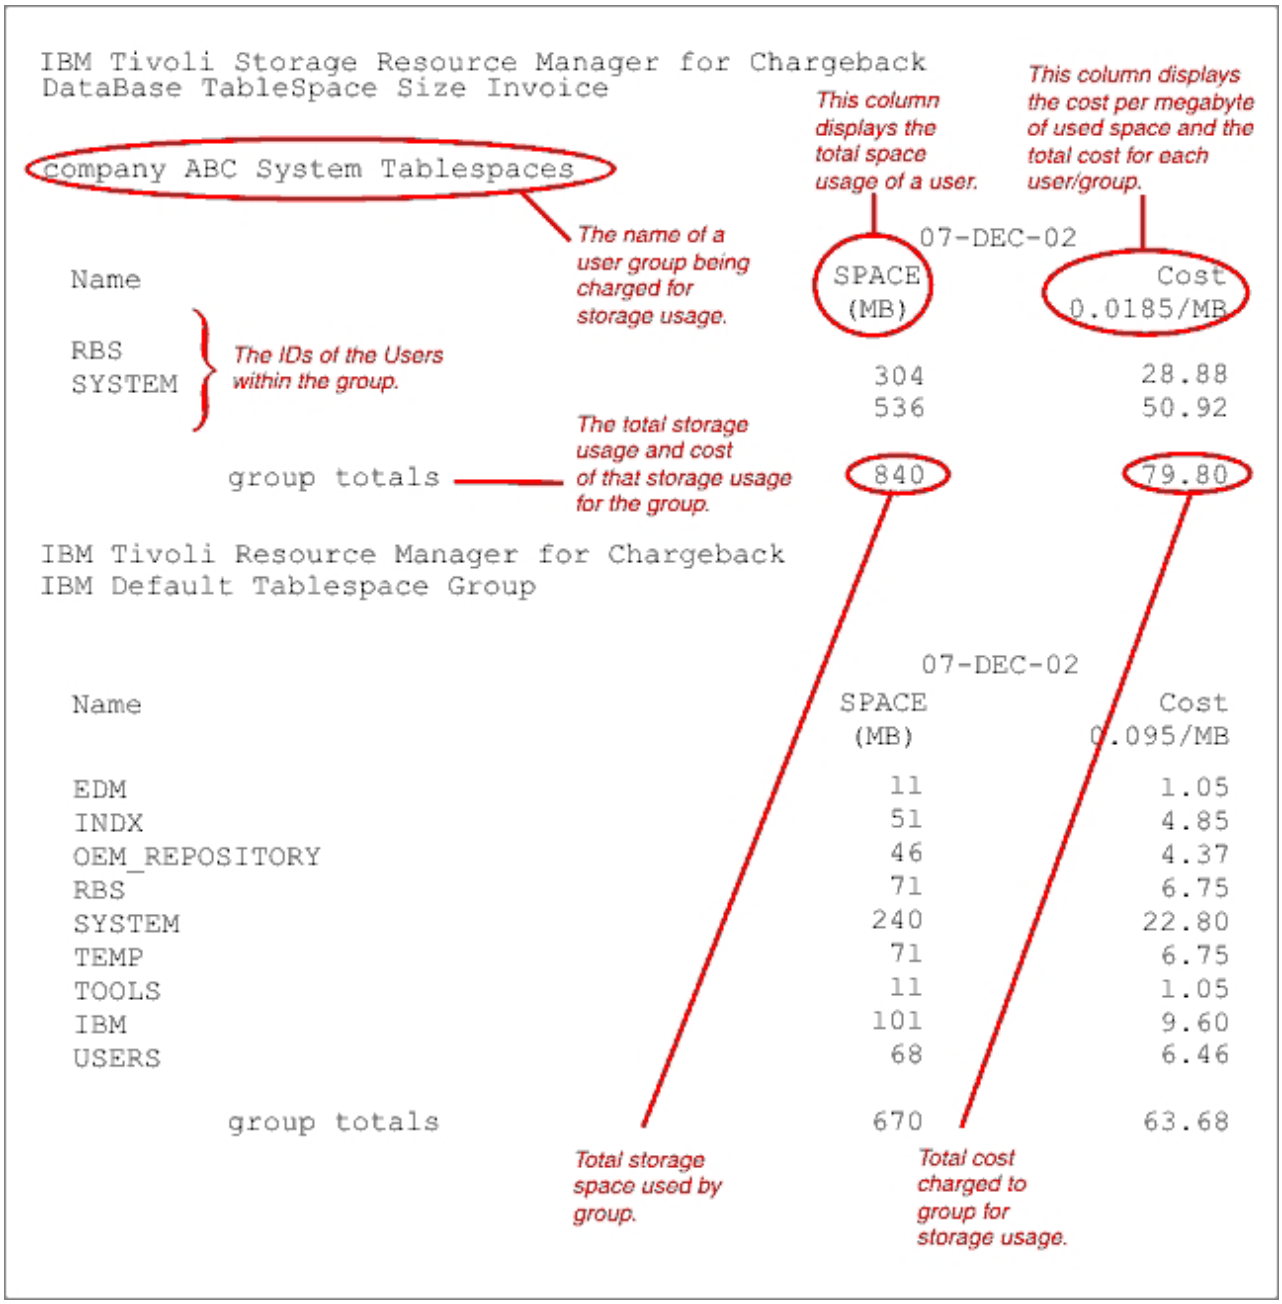

Figure 47. Sample Invoice — Total Size by Database-Tablespace

Chapter 5. Reporting

IBM Tivoli Storage Productivity Center provides over 300 reports for viewing information about storage resources. These reports are conveniently organized into different types and categories and provide both summary and detailed information about the storage resources in your environment.

Data that appears in reports is based on statistical information accumulated by data collection jobs and stored in the database repository. You can select the

specific systems, storage subsystems, switches, clusters, disks, filesystems, and even users upon which you want to report.

Reports are organized into report types and categories, which can be expanded in the navigation tree to reveal additional sub-categories of reporting. Some reports are also accessible through the topology viewer.

- **Reporting types:** system reports, batch reports, user ID reports, user-defined online reports
- **Report categories:** asset, availability, capacity, usage, usage violation, backup, storage subsystem performance, switch performance

Consider the following steps when determining how to best use Tivoli Storage Productivity Center reports:

1. Determine how you want to generate and view reports. See "Choosing a reporting type" for more information.
2. Identify which reports contain the most important information about your storage environment. See "Choosing a report category" on page 466 for more information.

Choosing a reporting type

Use pre-defined (system reports), saved reports, batch reports, rollup reports, and user-defined reports, tables, and charts to view the information collected by monitoring and quota jobs. IBM Tivoli Storage Productivity Center includes these different reporting types to provide you with the flexibility to view data about your storage resources according to the needs of your environment.

The following table describes the reporting types available in Tivoli Storage Productivity Center. Use this table to help determine how to best deliver and view reports about your storage resources.

Reporting type	Description	Location in the navigation tree
User-defined online reports	These reports can be generated at any time during a product session and are only available for that session. Each manager within Tivoli Storage Productivity Center has a set of reports specific to that manager. For example, to view detailed reports about storage subsystems, access the reporting function within Disk Manager. See the following topics for more information about online reports: • "User-defined online reports" on page 485 • "Data Manager reports" on page 496 • "Using Data Manager for Databases reports" on page 613 • "Disk Manager reports" on page 642 • "Fabric Manager reports" on page 653	Data Manager > Reporting Data Manager for Databases > Reporting Disk Manager > Reporting Fabric Manager > Reporting

Reporting type	Description	Location in the navigation tree
System reports	System reports are automatically generated by Tivoli Storage Productivity Center and contain some of the most common information about your monitored storage resources. These pre-defined reports enable you to view your storage enterprise data quickly and efficiently with one click. Data for these system reports are gathered every time that monitoring jobs are run against your storage resources. See the following topics for more information about online reports: • "System reports - Data" on page 661 • "System reports - Data Manager for Databases" on page 673 • "System Reports - Disk Manager" on page 679 • "System reports - Fabric Manager" on page 694	IBM Tivoli Storage Productivity Center > Reporting > System Reports
User ID (saved) reports	Create saved reports when you want specific reports generated for you whenever you log in to a product session. This enables you to save reports that you can generate at any time without having to define your report criteria again. Saved reports are only visible to the user ID that created them. See "User ID (saved) reports" on page 710 for more information.	IBM Tivoli Storage Productivity Center > Reporting > My Reports > *user ID* Reports where *user ID* represents the IBM Tivoli Storage Productivity Center ID of the user that saved the report
Batch reports	A batch report represents any Tivoli Storage Productivity Center report that you want to run on a regular basis and save its data to a file. You can view the information in a generated batch report file directly or import it into an application of your choice. To create a batch report, you must determine which Tivoli Storage Productivity Center report you want to generate, the schedule on which to generate the report, and the file to which to save its data. You can save the data generated for a batch report to any of the following file formats: CSV, HTML, formatted text, PDF chart, and HTML chart. See "Batch reports" on page 699 for more information.	IBM Tivoli Storage Productivity Center > Reporting > My Reports > Batch Reports

Reporting type	Description	Location in the navigation tree
Rollup Reports	Use rollup reports to view the storage information collected by the master Tivoli Storage Productivity Center server from its associated subordinate servers. Rollup reports enable you to have a network-wide perspective of storage usage in an environment where multiple Tivoli Storage Productivity Center servers are deployed and collecting storage metrics. See the following topics for more information about online reports: • "Rollup reports" on page 711 • "IBM Tivoli Storage Productivity Center Servers" on page 148 • "Collecting storage statistics from IBM Tivoli Storage Productivity Center servers (TPC server probes)" on page 207	IBM Tivoli Storage Productivity Center > Reporting > Rollup Reports
Data Source reports	Use these reports to view detailed information about the agents from which Tivoli Storage Productivity Center collects information about storage resources. Data sources include CIMOM agents, Data agents, Fabric agents, VMWare, and subordinate Tivoli Storage Productivity Center servers. See Data Source reports for more information.	IBM Tivoli Storage Productivity Center > Reporting > Data Source Reports

Choosing a report category

IBM Tivoli Storage Productivity Center provides a number of categories for reports that help you identify which reports contain the information you want to view about your storage environment. Each category of report contains information gathered by different data collection and quota jobs.

The following table describes the reporting categories. Use this table to help determine which reports are most useful to you and which data collection jobs to run for those reports.

Report Category	Description	What data collection job should I run for this report?
Asset reports	View itemized statistics about your storage resources in a hierarchical view that enables you to drill down to view those resources in greater detail. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > Asset • Data Manager for Databases > Reporting > Asset • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) • Data Manager for Databases > My Reports **Example:** View information about a storage subsystem and drill down into its related storage pools, disks, volumes, and disk groups. See "Asset reports" on page 499 for more information about Asset reports.	Probes, Discovery Data for System-wide Users and Monitored Directory reports are gathered by Scans.
Availability reports	View availability statistics about your storage resources, both from a network point of view and from a computer uptime perspective. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > Availability • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) **Example:** Use a Computer Uptime report to determine the percentage of time that a computer was up, the number of times it was rebooted, and the total time it was down within a user-defined period of time. See "Availability Reports" on page 515 for more information about Availability reports.	Ping data: Pings Computer Uptime: Probes

Report Category	Description	What data collection job should I run for this report?
TPC-wide Storage Space	Examine storage capacities of storage resources at a file system, computer, or network level. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > TPC-wide Storage Space • Data Manager for Databases> Reporting > Capacity • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) • Data Manager for Databases > My Reports **Example:** Identify how much storage space is being used on the file systems within your environment, and how much unused storage space is available. See "TPC-wide Storage Space Reports" on page 517 for more information about Capacity reports.	Probes
Usage reports	View information about the usage and growth of consumption of your storage resources, as well as select files against which to run an Tivoli Storage Productivity Center archive/backup job. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > Usage • Data Manager for Databases > Reporting > Usage • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) • Data Manager for Databases > My Reports (if you save the definition of an online report it appears in this location) **Example:** Determine how much of the storage being used in your environment can be considered wasted space (orphaned, obsolete, or forbidden files). See "Usage Reports" on page 573 for more information about Usage reports.	Scans

Report Category	Description	What data collection job should I run for this report?
Usage Violation reports	View quota and constraint violations as well as select files for archive/backup jobs. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > Usage Violations • Data Manager for Databases > Reporting > Usage Violations • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) • Data Manager for Databases > My Reports (if you save the definition of an online report it appears in this location) **Example:** Identify the users or groups of users that are using storage space on machines where they should not. See "Usage Violation reports" on page 581 for more information about Usage Violation reports.	Scans, Quotas, Constraints
Backup reports	Investigate the sizes of data that would be required to backup to achieve full backups, determine the amount of space required to perform an incremental backup of any or all of your data at any given point in time, and select files for archive/backup jobs. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > Backup • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) • Data Manager for Databases > My Reports (if you save the definition of an online report it appears in this location) **Example:** Determine if you have enough storage space to perform a backup of your data (incremental or full). See "Backup reports" on page 583 for more information about Usage Violation reports.	Scans

Report Category	Description	What data collection job should I run for this report?
Monitored Computer Storage Space	Use these host-based reports to view the data that is collected by probes for computers, Network Attached Storage, and hypervisors. These reports contain detailed information about the consumption of file system space and enable you to relate that space back to external or internal storage without counting shared space multiple times. **Location:** These reports are available in the following locations on the navigation tree: • Data Manager > Reporting > Monitored Computer Storage Space • IBM Tivoli Storage Productivity Center > Reporting > My Reports (if you save the definition of an online report it appears in this location) See "Monitored Computer Storage Space reports" on page 584 for more information about Monitored Computer Storage Space reports.	Probes
Storage subsystem reports	View storage capacity for a computer, file system, storage subsystem, volumes, array sites, ranks, and disks. These reports also enable you to view the relationships among the components of a storage subsystem. **Location:** Disk Manager > Reporting > Storage Subsystems **Example:** Determine if you have enough storage space to perform a backup of your data (incremental or full). See "Disk Manager reports" on page 642 for more information about Storage Subsystem reports.	Probes
Storage Subsystem Performance reports	View performance data for storage subsystems. **Location:** Disk Manager > Reporting > Storage Subsystem Performance **Example:** Check the performance rate for storage subsystem activities, such as read I/O rate, read data rate, total port I/O rate, etc. See "Disk Manager reports" on page 642 for more information about Storage Subsystem Performance Monitor reports.	Subsystem Performance Monitors

Report Category	Description	What data collection job should I run for this report?
Switch Performance reports	View performance data for fabric switches. **Location:** Fabric Manager > Reporting > Switch Performance **Example:** Check the performance rate for switch activities, such as port send packet rate, port send data rate, total port data rate, etc. See "Fabric Manager reports" on page 653 for more information about Switch Performance Monitor reports.	Switch Performance Monitors
Rollup reports: Asset, Capacity, Database Asset, Database Capacity	View the storage information that has been collected by IBM Tivoli Storage Productivity Center servers through a single interface. These rollup reports enable you to have a consolidated perspective of storage usage in an environment where multiple IBM Tivoli Storage Productivity Center servers are deployed and collecting storage metrics. **Location:** IBM Tivoli Storage Productivity Center > Reporting > Rollup Reports **Example:** Check the disk capacity of all the computers monitored by different servers in your environment. See "Rollup reports" on page 711 for more information about rollup reports.	IBM Tivoli Storage Productivity Center Server Probes
Data Source reports	Use data source reports to view detailed information about the agents from which Tivoli Storage Productivity Center collects information about storage resources. Data sources include CIMOM agents, Data agents, Fabric agents, VMWare, and subordinate Tivoli Storage Productivity Center servers. **Location:** IBM Tivoli Storage Productivity Center > Reporting > Data Source Reports **Example:** View detailed information about the CIMOM Agents that have been added to Tivoli Storage Productivity Center as data sources. See "Data source reports - overview" on page 766 for more information about data source reports.	These reports are available for data sources that have been associated with Tivoli Storage Productivity Center.

What can I find out using reports?

Learn about some of the common storage questions that you can answer using IBM Tivoli Storage Productivity Center reports.

The following table lists some of the common questions you might have while managing your storage resources and shows the reports and data collection jobs you can run to answer those questions.

What do you want to know?	Which reports answer this question?	What job should I run for this report?
What hardware is currently being used within our organization?	Data Manager > Asset > By <storage_resource_type>	Probes
How much storage space are we using?	Data Manager > Capacity > Filesystem Used Space	Probes
Of the space that is being used, how much is wasted space? (orphaned, obsolete, or forbidden files)	Data Manager > Usage > Files Data Manager > Usage Violations > Constraint Violations	Scans
How much unused or unallocated storage space do we have?	Data Manager > Capacity > Filesystem Free Space	Probes
Do we have enough space to do a backup of our data? (incremental or full)	Data Manager Backup > Backup Storage Requirements	Scans
What files in the environment are at risk if a disaster hits?	Data Manager > Backup > Most at Risk Files	Scans
What ″forbidden files″ are taking up storage space on our network?	Data Manager > Usage Violations > Constraint Violations	Scans
What user or group of users is using space on machines they shouldn't?	Data Manager > Usage Violations > Quota Violations	Scans
How fast are our storage resources being consumed?	Data Manager > Capacity > Filesystem Used Space	Probes
What user or business unit is consuming the most storage resources?	Data Manager > Usage > Users Data Manager > Usage > OS User Groups	Scans
Where is disk space reaching a critical shortage?	Data Manager > Capacity > Disk Capacity	Probes
What machines have been experiencing the most use/traffic?	Data Manager > Usage > Access Load	Scans
Based on historical consumption, how much storage space should be added and when should it be added?	Data Manager > Capacity > Filesystem Free Space	Probes
What machines have been experiencing the least use/traffic?	Data Manager > Usage > Access Load	Scans
What files have not been used within a certain period of time?	Data Manager > Usage > Files > Most Obsolete Files	Scans
What are the largest files or directories within the environment?	Data Manager > Usage > Files > Largest Files	Scans
What are the most modified files within the environment?	Data Manager > Usage > Access Load > Modification Time	Scans
What computers are "up" or "down" within my environment?	Data Manager > Availability > Pings	Pings

What do you want to know?	Which reports answer this question?	What job should I run for this report?
What hardware is most prone to error? (i.e., disk defects)	Data Manager > Asset > System-wide > Disks	Probes
What file owners no longer exist in the environment?	Data Manager > Usage > Files > Orphan files	Scans
How much storage space is an Instance consuming?	• **Oracle:** Data Manager for Databases > Reporting > Capacity > Oracle > Total Instance Storage • **SQL Server:** Data Manager for Databases > Reporting > Capacity > SQL Server > Total Instance Storage • **Sybase:** Data Manager for Databases > Reporting > Capacity > Sybase > Total Instance Storage • **UDB:** Data Manager for Databases > Reporting > Capacity > UDB > Total Instance Storage	Probes
What user or business unit is consuming the most storage resources and on what instances are they consuming those resources?	**Oracle:** • Data Manager for Databases > Reporting > Usage > Oracle > Users > Group by User Name, *or* • Data Manager for Databases > System Reports > Data > Oracle - User Database Space Usage **SQL Server:** • Data Manager for Databases > Reporting > Usage > SQL Server > Users > Group by User Name, *or* • Data Manager for Databases > System Reports > Data > SQL Server - User Database Space Usage **Sybase:** • Data Manager for Databases > Reporting > Usage > Sybase > Users > Group by User Name, *or* • Data Manager for Databases > System Reports > Sybase - User Database Space Usage **UDB:** • Data Manager for Databases > Reporting > Usage > UDB > Users > Group By User Name, *or* • Data Manager for Databases > System Reports > UDB - User Database Space Usage	Scans
What tablespace is reaching a critical shortage of free space?	• **Oracle:** Data Manager for Databases > Reporting > Capacity > Oracle > Data File Free Space > By Tablespace • **UDB:** Data Manager for Databases > Reporting > Capacity > UDB > Container Free Space > By Tablespace	Probes

What do you want to know?	Which reports answer this question?	What job should I run for this report?
What tablespaces contain the most free space?	• **Oracle:** Data Manager for Databases > Reporting > Capacity > Oracle > Data File Used Space > By Tablespace • **UDB:** Data Manager for Databases > Reporting > Capacity > UDB > Container Used Space > By Tablespace	Probes
What database is reaching a critical shortage of free space?	• **SQL Server:** Data Manager for Databases > Reporting > Capacity > SQL Server > Data File Free Space > By Database • **Sybase:** Data Manager for Databases > Reporting > Capacity > Sybase > Fragment Free Space > By Database	Probes
What database contains the most used space?	• **SQL Server:** Data Manager for Databases > Reporting > Capacity > SQL Server > Data File Used Space > By Database • **Sybase:** Data Manager for Databases > Reporting > Capacity > Sybase > Fragment Used Space > By Database	Probes
Based on historical consumption of storage space for a RDBMS, how much storage space should be added and when should it be added?	• **Oracle:** Data Manager for Databases > Reporting > Capacity > Oracle > Data File Free Space > By Database > History Chart:Free Space % for selected • **SQL Server:** Data Manager for Databases > Reporting > Capacity > SQL Server > Data File Free Space > By Instance > History Chart:Free Space % for selected • **Sybase:** Data Manager for Databases > Reporting > Capacity > Sybase > Device Free Space > By Device > History Chart:Free Space % for selected • **UDB:** Data Manager for Databases > Reporting > Capacity > UDB > Container Free Space > By Instance > History Chart:Free Space % for selected	Probes
What are the largest tables within the environment?	• **Oracle:** Data Manager for Databases > Reporting > Usage > Oracle > Tables > Largest Tables • **SQL Server:** Data Manager for Databases > Reporting > Usage > SQL Server > Tables > Largest Tables • **Sybase:** Data Manager for Databases > Reporting > Usage > Sybase > Tables > Largest Tables • **UDB:** Data Manager for Databases > Reporting > Usage > UDB > Tables > Largest Tables	Scans
What segments in an Instance contain the most extents?	**Oracle:** Data Manager for Databases > Reporting > Usage > Oracle > Segments > Most Extents	Scans

What do you want to know?	Which reports answer this question?	What job should I run for this report?
How much allocated space for an object is being wasted?	**Oracle:** My Reports > System reports > Oracle - Segments with Wasted Space	Scans
How many free extents does a tablespace have?	**Oracle:** Data Manager for Databases > Reporting > Asset > By DBMS Type > Oracle > instance_name > By Tablespace	Probes
How many tables, indexes, and datafiles/containers are associated with a tablespace?	• **Oracle:** Data Manager for Databases > Reporting > Asset > System-wide > Oracle > Tablespaces • **UDB:** Data Manager for Databases > Reporting > Asset > System-wide > UDB > Tablespaces	Scans
How many tables, indexes, and datafiles/fragments are associated with a database?	• **SQL Server:** Data Manager for Databases > Reporting > Asset > System-wide > SQL Server > Databases • **Sybase:** Data Manager for Databases > Reporting > Asset > System-wide > Sybase > Databases	Scans
How can I determine which users have violated a Quota on a specific Instance?	**All DBMSs:** Data Manager for Databases> Reporting > Quota Violations > Instance Quota > By Instance	Scans
What are the sizes and locations of Control Files and log files for an Instance within my environment?	• **Oracle:** Data Manager for Databases > Reporting > Asset > By DBMS Type > Oracle > instance_name • **SQL Server:** Data Manager for Databases > Reporting > Asset > By DBMS Type > SQL Server > instance_name > Databases	Probes
How can I determine how close a user is to exceeding a quota?	**All DBMSs:** Data Manager for Databases > Usage Violations > Quota Violations > All Quotas > By User	Scans

Report Considerations and Limitations

Learn about some of the considerations you should know before using IBM Tivoli Storage Productivity Center reports.

Data collection for reports

Some columns within reports might not be populated with data for the storage resources that you are monitoring. Columns appear empty if you have not configured Tivoli Storage Productivity Center to discover or collect the information intended for that column. Use discovery, probe, ping, or scan jobs to ensure that you collect information for all the columns that appear in reports.

Storage totals for Network-wide reports

Tivoli Storage Productivity Center provides a number of network-wide reports that display storage totals based on the data collected from the monitored computers and storage subsystems in your network. You can click the magnifying glass icon next to the row in a network-wide report to ″drill down″ and view the details about the computers that are included in those storage totals. However, the reports that you drill down into display information about computers only. Storage information about storage

subsystems is not included in the values for reports that are shown when you drill down from a network-wide report.

For example, the Data Manager > Reporting > Capacity > Disk Capacity > Network-wide report displays disk capacity totals for the computers and storage subsystems in your network that are monitored by Tivoli Storage Productivity Center. However, when you click the magnifying glass icon next to a row in that report, the resulting report displays disk totals for computers only. This means that the totals displayed in the network-wide report might appear larger than those reflected in the report that you drill down into if you have disks that are part of both computers and storage subsystems within your network.

IBM XIV Storage System reporting limitation for disk drives

Tivoli Storage Productivity Center does not collect physical serial number information about the disk drives that are associated with XIV Storage System devices. Therefore, any reports that display information about XIV Storage System do not include the physical serial numbers of the disk drives.

Hitachi Data Systems storage subsystems limitations

- There is a limitation within reports when Tivoli Storage Productivity Center collects information about Hitachi Data Systems (HDS) storage subsystems that are configured as storage virtualizers. Specifically, any reports that display a value for Disk Unallocated Space do not count the storage consumed by volumes residing on external disks when the HDS is used as an virtualizer.
- Tivoli Storage Productivity Center does not support Logical Unit Size Expansion (LUSE) with Hitachi storage subsystems. Because of this, Tivoli Storage Productivity Center might report a different number of volumes for a subsystem than is reflected in the HiCommand Device Manager (HDvM) when LUSE is used.

Time zones for computers monitored by Data agents and 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 Data Manager > Reporting > Assets > By Computer report:

- 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

SQL access to data collected by Tivoli Storage Productivity Center for use in external reports and applications

You can use the data collected by data collection jobs to create custom reports in external tools. Tivoli Storage Productivity Center provides a set of DB2® views that represent key data that has been stored in the database repository. You can use Structured Query Language (SQL) commands to retrieve the information in these views for use in external tools such as Business Intelligence and Reporting Tools (BIRT), Microsoft Excel, and Tivoli Common Reporting. See *Planning to retrieve information from views in the database repository* topic in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and*

Configuration Guide for more information about how to use SQL to retrieve data for use in custom reports and tools.

Working with reports

Learn how to drill down through reports, navigate report tabs, view historical trends in reports, and view reports as tables of data or as graphical charts. The ability to customize your view of reports provides you with the flexibility to get a very detailed look at your information or a very high level overview.

Generating and viewing reports

Learn how to generate and view reports.

Each of the report types available in IBM Tivoli Storage Productivity Center are generated in different ways to provide you with the flexibility to meet your environment's changing needs. Click the following links to learn more about how to generate the different report types and determine what works best for your organization.

Report type	How to generate
User-defined online reports	"Generating reports" on page 485
System reports	"Generating and viewing system reports" on page 660
User ID (saved) reports	"Creating a user-defined report" on page 711
Batch reports	"Creating batch reports" on page 700
Rollup reports	"Generating rollup reports" on page 712

Sorting report columns

Learn how to sort the columns in reports.

Reports appear in a tabular format and are comprised of rows and columns. You can scroll the report up and down and to the right and left to view the entire report. Use the **View** menu to hide/show the navigation tree to increase the viewable area of the report. You can also drag the divider bar in the middle of the screen back and forth to reallocate the amount of space that is given each pane.

To sort the columns or change the order of the columns in a report:

- **Re-sorting columns.** To resort a column, click on the column header. The column will resort in ascending or descending order.
- **Moving columns.** To change a column order, press <Alt>, and then click and drag the column you want to move.

Printing reports or charts

Learn how to print reports to a printer, a PDF, an HTML file, a CSV file, or a formatted text file.

You can print any IBM Tivoli Storage Productivity Center report or chart displayed in the content pane.

To print a report or chart, complete the following steps:

1. Generate the report you want.

2. Click **File > Print** or click  on the tool bar. A print window opens.
3. Select the output that you want:
   - Select **Printer** to send the output directly to a printer
   - Select **PDF file** to save the output as a PDF file
   - Select **HTML file** to save the output as an HTML file
   - Select **CSV file** to save the output as a comma delimited file. Select **Include Headers** to include column headers in the report output file. Select **Include Totals** to include the totals column in the report output file.
   - Select **Formatted file** to save the output as a formatted text file
4. Click **Continue**.
   - If you selected **Printer**, the Page Setup window opens. Specify any custom settings for the print job on the Page Setup window and click **OK**. The report or chart is printed to your default printer.
   - If you selected **PDF file, HTML file, CSV file** or **Formatted file**, the Save As window opens. Select the location where you want to store the file and enter a name for the file in the **File Name** field. Click **Save**. Depending on the file type you selected, the report or chart is saved as a PDF or HTML file.

Including charts in reports

Learn how to create charts from the reports that you generate.

There are a large number of charts and chart types that can be generated from reports. All reports that support charting will have a  button at the top of the report. Any row in a report that supports charting on that row will have a  beside it.

Many charts can be generated simultaneously from the same report. After generating a chart, return to the **Report** tab and choose another chart or another set of objects to generate a chart upon. Another **Chart** tab will be generated to the right of the already existing tabs and you will be positioned to the new chart tab.

See “Generating charts” for information on how to generate charts from reports.

You can also customize the appearance of charts to determine how you want to present data. For example, you can create pie chart, bar chart, stacked bar chart, determine how many charts appear per page, specify time increments, and so on. See “Customizing charts” on page 479 for information on how to customize charts.

Generating charts

Learn how to generate charts.

To view a chart of all of the rows in the table:

1. Click  , which is at the top of the report. All reports that support charting will have this multi-colored pie chart icon at the top of the report.
2. Select the type of chart you would like to generate. The types of charts available will vary depending on the report. After you select the chart type, a new tab for that chart will appear in the content pane.

To generate a chart for just the element in that row:

- Click the chart icon that appears to the left of that report row.

To generate a chart for multiple, non-consecutive rows in a report:

1. Select rows in the report by holding down the control key (Ctrl key) on the keyboard while clicking on the desired rows.
2. Click any  button that appears to the left of the selected rows and select the type of chart you would like to generate. After you select the chart type, a new tab for that chart will appear in the content pane.

To generate a chart for a range of rows in a report:

1. Click on the first row in the range and then hold down the Shift and Ctrl while clicking on the end row in the range.
2. Click any  button that appears to the left of the selected rows and select the type of chart you would like to generate. After you select the chart type, a new tab for that chart will appear in the content pane.

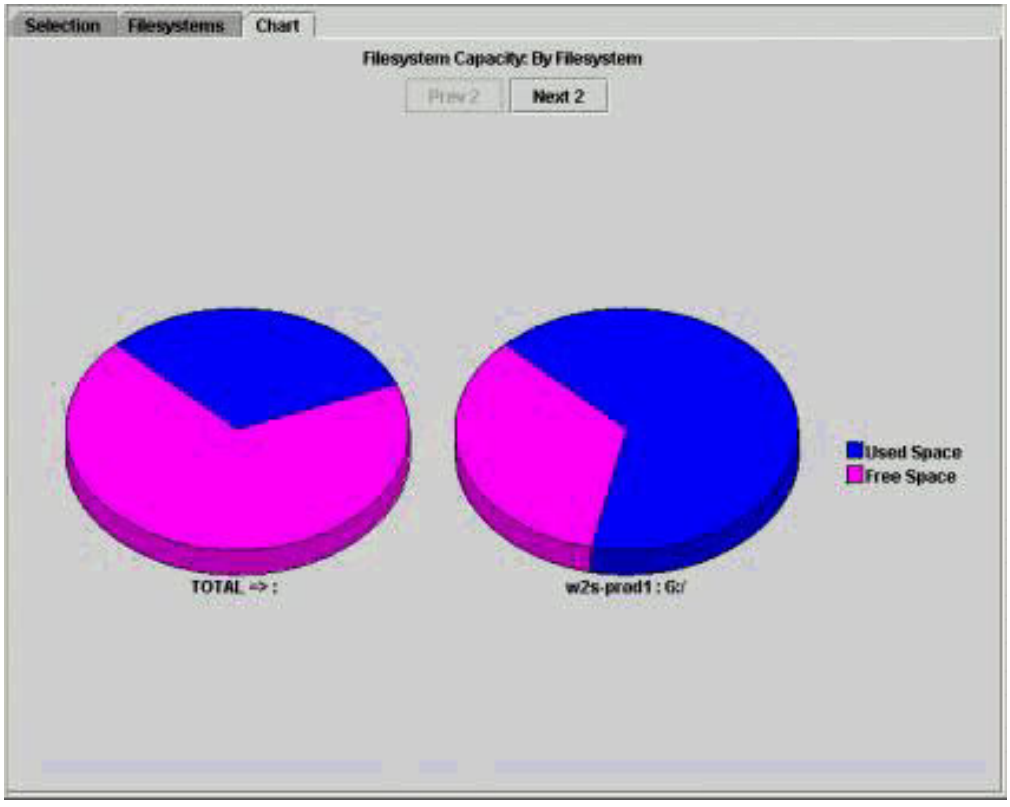

Figure 48. Chart Generation

Customizing charts

Learn how to customize the appearance of charts.

While viewing a chart, you can customize its appearance using the Chart Customization window. This includes:

- Entering the title that appears at the top of a chart

- Determining the maximum number of charts or series that appear per screen
- Specifying how to show values for data points or areas
- Selecting a type for a chart such as a pie, bar, stacked bar, and 100% stacked bar
- Determining the orientation (vertical or horizontal) of a chart
- Specifying history chart settings such as trends and time increment values

To access this window from a chart, right-click on the chart and select **Customize this chart** from the pop-up menu. The Chart Customization window appears.

Use this window to customize the chart generated from a report, a report created in the Reporting facility, a batch report, user reports, and dashboard reports. You can use the settings on this window to:

- Customize the display of the current chart.
- Customize the display of the current chart and save those changes. If you save the customization settings for a report chart, that chart will use those settings whenever it is generated.

If a field on this window does not apply to a selected report it will appear as grayed out and cannot be edited.

Field	Description
Chart title	Enter the name you want to appear at the top of the generated chart.
Maximum number of charts or series per screen	Select the maximum number of charts or series per page that should appear in the generated report from this list. For example, how many pie charts will appear on one page. When the number of charts that appear on one page exceeds the number in this field, the additional charts will be shown on a separate page. The most charts or series you can define to appear on a page is 30.
Show values for data points or areas	Determine how to display the data points/areas on a chart. You can choose to display the data points areas: • Only when hovering - display values for data points on an online chart when the mouse hovers over an appropriate place on that chart • Always - always display and print values for data points
Graph charts	Use the fields in this section to customize the appearance of graphs and charts.
Chart type	Select a type for the generated chart including: • Pie Chart • Bar Chart • Stacked Bar • 100% Stacked Bar

Field	Description
Chart orientation	Determine whether the chart graphic will appear vertically or horizontally on a page. You can set the orientation for the following chart types only: Bar Chart, Stacked Bar, and 100% Stacked Bar.
Maximum number of series per graph	Select the maximum number of series that should appear in one graph or chart. For example, how many slices will appear in one pie chart. When the number of series that appear in one graph exceeds the number in this field, the additional charts will be shown on a separate page. The default maximum is 12. The most series that can appear in one chart is 30. **Note:** Disk Capacity batch reports always display one series per chart, and one chart per page.
History Charts	Use the fields in this section to customize the appearance of history charts.
Show Trends	Check this box if you want IBM Tivoli Storage Productivity Center to show a trend for the data in a history chart. The trend attempts to show future storage usage in a chart based on earlier usage.
Time increment for x-axis	Select which value to use for the time increment that appears on the x-axis of a history chart. You can select By Scan, Weekly, Monthly. For example, if you choose Monthly, the x-axis will measure storage usage in monthly increments.
Restore Defaults	Reset the values in the fields on this window to their original default settings.
Apply	Apply the settings defined in this window to the current chart.
Save	Apply the settings defined in this window to the current chart and save those settings. When the chart is generated in the future, it will automatically use the saved settings.
Cancel	Exit the window without applying or saving the settings in the window.

Drilling down a report row

Learn how to drill down into report rows to view more detailed information about your storage resources.

Many of the rows that are presented in the generated reports allow you to drill down and get more reports of more detailed information. For example, if you are

looking at a report on computers, you might be able to drill down on a row for a specific computer and generate a report for the file systems contained on that computer.

When detailed drill down is allowed within a report, there is a  icon at the left of that row. To drill down, simply click on  and a new tab with the drill down report will be generated and added to the tab dialog. You will then be positioned on that tab viewing the new detail report that you have generated.

Navigating tabs

Learn how to navigate report tabs to view detailed sub reports.

As you generate reports and drill down on rows to generate detailed sub reports, the number of tabs in the Reporting window will grow. You can move about from tab to tab simply by clicking on the tab name.

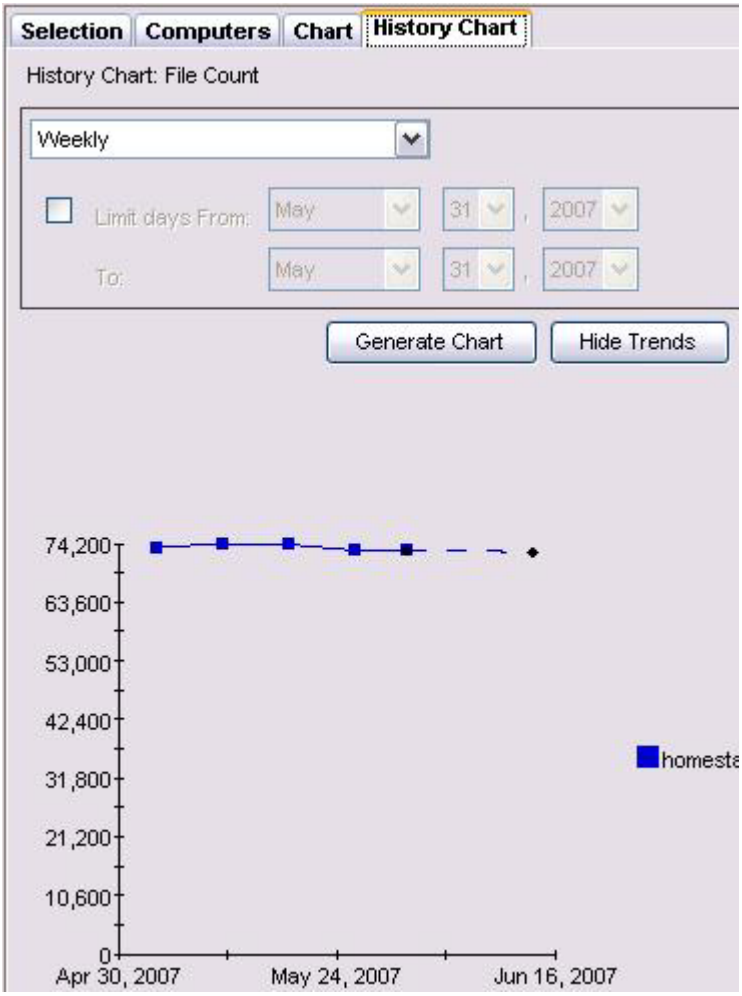

Figure 49. Tabs

You can also return to the **Selection** tab and change your selection criteria to start all over again with a newly generated report at any time. If you change your

selection criteria and generate a new report, all subsequent or dependent tabs (for example, drill down and chart tabs) are closed.

To close a tab, right-click on the tab folder top and select **Close** from the pop-up menu.

Trending with historical data

Learn how to view trends related to the capacity and usage of your storage resources.

The History Aggregator in IBM Tivoli Storage Productivity Center defines and runs jobs to sum data in the enterprise 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.

To view a historical data, simply right-click on the row or rows for which you want to view the data. A pop-up menu is displayed, enabling you to view historical charts.

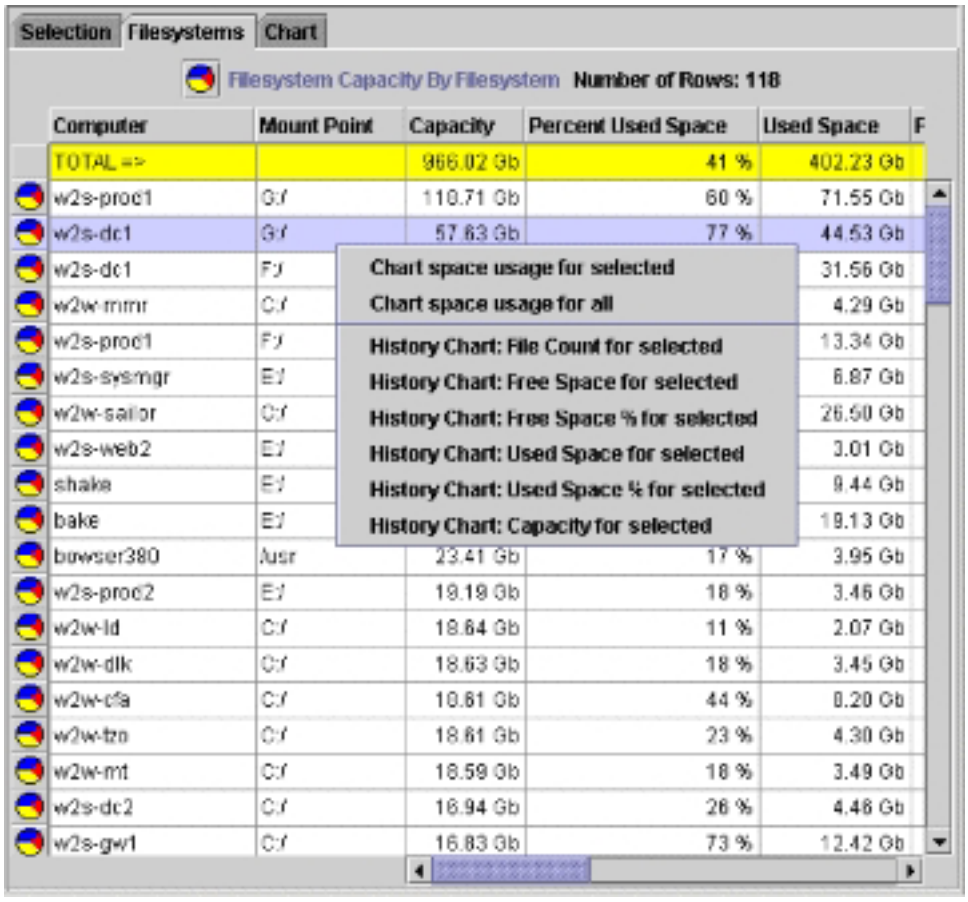

Figure 50. Trending with Historical Data

Depending on the type of report you are generating, you can generate a history chart for the items you select:

History chart	Description
File Count for Selected	Shows the number of files for selected item(s) over a specified period of time: daily, weekly or monthly
Free Space for Selected	Shows the free space available over a specified period of time: daily, weekly or monthly
Free Space % for Selected	Shows the percentage of free space available over a specified period of time: daily, weekly or monthly
Used Space for Selected	Shows used space over a specified period of time: daily, weekly or monthly
Used Space % for Selected	Shows percentage of used space over a specified period of time: daily, weekly or monthly
Capacity for Selected	Shows capacity for selected over a specified period of time: daily, weekly, or monthly
File Count	Shows the number of files accessed or modified over a period of: • 1 day • 1 day - 1 week • 1 week - 1 month • 1 month - 1 year • > year
File Count %	Shows the percentage of files accessed or modified over a period of: • 1 day • 1 day - 1 week • 1 week - 1 month • 1 month - 1 year • > year
Size	Shows the size of files accessed or modified over a period of: • 1 day • 1 day - 1 week • 1 week - 1 month • 1 month - 1 year • > year
Size %	Shows the percentage of file size accessed or modified over a period of: • 1 day • 1 day - 1 week • 1 week - 1 month • 1 month - 1 year • > year

Figure 51 on page 485 shows a history chart for used space for a file system:

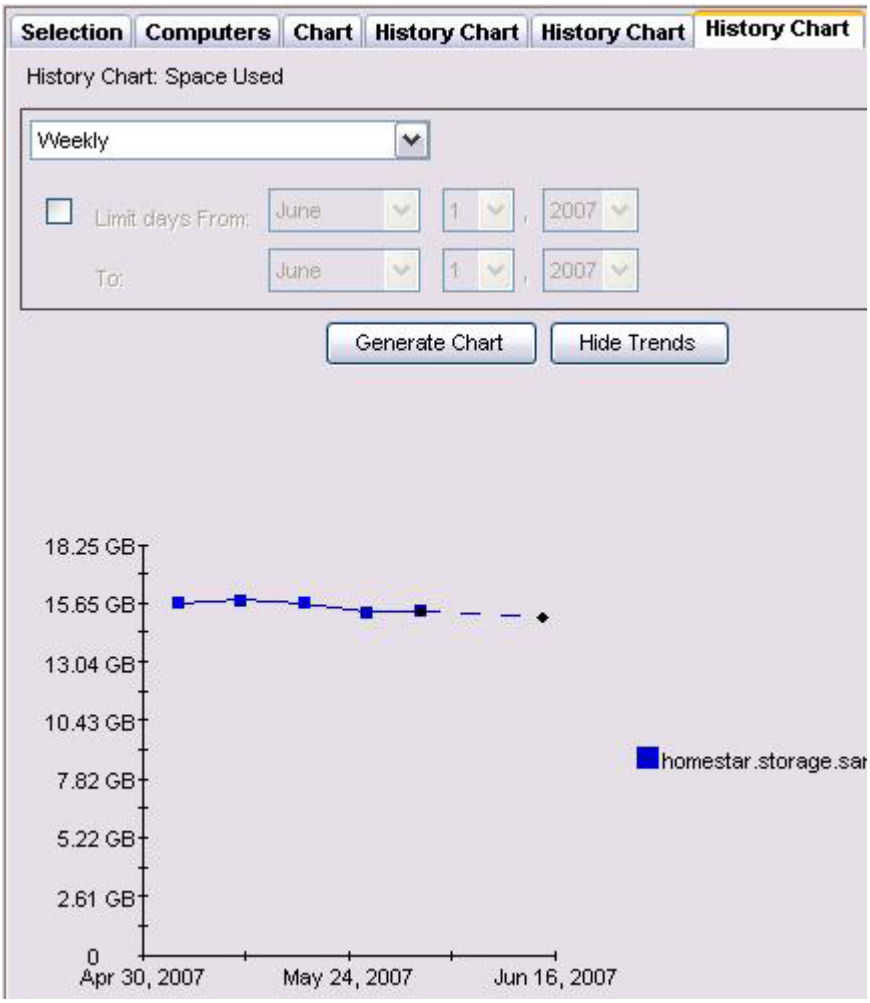

Figure 51. History Chart for File System Used Space

Select **File → Export data** to export history data to a comma delimited file (.csv). A dialog box is displayed allowing you to export the history data into a .csv file.

You can choose to include or not include headers in the comma delimited file you create. The column headings within the header reflect the type of report from which you are exporting history data.

User-defined online reports

Use the information in this section to learn how to generate and view online reports. These reports are available in the **Reporting** nodes under Data Manager, Data Manager for Databases, Disk Manager, and Fabric Manager.

Generating reports

Learn how to generate and view online reports in Data Manager, Data Manager for Databases, Disk Manager, and Fabric Manager.

1. Expand the navigation tree to display the type of report you want to generate. For example, if you want to generate a disk capacity report, expand **Data Manager → Reporting → Capacity → Disk Capacity**.

2. Highlight a By Clause to determine how you want to view a report. For example, to view capacity information according to computer group, click **Reporting → Capacity → Disk Capacity → By Computer Group**. The **Selection** page is displayed.
3. Use the **Selection** page to select the profile to use when generating a report and determine what columns appear within a report. See "Selecting resources to appear on a report" on page 488 for information.
4. Click **Selection...** to select the objects that you want to report upon from the Select Resources window.
5. Click **Filter...** to further filter the objects that appear in a report. Filters enable you to apply general rules to the report based on the rows in that report. See "Filtering the resources that appear on a report" on page 491 for more information on filtering.
6. Use the **Available Columns** and **Included Columns** list boxes to determine what columns are displayed in a generated report.
7. Click **Generate Report**. A new tab will be added to the tab dialog representing the report that you generated.
8. Click the new tab to view the report. Reports are tabular in format and composed of rows and columns. You can scroll the report up and down and to the right and left to view the entire report. Use the **View** menu to hide/show the navigation tree to increase the viewable area of the report or drag the divider bar in the middle of the screen back and forth to reallocate the amount of space that is given each pane.

Defining options for a report

Learn how to customize a report before generating it.

Depending on the report that you want to generate and the level upon which you have chosen to report (for example, by computer, by filesystem, network wide, and so on), you are typically presented with a **Selection** page. For example, if you want to generate the Most Obsolete Files report, expand **Data Manager → Reporting → Usage → Files → Most Obsolete Files → By Computer**. Use the **Selection** page to:

- Select the profile to use when generating a report
- Select the resources (for example, computers, computer groups, storage subsystems, and so on) to display in a generated report
- Apply filters to the columns within a report (for example, show only the systems that contain over 80% used space) to determine what resources appear within the rows of a report
- Determine what columns appear within a report

The following table describes the interface elements that might appear on the **Selection** page (depending on the report that you want to generate):

Interface element	Description
Profile (Data Manager only)	This field indicates the profile to use when generating the report. If you just want to use the supplied default profile, then leave the field as is and generate the report. However, if you want to use a customized profile that you have set up for this reporting, choose that profile from the list. Keep in mind that profiles control what statistics are gathered and also what files are included in that statistical analysis.

Interface element	Description
Return maximum of __ rows per *storage_resource*	This field controls how many rows appear in a generated report. The value that appears as the *storage_resource* in this field depends on the selected By clause. For example, if you select **Data Manager → Reporting → Usage → Files → Most Obsolete Files → By Computer**, you can enter **10** in the **Return maximum of __ rows per computer** field and to limit the number of rows in the report to only 10 obsolete files per computer. If you have set up scans to collect many more rows than that per computer, this is an effective and efficient way to control the size of the generated report. This field enables you to control the size of the report that is generated for different by clauses.
Limit to this time range:	The fields in this section determine the date range from which to gather availability information for the objects that you select.
Relate Computers to: (Disk Manager - Storage Subsystem reports only)	This field generates a report that shows the relationship between the computer and one of these options: storage subsystems, volumes, or disks. • Storage Subsystem. Choose this option to view the Storage Subsystems that the logical volumes within a computer are using. • volumes. Choose this option to show the volumes within a storage subsystem that the logical volumes within a computer are using. • Disks. Choose this option to show the disks within a disk array that the logical volumes within a computer are using.
Relate Filesystems/Logical Volumes to: (Disk Manager - Storage Subsystem reports only)	This field generates a report that shows the relationship between the file systems/logical volumes and one of these options: storage subsystems, volumes, or disks. • Storage Subsystem. Choose this option to view the storage subsystems on which a file system/logical volume is allocated. • volumes. Choose this option to view the volumes on which the file system/logical volume is allocated. • Disks. Choose this option to show the disks on which a file system/logical volume is allocated.
Relate Storage Subsystems to: (Disk Manager - Storage Subsystem reports only)	This field generates a report that shows the relationship between a storage subsystem and one of these options: computer or file system/logical volume. • Computer. Choose this option to view the computers that are using a storage subsystem. • Filesystem/Logical Volume. Choose this option to view the file systems/logical volumes that are using a storage subsystem.
Relate Volumess to: (Disk Manager - Storage Subsystem reports only)	This field generates a report that shows the relationship between a volume and one of these options: computer or file system/logical volume. • Computer. Choose this option to view the computers that are using a volume. • Filesystem/Logical Volume. Choose this option to view the file systems/logical volumes that are using a volume.
Relate Disks to: (Disk Manager - Storage Subsystem reports only)	This field generates a report that shows the relationship between a disk and one of these options: computer and file system/logical volume. • Computer. Choose this option to view the computers that are using a disk. • Filesystem/Logical Volume. Choose this option to view the file systems/logical volumes that are using a disk.

Interface element	Description
Selection...	This button is used to determine the storage resources (computers, file systems, logical volumes, storage subsystems, and so on) upon which you want to report. When you click **Selection...**, the Select Resources window opens. If you clear the check box for a resource on the Select Resource window, a push-pin icon displays on the **Selection...** button. If you are generating a Network-wide report, the **Selection...** button will be disabled because you have chosen to generate the report network wide, for all objects.
Filter...	Use this button to further select or restrict the values within the columns you want to include on a report. Filters enable you to apply general rules to the selected columns. When you click **Filter...**, the Edit Filter window opens. For example, select all computers with the name beginning with *w2s-**, or select all file systems with percent used space >= to a specific value. If a filter is defined for this report, a pushpin icon displays on the **Filter...** button. **Note:** For storage entities that contain an asterisk (*) within their names, the * is used as a wildcard character if you select **LIKE** from the **Operator** menu on the Edit Filter panel. See "Filtering the resources that appear on a report" on page 491 for more information.
Available Columns / Included Columns	These panels display the columns that can be included or are already included in the generated report. • The **Available Columns** list displays the columns that you can include in the generated report that are not already included. • The **Included Columns** list displays the columns that will appear in the generated report. By default, all the columns for a report are listed in this list box and will be displayed in the generated report. You can highlight the columns and use the up and down arrows to rearrange columns in the report. The order of the items in this list determines the order in which the columns will appear in the generated report. The greyed-out names shown at the top of the list are frozen columns of the report. They cannot be selected, removed, or reordered. To include a column in a report: 1. In the **Available Columns** list, click the name of the column. To select additional columns, press Shift and click the columns that you want to add or press Ctrl and click the first and last column to select a range of columns. 2. Click the right-arrow button (>>) to add the columns to the **Included Columns** list. To remove a column in a report: 1. In the **Included Columns** list, click the name of the column. To select additional columns, press Shift and click the columns that you want to add or press Ctrl and click the first and last column to select a range of columns. 2. Click the left-arrow button (<<) to add the columns to the **Available Columns** list.

Selecting resources to appear on a report

Learn how to select the storage resources that will appear in a report.

Use the Select Resource window to determine the objects (computers, filesystems, logical volumes, storage subsystems, etc.) upon which you want to report. To select the resources on which you want to report:

1. Select the type of report you want to generate. For example, **Data Manager** → **Reporting** → **Capacity** → **Disk Capacity** → **By Disk**.
2. Click **Selection**. The Select Resources window opens.

The objects that appear in this window are determined by the report you are generating and the By Clause you selected. For example, if you choose to generate the report **Data Manager** → **Reporting** → **Capacity** → **Disk Capacity** → **By Disk**, the Select Resource window will display a list of computers and subsystems that you can include in the report.

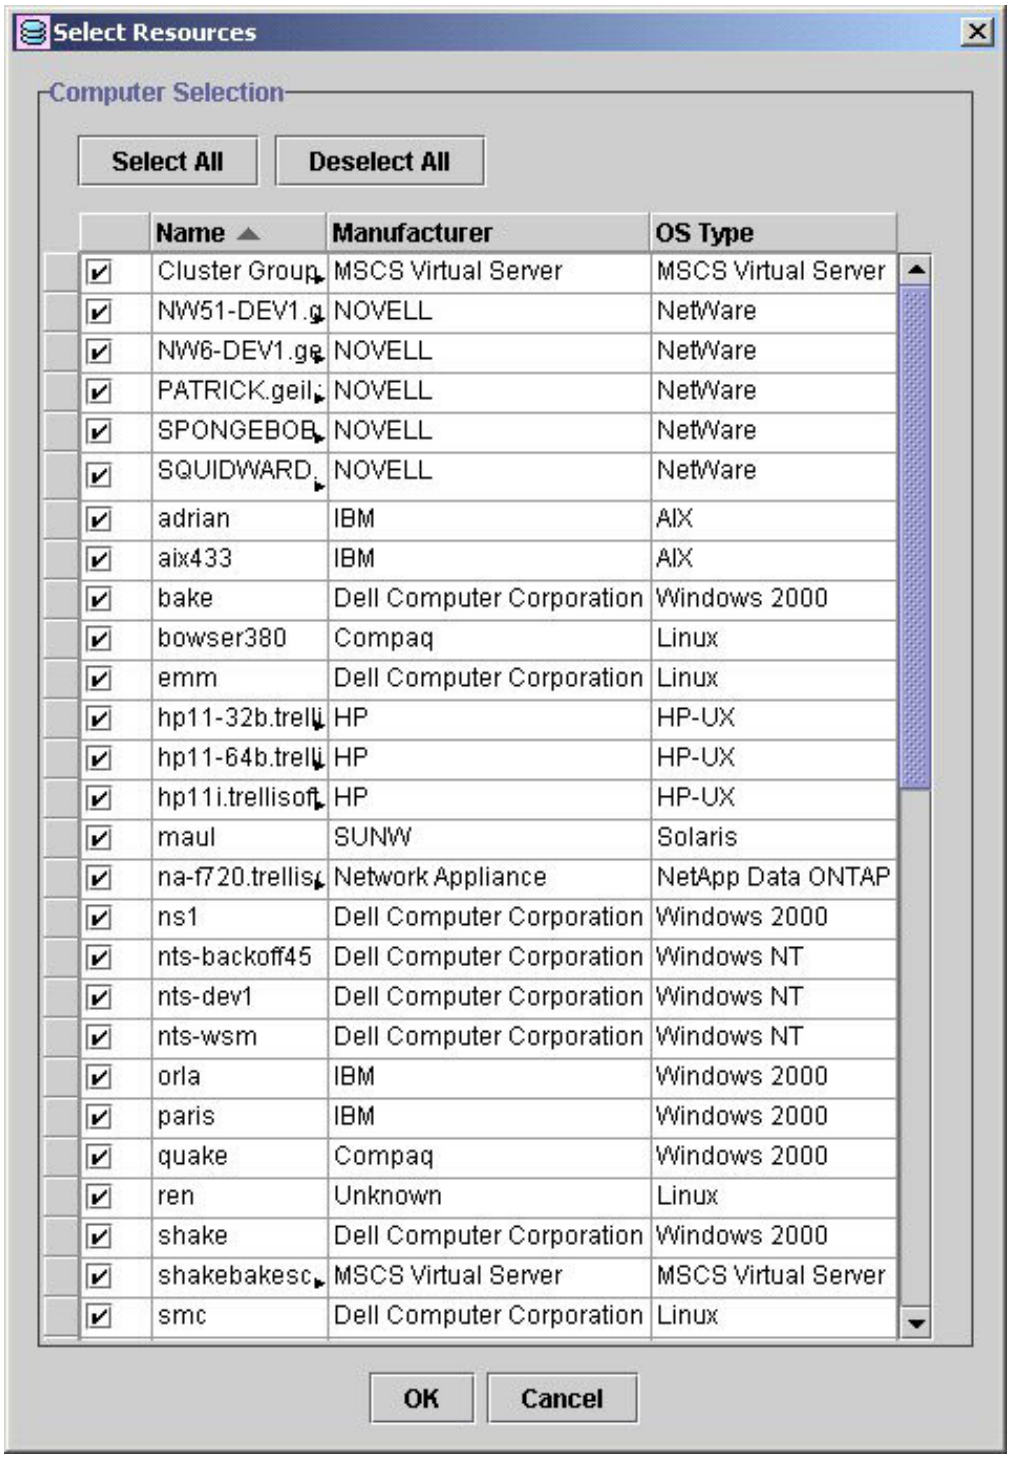

Figure 52. Selection Resources window

When the Select Resource window is initially opened, all objects in the list are selected.

- If you would like to generate the report for all the objects that have been selected, click **OK**.
- To select/deselect individual objects in the selection list, simply select the check box to the left of the object. When there is a check mark in the check box, then the object has been selected for inclusion in the report.

You can use **Select All** and **Deselect All** to quickly select and deselect all of the objects in the list.

Note: If you are generating a Network-wide report, the **Selection** button is disabled on the **Selection** tab because you have chosen to generate the report network wide, for all objects.

Filtering the resources that appear on a report

Learn how to apply conditions to the columns in a report to further filter the resources that appear in a report.

Use the Edit Filter window to apply conditions to the columns within a report to further select or restrict objects on which you want to report. Filters enable you to apply general rules to the report. For example, you can show only the file systems that contain over 80% used space, or show computers that match ″w2s-*″, or select all disks with percent used space >= to a specific value. To filter the objects that appear in a report:

1. Select the type of report you want to generate. For example, **Data Manager** → **Reporting** → **TPC-wide Storage Space** → **Disk Space** → **By Disk**.
2. Click **Filter...**. The Edit Filter window opens.

Note: For storage entities that contain an asterisk (*) within their names, the * is used as a wildcard character if you select **LIKE** from the **Operator** menu on this panel. For example, if you enter ***XYZ** in the **Value 1** field and select **LIKE** from the **Operator** menu, the resulting report includes any entities named ″*XYZ″, and also all entities whose names end with the letters ″XYZ″. To create a filter that includes only the storage entity named ″*XYZ″ in the resulting report, enter ***XYZ** in the **Value 1** field and select **IN** from the **Operator** menu.

The following table describes the interface elements that appear on the **Edit Filter** window:

Interface element	Description
Add	This button is used to specify multiple conditions for the filter. When you click **Add**, report columns will be displayed enabling you to define conditions for those columns.
Delete All	This button is used to delete all the conditions displayed on the table.
Case-sensitive string comparisons	This check box indicates that case-sensitive string comparisons should be done. Select this check box to ensure that values appearing in the generated report match the case of the values entered in the Value 1 and Value 2 fields.
Records must meet	This field specifies how the conditions in the filter are grouped together. • **All conditions**. All of the entered conditions (rows on the dialog box), must be true in order for the record to show up as a row in the report. • **At least one condition**. At least one of the entered condition must be true in order for the record to show up as a row in the report.

Interface element	Description
Column	This field is used to select specific columns on which you want to define conditions. The list on this field shows the columns available for the report. As you select a column from the list, it is displayed on the table, and the column name is removed from all other lists. Click  to delete the column.
Operator	This field is used to select the condition's operator. Not all operators apply to all the column data-types.
Value 1/ Value 2	This field is used to specify fixed value(s) to be compared against the record's column value. Depending on the operator and data type, this field may be absent or display-only. If the field is display-only, an **Edit...** button will be enabled. If both this field and the **Edit...** button are enabled, this field allows shorthand value entry. If you click **Edit** , a window opens to provide more assistance in editing the value. A **Value 2** field will be present only for (NOT) BETWEEN. For certain columns that contain string text, you can: • put the string in quotes for the **LIKE** operator, or • use the **=** operator without quotes instead of LIKE (the default) For example, if you want to filter the Device Path (computer) column on the **Reporting → Storage Subsystems → Computer Views → By Filesystem/Logical Volume** report, you can search/filter on a name such as ″**Disk 1**″ or ″**Disk 2**″ with the **LIKE** operator. For the **=** operator, you can enter **Disk 1** or **Disk 2** with no quotes. **Note:** An asterisk is prepended to the names of space-efficient VDisks in the IBM Tivoli Storage Productivity Center database repository. When you use a filter to determine which volume names appear in a report, you must include the ″*″ as the first character in the filter value to ensure that those VDisks are included. For example, if you want to include VDisks whose name begins with ″*vdisk1″, you must enter ***vdisk1** in the **Value 1** field and select **LIKE** or **IN** from the **Operator** menu. If you select the **LIKE** operator, the *vdisk1 VDisk is included in the report, as well as other VDisks whose names end in ″vdisk1″.
Edit	This button is used to edit values pertaining to size, date, and for specific values used for IN, LIKE, etc. operators. Depending on the condition you are creating, a dialog box displays to help you to enter values that correspond to the condition.

Interface element	Description
Selection...	This button is used to determine the objects (computers, file systems, logical volumes, storage subsystems, etc.) upon which you want to report. The Select Resources window opens. If you clear the check box for a resource on the Select Resource window, a pushpin icon displays on the **Selection...** button. If you are generating a Network-wide report, the **Selection...** button will be disabled because you have chosen to generate the report network wide, for all objects.
OK	This button is used to save the filter. You are returned to the **Selection** tab. The **Filter** button will display a pushpin to indicate that a filter has been defined for the report.
Cancel	This button is used to return to the **Selection** tab without saving the filter.

Using reporting groups

A reporting group is a user-defined set of storage assets, monitoring groups, and reporting groups upon which you want to report.

Use reporting groups to view storage information about:

- Resources that appear in different monitoring groups. Because a resource can appear in only one monitoring group at a time, you can use reporting groups to include resources from different monitoring groups and view reports on them.
- A subset of resources within a monitoring group. For example, you can create a reporting group for a specific subset of computers with the Payroll monitoring group and generate a report containing information about those computers.
- Resources that appear in other reporting groups.
- A subset of resources within other reporting groups.

Note: Unlike monitoring groups, a resource can appear in one or more reporting groups.

You can create the following types of reporting group:

- Data Manager: **Computer**, **Filesystem**. Within a computer reporting group, you can include computer monitoring groups, specific computers, and other reporting groups. Within a file system reporting group, you can include file system monitoring groups, specific file systems, and other reporting groups.
- Data Manager for Databases: **Computer**, **Databases-Tablespaces**. Within a computer reporting group, you can include computer monitoring groups, specific computers, and other computer reporting groups. Within a databases-tablespaces reporting group, you can include databases-tablespace monitoring groups, specific databases or tablespaces, and other databases-tablespaces reporting groups.
- Disk Manager: **Storage Subsystems**. Within a storage subsystem reporting group, you can include storage subsystem monitoring groups, specific storage subsystems, and other storage subsystem reporting groups.

When you generate a report, storage statistics about a reporting group will appear in the corresponding ″By Group″ clause for that report. The following lists the reports that reflect reporting group information.

Computer reporting groups (Data Manager and Data Manager for Databases)
The following reports display computer reporting group information:

- Data Manager > Reporting > Capacity reports, Data Manager for Databases > Reporting > Capacity reports
- Data Manager > Reporting > Usage reports, Data Manager for Databases > Reporting > Usage reports
- Data Manager > Reporting > Usage Violation reports, Data Manager for Databases > Reporting > Usage Violation reports
- Data Manager> Reporting > Backup reports
- IBM Tivoli Storage Productivity Center > Rollup Reports > Capacity
- IBM Tivoli Storage Productivity Center > My Reports > Batch Reports > Create Batch Report > Report page, select one of the following:
  - Capacity > capacity_report_type > By Computer Group
  - Usage > usage_reporting_type > By Computer Group
  - Usage Violations > usage_violation_reporting_type > By Computer Group
  - Backup > backup_reporting_type > By Computer Group
  - Rollup Reports > Capacity > capacity_report_type > By Computer Group
  - Rollup Reports > Database Capacity > rdbms_type > By Computer Group

Filesystem reporting groups (Data Manager)
The following reports display filesystem reporting group information:

- Data Manager > Reporting > Capacity reports
- Data Manager > Reporting > Usage reports
- Data Manager > Reporting > Usage Violation > Constraint Violations reports
- Data Manager > Reporting > Backup reports
- IBM Tivoli Storage Productivity Center > My Reports > Batch Reports > Create Batch Report > Report page, select one of the following:
  - Capacity > capacity_report_type > By Filesystem Group
  - Usage > usage_reporting_type > By Filesystem Group
  - Usage Violations > Constraint Violations > By Filesystem Group
  - Backup > backup_reporting_type > By Filesystem Group

Databases-Tablespaces reporting groups (Data Manager for Databases)
The following reports display databases reporting group and tablespaces reporting group information:

- Data Manager for Databases > Reporting > Capacity reports
- Data Manager for Databases > Reporting > Usage reports
- Data Manager for Databases > Reporting > Usage Violations > Database-Tablespace Quotas reports

Storage subsystem reporting groups
The following reports display storage subsystem reporting group information:

- Data Manager > Reporting > Storage Subsystems > Storage Subsystem Views > By Storage System Group
- Data Manager > Reporting > Capacity > Disk Capacity > By Storage Subsystem Group
- IBM Tivoli Storage Productivity Center > My Reports > Batch Reports > Create Batch Report > Report page, select one of the following:
  - Storage Subsystems > Storage Subsystem Views > By Storage Subsystem Group
  - Capacity > Disk Capacity > By Storage Subsystem Group

Creating a reporting group

To create a reporting group:

1. Expand **Reporting** → **Groups** under:
   - **Data Manager** to create computer and filesystem reporting groups.
   - **Data Manager for Databases** to create computer and databases-tablespaces reporting groups
   - **Disk Manager** to create storage subsystem reporting groups
2. Right click the type of group you want to create. A pop-up menu appears.
3. Select **Create <group type>** from the pop-up menu. The content pane displays a window for defining the group.
4. Define information about the group. See the IBM Tivoli Storage Productivity Center online Help for detailed information about the windows for creating a group.
5. Click **File** → **Save** to save the group definition. The group is saved to the database repository. You will now be able to generate and view reports based on these groups.

Note: You can define filters to select specific computers, filesystems, storage subsystems, tablespaces, and databases for a reporting group. Tivoli Storage Productivity Center uses these filters to dynamically include resources to the group as they are discovered within your environment. When defining dynamic filters for reporting groups, keep in mind that the filters you enter might or might not be case sensitive depending on the database under which the Tivoli Storage Productivity Center database repository is stored. For example:

- The dynamic filters you enter for reporting groups are case sensitive when the database repository is stored on Oracle, Sybase, and UDB/DB2.
- The dynamic filters you enter for reporting groups are not case sensitive when the database repository is stored on Microsoft SQL Server.

Keep in mind these case-sensitive issues when defining filters for reporting groups to prevent inaccurate reporting caused by unintended resources being included in reports.

Editing a reporting group

To edit an existing reporting group:

1. Expand **Reporting** → **Groups** under:
   - **Data Manager** to edit computer and filesystem reporting groups.
   - **Data Manager for Databases** to edit computer and databases-tablespaces reporting groups
   - **Disk Manager** to edit storage subsystem reporting groups

2. Expand the type of the reporting group you want to edit. For example, expand **Disk Manager** → **Reporting** → **Groups** → **Storage Subsystem** to edit a storage subsystem reporting group.
3. Click the reporting group you want to edit. The content pane displays a window for editing the group.
4. Edit information about the group. See the IBM Tivoli Storage Productivity Center online Help for detailed information about the windows for creating a group.
5. Click **File** → **Save** to save the updated group definition.

Note: You can define filters to select specific computers, filesystems, storage subsystems, tablespaces, and databases for a reporting group. Tivoli Storage Productivity Center uses these filters to dynamically include resources to the group as they are discovered within your environment. When defining dynamic filters for reporting groups, keep in mind that the filters you enter might or might not be case sensitive depending on the database under which the Tivoli Storage Productivity Center database repository is stored. For example:

- The dynamic filters you enter for reporting groups are case sensitive when the database repository is stored on Oracle, Sybase, and UDB/DB2.
- The dynamic filters you enter for reporting groups are not case sensitive when the database repository is stored on Microsoft SQL Server.

Keep in mind these case-sensitive issues when defining filters for reporting groups to prevent inaccurate reporting caused by unintended resources being included in reports.

Data Manager reports

Use Data Manager reports to view both overview and detailed information about your storage resources, including computers, disks, filesystems, storage subsystems, clusters, and users. These reports are constructed from the statistical information gathered by data collection jobs and accumulated in the database repository.

Note: You need to run appropriate data collection jobs on the storage resources before you generate reports. If you do not run the data collection jobs before generating reports, you might see inconsistent or inaccurate results.

Data Manager - Report categories

Use the report categories within Data Manager to collect information about storage assets, availability, backup requirements, capacity, usage, usage violations, and monitored computer storage space.

The following table describes the categories of reports available in Data Manager.

Report Category	Use these reports to do the following	What job collects information for these reports?
Asset reports	View information about the agents, computers, disk controllers, hard disks, clusters, fabrics, storage subsystems, volumes, tape libraries, file systems, and general hardware inventory of storage assets within your environment. See "Asset reports" on page 499 for more information about Asset reports.	Probes, discovery Data for System-wide Users and Monitored Directory reports are gathered by scans.

Report Category	Use these reports to do the following	What job collects information for these reports?
Availability reports	View information about the availability of your storage, both from a network point of view and from a computer up-time perspective. This allows you to see what percentage of the time the storage in your enterprise (or on particular server or group of servers), is off-network due to network problems or perhaps is unavailable due to system downtime. See "Availability Reports" on page 515 for more information about Availability reports.	Ping data: pings Computer uptime — probes
TPC-wide Storage Space reports	View storage capacity at a file system, computer, or entire network level. For a specific computer, group of computers, or all the computers in your enterprise, you can view the following: • How much storage you have • How much storage is being used • What percentage of storage is free for expansion See "TPC-wide Storage Space Reports" on page 517 for more information about Capacity reports.	Probes
Usage reports	View information about the usage and growth of the consumption of your storage. Use these reports to do the following: • View which servers and file systems are experiencing the heaviest (or lightest) load of data and storage access. • Identify wasted space by pinpointing files that are no longer needed or have not been accessed for the longest time. • Associate users with the amount of storage they are consuming on specific file systems, groups of file systems, specific computers, groups of computers, and throughout the entire network. • Select the files against which to run an IBM Tivoli Storage Manager archive/backup job. Tivoli Storage Manager protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage. See "Usage Reports" on page 573 for more information about Usage reports.	Scans

Report Category	Use these reports to do the following	What job collects information for these reports?
Usage Violation reports	View information about the storage usage violations committed by users and user groups within your environment. Use these reports to do the following: • Create and enforce corporate storage usage policies. You can report on violations of quotas and constraints. • Select files in Constraint Violation reports against which to run an Tivoli Storage Manager archive or backup job. • Add files from Constraint Violation reports to an existing Tivoli Storage Manager archive or backup job. See "Usage Violation reports" on page 581 for more information about Usage Violation reports.	Scans, Quotas, Constraints
Backup reports	Investigate the sizes of data that would be required to backup to achieve full backups, determine the amount of space required to perform an incremental backup of any or all of your data at any given point in time, and select files for archive/backup jobs. See "Backup reports" on page 583 for more information about Usage Violation reports.	Scans
Monitored Computer Storage Space reports	Use these host-based reports to view the data that is collected by probes for computers, Network Attached Storage, and hypervisors. These reports contain detailed information about the consumption of file system space and enable you to relate that space back to external or internal storage without counting shared space multiple times. See "Monitored Computer Storage Space reports" on page 584 for more information about Monitored Computer Storage Space reports.	Probes

Using By clauses

Use the By clauses associated with Data Manager reports as a method for viewing the data in those reports from different perspectives.

As you drill down through Data Manager reports you will notice a set of options that repeats for each report type. These *By Clauses* allow you to generate variations of the reports with different perspectives and groupings of the data.

The available By Clauses are described in the following table:

By Clause	Description
By Directory	Use this clause to select the directories that you want to report upon and the report that will be generated based upon those directories.
By Directory Group	Use this clause to select the directory groups that you want to report upon and the report that will be generated based upon those directory groups.

By Clause	Description
By Disk	Use this clause to select the disks that you want to report upon and the report that will be generated based upon those disks.
By Filesystem	Use this clause to select the file systems that you want to report upon and the report that will be generated based upon those file systems.
By Filesystem Group	Use this clause to select the file system groups that you want to report upon and the report that will be generated based upon those file system groups.
By Cluster	Use this clause to select the clusters or cluster resource groups that you want to report on and the report that will be generated based upon those clusters. Cluster resource groups in an HACMP or MSCS clustered environment are listed as computers.
By Computer	Use this clause to select the computers that you want to report upon and the report that will be generated based upon those computers.
By Computer Group	Use this clause to select the computer groups that you want to report upon and the report that will be generated based upon those computer groups.
By Storage Subsystem	Use this clause to select the storage subsystem that you want to report upon and the report that will be generated based upon those storage subsystems.
By Domain	Use this clause to select the Window Domains that you want to report upon and the report that will be generated based upon those domains.
Network Wide	The reports and charts that you see will be generated on all data, network wide. This allows you to get a global view of the storage and storage usage for your whole enterprise.

Asset reports

Use asset reports to view information about the agents, clusters, cluster resource groups, computers, disks, storage subsystems, controllers, filesystems, and general hardware inventory of the storage assets within your environment.

While running probes, IBM Tivoli Storage Productivity Center itemizes the information about your storage assets and provides a hierarchical view of that information so you can drill down to view your assets in greater detail. The following *by* clauses also enable you to change your view of asset reports:

Asset Reporting Type	Description
By Cluster	View your storage assets based on HACMP or MSCS cluster membership. You can drill down into a cluster and view its nodes and cluster resource groups. For each cluster resource group, you can drill down to expose its controllers, disks, file systems, logical volumes, and exports or shares. For a cluster node, local resources can be viewed.
By Computer	Drill down through your storage assets in a hierarchical manner. You can drill down into a computer and expose its controllers, disks, file systems, logical volumes, and directories. You can then drill down on a specific controller to see the disks under it and/or drill down on a disk to see the file systems under it, and so on.

Asset Reporting Type	Description
By Hypervisor	Drill down through your hypervisors in a hierarchical manner. You can use this report to drill down and view details about the virtual machines, controllers, disks and file systems or logical volumes.
By OS Type Windows NT Windows 2000 Solaris Linux HP-UX AIX Netware Unknown	View your storage assets based upon the operating systems that are running on those storage servers. This gives you an easy way to see what your OS dependencies are throughout your organization. By drilling down on any or each of these different supported Operating System types, you can see how many of each server type you have and you can drill down into them to investigate the storage underlying them.
By Storage Subsystem	View storage capacity at a computer, file system, storage subsystem, volume, and disk level. These reports also enable you to view the relationships among the components of a storage subsystem. Reports can be generated for IBM Tivoli Storage Enterprise Storage Server subsystems as well as IBM SAN Volume Controller clusters.
Storage Subsystem name	View information about a storage subsystem. You can also view storage subsystems based on the following report types:
Disk Groups	View information about the disk groups related to a storage subsystem, as well as the volume spaces and disks associated with those disk groups.
Volume Spaces	View information about the volume spaces related to a storage subsystem, as well as the disks and volumes associated with those volume spaces.
Disks	View information about the disks related to a storage subsystem and the volumes associated with those disks.
System Wide	View your storage assets as a whole, across the entire system. This lets you get lists of all of your disks, all of your computers, all of your file systems, the Data agents installed in your environment, as well as viewing aggregated reports based upon certain filters, such as grown defects per disk.
Agents	View information about IBM Tivoli Storage Productivity Center agents based upon the following report types: By Agent, By OS (operating system) Type, By Status, and By Version.
Computers	View information about your monitored computers based upon the following report types: • By Boot Time: sort the report based on when the computer was last booted. • By Probe Time: sort the report based on when the last probe ran on the computer. • By Discovered Time: sort the report based upon when the computer was first detected on the network. Use this information to help determine what is the oldest equipment in your environment and most likely to need replacement sooner.
Storage Subsystems	View information about your monitored storage subsystems based upon the following report types: • By Storage Subsystem: sort the report based on the names of storage subsystems. • By Disk Space: sort the report based on total disk space in storage subsystems. •
Unmanaged Virtual Machines	Use this report to view detailed information about the virtual machines in your environment. This information is provided by the ESX Servers. These virtual machines are discovered through the hypervisors but do not have a Data agent installed on those machines.

Asset Reporting Type	Description
Virtual Machines With No VMWare Agent	Use this report to view detailed information about the virtual machines that have Data agents installed but are not discovered through the VMWare data sources.
Disk/Volume Groups	Data Manager provides support for the following volume managers: • Veritas Volume Manager (Sun Solaris) • Logical Volume Manager (AIX) These volume managers allow you to create groups of logical volumes and disks. You can sort the content of this report using the following By clauses: • By Computer • By Disk/Volume Group Space • By Probe Time • By Available Disk/Volume Group Space • By Discovered Time •
Disks	View information about the monitored disks in your environment. You can sort the content of this report using the following By clauses: • By Available Disk Space • By Grown Defects • By Primary Defects • By Recovered Defects • By Unrecovered Defects • By Probe Time • By Discovered Time •
File Systems or Logical Volumes	View information about the monitored file systems or logical volumes in your environment. You can sort the content of this report using the following By clauses: • By Available File System Space • By File System Probe Time • By File System Scan Time • By File System Discovered Time • By Removed File Systems: View file systems that have been removed. • By Logical Volumes without File Systems: View logical volumes that do not have file systems on them. • Unused Virtual Disk Files: View a list of virtual disk files which are not assigned to any of the virtual machines. This represents wasted storage capacity. • Mapping to Hypervisor Storage: View this report to Identify the paths from the file systems on virtual machines to the backend storage volumes on hypervisors or mapped to hypervisors. •
Volumes	View information about the monitored volumes in your environment. You can sort the content of this report using the following By clauses: • By Storage Subsystem • By Volume Space • By Probe Time •

Asset Reporting Type	Description
Exports or Shares	View all of your exports and shares sorted by the time when they were discovered by IBM Tivoli Storage Productivity Center.
Monitored Directories	View all of your monitored directories in your environment. You can sort the content of this report using the following By clauses: • By Space Used • By Scan Time • By Discovered Time • Removed Directories: View the monitored directories that have been deleted or moved to a different path. The retention settings defined in the **Administrative** node of the navigation tree determine how long directories that have been moved or deleted from the system will appear in this report.
Unmanaged Computers	View all of the computers that are discovered, via discovery, which are not being scanned or probed. These computers are not being monitored by Data Manager, therefore are considered to be ″unmanaged″. It is most likely that these are new computers, that have been placed on the system and have not yet been scheduled for scanning or probing (monitoring).
Users	View information about the users in your environment sorted by the time when they were discovered by IBM Tivoli Storage Productivity Center.
OS User Groups	View information about the OS users in your environment sorted by the time when they were discovered by IBM Tivoli Storage Productivity Center.

By Cluster:

Use the By Cluster reports to drill down into a cluster to view its nodes and cluster resource groups.

- For a node, you can view its local resources, including controllers, disks, file systems, logical volumes, exports, shares, and directories.
- For a cluster resource group, you can view controllers, disks, file systems, logical volumes, exports, and shares.
- For a node, the number of processors that are reported are the logical processors. The amount of RAM that is reported is the available RAM.

1. Expand **Asset → By Cluster** node to display the nodes for each cluster that has been discovered and monitored by Data Manager.
2. Click a node to view cluster information, or further expand the cluster node to drill down into greater detail.

Cluster Information

Click the name of a cluster that appears under **Asset → By Cluster** to view information about that cluster such as:

- Name
- Domain
- IP address
- Discovered Time

Controller Information

As you drill down into Asset Reporting By Cluster for a cluster node or cluster resource group, you can expand the assets under a controller, or you can select a controller to view its detailed information.

The name of the controller's computer is shown, as well as the driver name, instance number, bus number, address, and type of driver (for example, SCSI, IDE, etc).

Note: Host Bus Adapter (HBA) controllers must be connected to the SAN and see volumes in order to be displayed under the Controllers node.

Disk Information

As you drill down into Asset Reporting By Cluster for a cluster node or cluster resource group, you can expand the assets under a disk, or you can select a disk to view its detailed information.

The disk detail information is shown in a tab dialog that is made of up three pages: **General**, **Paths**, and **Latest Probe**.

General page

Use the General page to view the computer name, the path name, SCSI Target ID, logical unit number, and the number of access paths. The disk information shown includes the hardware manufacturer of the disk, the model number of the disk, the firmware revision, the disk serial number, and the manufacture date along with the time when Data Manager discovered the disk. In the bottom area of the page is a breakdown of how the disk is allocated, including the disk allocation point and the length of the allocation.

Latest Probe page

Use the Latest Probe page to view the information that was obtained by Data Manager about this disk during the most recent probe of that disk. This information on this page includes:

- Sector size and sector information
- the number of heads
- the number of cylinders
- the logical block size
- the capacity of the disk
- the RPM or revolutions per minute that the disk spins at
- the elapsed time the disk has been powered on
- spare blocks
- the current write protection setting
- the cache enabled setting
- whether failure is predicted for this disk
- the time and date of the last Probe
- the unallocated space
- the number of primary defects
- the number of grown defects
- the NonMedium errors
- the recorded operating temperature

Some of these values are manufacturer specific, so if they are unavailable, then you will see **N/A** for a value. Additional information about the breakdown of recovered errors, unrecovered errors, and the number of bytes broken down by Write, Read, and Verify is shown in this report. Also included are three SCSI disk bit settings that you can turn on/off and the change will actually be made to the disk settings the next time a probe is run on that disk. These include on/off settings for:

- Automatic read reallocation
- Automatic write reallocation
- Read continuous (no error correction)

Probe History page

Use the Probe History page to view the history of probes that have

been run on the disk you are reporting on. Each row on this report represents one time that a probe has been run against the disk. Each row shows what time the probe ran, the type of record (which could be daily, weekly or monthly), bytes written, bytes read, number of grown defects.

Filesystem and Logical Volume Information

The Asset Reporting By Cluster Filesystem or Logical Volume information report contains two pages: **Filesystem** and **Logical Volume**.

Filesystem page

Use this page to view a pie chart that depicts used space and free space on the file system and includes the following information:

- Computer name
- File system type
- Use count
- Mount point
- Physical size of the file system
- Capacity of the file system
- The time and date of the last probe and scan
- When the file system was discovered on the system by Data Manager
- If the file system has been removed then a removed date/time will be shown

Logical Volume page

Use this page to view:

- The path name of the logical volume
- The capacity of the logical volume
- The use count
- The type of logical volume
- Whether it is a swap volume and mirrored
- A tree displaying the mapping of the logical volume on the disk area(s) that make up the logical volume

Exports or Shares Information

The Asset Reporting By Computer Exports or Shares information report includes the following information:

- the name of the computer the share is on
- the export name that it will be known by
- the path on the computer the share resides upon
- the protocol being used and the date and time that Data Manager discovered the export or share

Monitored Directory Information

To view a report on a directory that you are monitoring, expand the **Monitored Directories** node on the navigation tree and select the directory on which you want to report. This information is available for clustered nodes, but not cluster resource groups.

The Asset Reporting By Cluster Monitored Directory report includes the following information:

- computer name the directory resides on
- name of the directory
- owner
- OS Group

- total size of the directory
- number of files in the directory
- count of subdirectories in the directory
- average file size
- access time
- modification time and creation time of the directory
- last scan time and the time that Data Manager discovered the directory
- file attributes that are set
- list of the Directory Groups that the directory is a part of

By Computer:

Use the By Computer reports to drill down into a computer and view detailed information about its controllers, disks, file systems, logical volumes, and directories.

Computer Information

As you drill down into Asset Reporting By Computer, you can expand the assets under a computer, or you can select a computer to view the details for that computer. The detailed report on the right shows the computer's:

- Computer
- Host ID: the identifier of the host machine in hexadecimal format. Currently, IBM Tivoli Storage Productivity Center collects Host ID information for Solaris machines only. All other platforms will display GUID information in the Host ID column.
- Group and Domain
- Network Address, IP address
- Time zone
- Hardware manufacturer and model numbers
- Processor type
- Count and speed
- Amount of RAM
- Operating System (OS) information such as:
  - type of OS and version
  - the storage attributes of the server
  - the computer's disk capacity, unallocated space, and free space

Controller Information

As you drill down into Asset Reporting By Computer, you can expand the assets under a controller, or you can select a controller to view the details for that controller. The name of the computer's controller is shown, as well as the driver name, instance number, bus number, address, and type of driver (for example, SCSI, IDE, and so on).

Note: Host Bus Adapter (HBA) controllers must be connected to the SAN and see volumes in order to be displayed under the Controllers subnode.

Disk Information

As you drill down into Asset Reporting By Computer, you can expand the assets under a disk, or you can select a disk to view the details for that disk. The disk detail information is shown in a tab dialog that is made of up four tabs: **General**, **Paths**, **Latest Probe**, and **Probe History**.

General page

The General page includes the computer name, the path name, SCSI Target ID, logical unit number, the number of access paths.

The disk information shown includes the hardware manufacturer of the disk, the model number of the disk, the firmware revision, the disk serial number, and the manufacture date along with the time when Data Manager discovered the disk.

In the bottom area of the tab is a breakdown of how the disk is allocated, including the disk allocation point and the length of the allocation.

Paths The **Paths** page shows the following information:

- host
- OS type
- path
- controller
- instance
- bus number
- SCSI target id
- logical unit number

Latest Probe page

The **Latest Probe** page shows the information that was obtained by Data Manager about this disk during the most recent probe of that disk. This information on this tab includes:

- sector size and sector information
- the number of heads
- the number of cylinders
- the logical block size
- the capacity of the disk
- the RPM or revolutions per minute that the disk spins at
- the elapsed time the disk has been powered on
- spare blocks
- the current write protection setting
- the cache enabled setting
- whether failure is predicted for this disk
- the time and date of the last Probe
- the unallocated space
- the number of primary defects
- the number of grown defects
- the NonMedium errors
- the recorded operating temperature

Some of these values are manufacturer specific, so if they are unavailable, then you will see **N/A** for a value.

Additional information about the breakdown of recovered errors, unrecovered errors, and the number of bytes broken down by Write, Read, and Verify is shown in this report.

Also included are three SCSI disk bit settings that you can turn on/off and the change will actually be made to the disk settings the next time a probe is run on that disk. These include on/off settings for:

- Automatic read reallocation
- Automatic write reallocation
- Read continuous (no error correction)

Probe History page

The **Probe History** page shows the history of probes that have been run on the disk you are reporting on. Each row on this report represents one time that a probe has been run against the disk. Each row shows what time the probe ran, the type of record (which could be daily, weekly or monthly), bytes written, bytes read, number of grown defects.

Volume Group Information

Data Manager provides support for the following volume managers:

- Veritas Volume Manager (Sun Solaris)
- Logical Volume Manager (AIX)

These volume managers allow you to group logical volumes and disks. As you drill down into Asset Reporting By Computer for an AIX or Sun Solaris machine, you can expand the assets under a volume group to see the group's disks and file systems or logical volumes.

You can also select a volume group to view detailed information about it. The group's name is shown, as well as the type (LVM or Veritas), capacity, freespace, number of volumes, number of disks, discovered time, and probe time.

File System / Logical Volume Information

The Asset Reporting By Computer File System / Logical Volume information report contains two pages: **File System** and **Logical Volume**.

File System page

This page shows a pie chart depicting used space and free space on the file system and includes the following information:

- computer name
- file system type
- use count
- mount point
- physical size of the file system
- capacity of the file system
- the time and date of the last probe and scan
- when the file system was discovered on the system by Data Manager
- if the file system has been removed, then a removed date/time will be shown

Also shown for UNIX based systems are:

- the maximum file count
- used inodes
- free inodes. Inodes are areas reserved for index entries within the UNIX file systems. For Windows machines, these values will be zero.
- the amount of used space and free space is shown

Logical Volume page

This information on this page includes:

- the path name of the logical volume
- the capacity of the logical volume
- the use count
- the type of logical volume
- whether it is a swap volume and mirrored

- a tree displaying the mapping of the logical volume on the disk areas that make up the logical volume

Exports/Shares Information

The Asset Reporting By Computer Exports/Shares information report includes the following information:

- the name of the computer the share is on
- the export name that it will be known by
- the path on the computer the share resides upon
- the protocol being used and the date/time that Data Manager discovered the export/share

Monitored Directory Information

To view a report on a directory that you are monitoring, expand the **Monitored Directories** node on the navigation tree and select the directory on which you want to report.

Note: This information is not shown for cluster resource groups.
The Asset Reporting By Computer Monitored Directory report includes the following information:

- computer name the directory resides on
- name of the directory
- owner
- OS Group
- total size of the directory
- number of files in the directory
- count of subdirectories in the directory
- average file size
- access time
- modification time and creation time of the directory
- last scan time and the time that Data Manager discovered the directory
- file attributes that are set
- list of the Directory Groups that the directory is a part of

The Asset Reporting By Computer Monitored Directory report includes the following information:

- computer name the directory resides on
- name of the directory
- owner
- OS Group
- total size of the directory
- number of files in the directory
- count of subdirectories in the directory
- average file size
- access time
- modification time and creation time of the directory
- last scan time and the time that Data Manager discovered the directory
- file attributes that are set
- list of the Directory Groups that the directory is a part of

By OS Type:

Use the By OS Type Asset report to view your storage assets based upon the operating systems under which they are running. This provides you with a an easy way to see what your OS dependencies are throughout your organization. You can drill down into any of the storage resources that appear under an operating system node to access all of the same sub-reports that as described for other asset reports.

By Storage Subsystem:

Use the By Storage Subsystem reports to drill down through your storage subsystem assets in a hierarchical manner. Data Manager provides reporting for individual storage subsystems (for example: IBM ESS, IBM FAStT, HP StorageWorks, IBM Tivoli Storage DS6000, IBM XIV Storage System, IBM Tivoli Storage DS8000, and Hitachi Data Systems), any disk array subsystems whose SMI-S Providers are CTP certified by SNIA for SMI-S 1.0.2 or later, and IBM SAN Volume Controller clusters.

For Disk Array subsystems, you can view information about:

- Disk Groups
- Array Sites
- Ranks
- Storage Pools
- Disks
- Volumes

For Virtualization systems, you can view information about:

- Managed Disk Groups
- Managed Disks
- Virtual Disks
- Virtual Disk Copies

Keep in mind the following considerations:

- The total capacity value for storage pools in Tivoli Storage Enterprise Storage Server subsystems includes the formatted (usable) capacity for storage pools only.
- For IBM SAN Volume Controller clusters, you can drill down into storage subsystems and expose managed disks, managed disk groups, and virtual disks.
  - Information displayed for managed disks is the same as for storage subsystem disks. See "By Storage Subsystem" for a description of the available views.
  - Information displayed for virtual disks is the same as for storage subsystem volumes. See "By Storage Subsystem" for a description of the available views.
- For IBM FAStT, HP StorageWorks, and IBM XIV Storage System you can drill down into storage subsystems and expose storage pools, disks, and volumes.

  Note: The total capacity value for storage pools in these storage subsystems includes raw (unconfigured) capacity—it does not include formatted/usable capacity.
- For Hitachi Data Systems TagmaStore subsystem, the relationship between a disk or volume is not displayed if the volume is using imported extents (virtual disks).

  Note: The total capacity value for storage pools in HDS subsystems includes the formatted/usable capacity for storage pools only.
- Some of these reports will contain inaccurate information if the following configuration is used:
  - An IBM SAN Volume Controller is used to virtualize a disk array subsystem. In this situation, disks and volumes appear to belong to both the IBM SAN

Volume Controller and the disk array subsystem, and is counted twice in system-wide reports. To avoid this situation, set up CIMOM access for either the IBM SAN Volume Controller or for the disk array subsystems that it virtualizes.

– You have configured Data Manager to access the CIMOMs of both an IBM SAN Volume Controller and the disk array subsystems.

Disk Groups

When you expand the **Disk Group** node, the navigation tree displays a list of disk groups for the storage subsystem. You can expand the name of a disk group to view the disks and volume spaces within that disk group. Drill down through both Volume Spaces and Disks subnodes to view:

- the disks and volumes associated with a volume space
- the volumes with segments on a disk
- the disks containing segments from a volume
- detailed information about a disk

Array Sites

(IBM Tivoli Storage DS8000 only) When you expand the **Array Site** node, the navigation tree displays a list of array sites for the storage subsystem. Click **Array Site** to view a description of the information displayed about array sites. You can expand the name of an array site to view the ranks, storage pools, and disks associated with that array site.

Storage Pools

When you expand the **Storage Pool** node, the navigation tree displays a list of storage pools for the storage subsystem. You can expand the name of a storage pool to view the disks and volumes within that storage pool, including:

- the volumes with segments on a disk
- the disks containing segments from a volume
- detailed information about a disk

Disks When you expand the **Disks** node, the navigation tree displays a list of disks for the storage subsystem. You can expand the name of disk to view the volumes associated with that disk. Click the name of a volume associated with a disk to view detailed information about that volume.

Volume Spaces

When you expand the **Volume Spaces** node, the navigation tree displays a list of volume spaces for the storage subsystem. You can expand the name of a volume space to view the disks and volumes within that volume space. Drill down through both the Disks and volumes subnodes to view:

- the volumes with segments on a disk
- the disks containing segments from a volume
- detailed information about a disk
- detailed information about a volume

Volumes

When you expand the volumes node, the navigation tree displays a list of volumes for the storage subsystem. Click the name of a volume to view detailed information about that volume.

Note: Volume Reporting is not available for IBM FAStT FlashCopy logical drive and repository logical drive.

Managed Disk Groups

When you expand the Managed Disk Group node, the navigation tree displays a list of disk groups that are managed by the IBM SAN Volume Controller cluster. You can expand the name of a disk group to view its physical disks and virtual disks.

Managed Disks

When you expand the Managed Disks node, the navigation tree displays a list of disks managed by the IBM SAN Volume Controller cluster. You can expand the name of a disk to view the virtual disks associated with it.

Virtual Disks and Virtual Disk Copies

There are three nodes in the navigation tree that are related to virtual disks (VDisks) and might appear under nodes for the IBM SAN Volume Controllers in your environment:

- **Virtual Disks** (at the same level as **Managed Disks** and **Managed Disk Groups**): Expand this node to view a list of the logical disks that a cluster provides to the SAN. You can further expand each virtual disk (VDisk) to view its mirror copies (primary copies and additional copies).
- **Virtual Disks** (under **Managed Disks** and **Managed Disk Groups**): Expand this node to view the VDisks for which the primary copies are on the selected Managed Disk.
- **Virtual Disk Copies** (under **Managed Disks** and **Managed Disk Groups**): Expand this node to view all the virtual disk copies that are on the selected Managed Disk. The primary copies are not listed under this node because they are displayed under the **Virtual Disks** node.

  Note:

  – The **Virtual Disk Copies** node appears in the navigation tree only if a monitored IBM SAN Volume Controller supports virtual disk mirrors. If it supports virtual disk mirrors, all virtual disks will have a related primary virtual disk copy regardless of whether the virtual disk is mirrored. This is because all virtual disks on an IBM SAN Volume Controller that supports virtual disk mirrors are primary copies that can be updated to have actual mirrored disks or virtual disk copies.

  – The icon for all virtual disk copies appear as folders in the navigation tree because they represent logical extensions of a virtual disk.

System-wide:

Asset reporting System-wide enables you to do the following:

- View your storage assets as a whole, across the entire system. This lets you get lists of all of your disks, computers, storage subsystems, disk/volume groups, disks, file systems, logical volumes, LUNs, exports/shares, monitored directories, users, OS users, etc. Additionally, you can view aggregated reports based upon certain filters, such as grown defects per disk.
- View detailed information about Data agents. Use these reports to monitor your agents and keep track of their status.
- The Selection button is disabled for System-wide reports because all the monitored storage entities are automatically displayed for each report type. For example, all the computers monitored by IBM Tivoli Storage Productivity Center are displayed in the **System-wide > Computers** report.

Agents report

The **System-wide → Agents** report enables you to track and monitor the

Data agents deployed in your storage environment. With this report, you can export a comma separated file containing the name and status of an agent. This report enables you to:

- view information about the agents deployed in your environment
- perform actions against one or more agents shown in the report

Each Data agent that appears in the report has a drill down icon beside next to it which enables you to display detailed status for the agent. The drill down icon will display the same agent detail that is displayed when you click on **Administrative Services** → **Agents** → **Data** node in the navigation tree.

You can also select one or more rows on the report and right-click to display a pop-up menu which enables you to perform an action against multiple agents at the same time. The pop-up menu will be customized to offer appropriate options depending on whether a single row or multiple rows are selected and also depending on the status of the agent(s) selected. The pop-up menu contains all the other operations available from the **Administrative Services** → **Agents** → **Data** → **<agent_name>** node in the navigation tree:

Pop-up Menu Action	Description
View Log	View the log file for the agent. The log file contains informational, warning, and error messages for the previous sessions of the agent. Use the content of the log file to troubleshoot any errors that might occur during agent startup, processing, or shutdown. By default, information from the last 5 sessions of the agent appears in the log window. A session starts when the agent is started and ends when the agents goes down. You can increase or decrease how many sessions for the agent are kept in the history by changing the value for the `logFilesKept` parameter in the `agent.config` file. Log files for an agent are located in the agent install directory. The default agent install directory for agent log files is: • (Windows) `\program files\IBM\TPC\ca\subagents\TPC\Data\log\computername\` • (UNIX, Linux) `/opt/IBM/TPC/ca/subagents/TPC/Data/log/computername/` The `agent.config` file is located in the agent install directory. The default agent install directory for the `agent.config` file is: • (Windows) `\program files\IBM\TPC\ca\subagents\TPC\Data\config\` • (UNIX, Linux) `/opt/IBM/TPC/ca/subagents/TPC/Data/config/`
Configure Tracing	View or modify agent tracing configuration.
Read Config	Have an agent read the `agent.config` file. The `agent.config` file contains parameters that you can change to customize the behavior of an agent. Select this option for an agent if you have made changes to the `agent.config` file and you want that agent to use those changes.

Pop-up Menu Action	Description
Check	- Check to: 1) determine if an agent is up or down, and 2) determine if the version of the agent is out of date. If the agent needs to be upgraded, clicking this option will automatically install the required upgrades. **Note:** A check will also enable an agent that is currently disabled.
Shutdown	Shutdown this agent. You can choose to shutdown in the following ways: • **Normal**: Shut down the agent software and let all running processes complete. The agent will continue to accept new jobs that are submitted and won't shut down until 1) all running processes are complete and 2) there are no new jobs submitted. • **Immediate**: Shut down the agent software when the currently running processes complete. No new jobs will be accepted by the agent and shutdown will occur immediately after the last job completes. • **Abort**: Shutdown the agent software and stop whatever processes are currently running.
Disable	• **Disable**: Select this option to designate that an agent is ″unavailable″ and to indicate that the server should not contact the agent for any job processing. Use this option depending on the situation within your environment. For example, you might want to disable an agent when: – The agent machine is undergoing maintenance and will be unavailable. This prevents the server from flagging the agent as ″down″ if it cannot reach the agent a certain amount of times as defined by the `agentErrorLimit` parameter in the `server.config` file. – The agent machine is performing some resource-intensive processing and you want do not want to add any Data Manager jobs to that processing load. While disabled, the agent will appear with a red outlined circle with a line through it in the user interface and the server will not attempt to contact the agent.
Enable	Select this option to re-enable an agent that is currently disabled. Once enabled, the server will resume communication with the agent. **Note:** If the server cannot contact an agent equal to the number defined for the `agentErrorLimit` parameter in the `server.config` (the default is 3), the server will automatically flag the server as ″down″. Select the **Enable** option to reestablish communication between the agent and server and reset the `agentErrorLimit` ″counter″ to `0`. The default directory for the `server.config` file is: (Windows) `\program files\IBM\TPC\Data\config` (UNIX, Linux) `/opt/IBM/TPC/Data/config` **Note:** A **Check** action will also enable an agent that is currently disabled.
Delete	Erase agent from repository

Computers report

Use this report allows you to see all the computers in the system, sorted in order of the time they were last booted, probed, or discovered. This is valuable in environments where computers become more unreliable after they have been up and running for long periods of time.

Storage Subsystem report
Use this report to view all your monitored storage subsystems and SAN Volume Controllers, sorted by their total capacity (By Capacity) or alias (By Storage Subsystem).

Unmanaged Virtual Machines report
Use this report to view detailed information about the unmanaged virtual machines in your environment. Each row in the report gives you information for each unmanaged virtual machine.

Virtual Machines With No VMware report
Use this report to view detailed information about the virtual machines that have Data agents installed but are not discovered through the VMWare data sources. Each row in the report gives you information for each virtual machine.

Disk/Volume Groups report
Data Manager provides support for the following volume managers:

- Veritas Volume Manager (Sun Solaris)
- Logical Volume Manager (AIX)

These volume managers allow you to create groups of logical volumes and disks. Use this report to view all your disk and volume groups, sorted by their computer (By Computer), storage capacity (By Capacity) total freespace (By Freespace), time when they were last probed (By Probe Time), and time when they were discovered (By Discovered Time).

Disks report
Use this report to view all the disks in the system, sorted in order of the amount of unallocated space they have on them (By Unallocated Space), how many recovered and unrecovered errors they have, time when they were probed (By Probe Time), and time when they were discovered (By Discovered Time).

The **Asset Reporting → System-wide → Disks by Grown Defects** report enables you to view the following:

- All the disks in the system sorted in order of the number of grown defects they have on them
- What disks in your system are experiencing problems and may be in danger of an imminent failure

The **Asset Reporting → System-Wide → Disks by Primary Defects** report will return no data if disk defects are not detected. This report enables you to view the following:

- All the disks in the system sorted in order of the number of primary defects they have on them
- What disks in your system are in good shape and which ones are more or less defective

File systems / Logical Volumes reports
Use the **Asset Reporting → System-Wide → File Systems** or **Logical Volumes by Free Space** reports to view the following:

- All the file systems in the system sorted in order of the amount of free space they have on them

- Which file systems have less free space and may be in danger of running out of free space
- Which file systems have perhaps too much free space and where space consolidation may be necessary

Volumes reports

Use the **Asset Reporting → System-Wide → Volumes** report to view the following:

- Information about the storage subsystems that host volumes
- Storage capacity of the volumes hosted by storage subsystems
- Information about the storage subsystems that host volumes sorted by the date/time when they were probed

Exports or Shares reports

Use this report to view all the Exports or Shares in your environment, sorted in the order that they were discovered.

Monitored Directories reports

Use this report to view all the monitored directories in the system, sorted in order of the amount of storage space they are taking up (By Space Used), the time when they were last scanned (By Scan Time), the time when they were discovered (By Discovered Time), and when they were removed (Removed Directories).

Unmanaged Computers reports

Use this report to view all of the new computers that have been introduced into your system and that do not have Data agents installed on them. It also lists other computers that you have not put under Data Manager control.

Users reports

Use this report to view all the users in your environment sorted in the order that they were discovered.

OS User Groups reports

Use this report to view all the OS User Groups in your environment sorted in the order when they were discovered.

Availability Reports

Use Availability reports to monitor and report on the availability of your storage, both from a network point of view and from a computer uptime perspective.

These reports enable you to see what percentage of the time the storage in your enterprise (or on particular server or group of servers), is off-network due to network problems or perhaps is unavailable due to system downtime. In today's fast-paced business environment, 24x7 application availability is critical, the Availability reports provide a powerful tool for monitoring the availability of the storage and data that back up your applications.

Reporting Type	Description
Ping	Use the ping report to monitor how available your storage is on the network. To do this, regularly scheduled pings are sent out from the Data server to all of the monitored Data agent machines in your environment. The Data server checks not only for whether the Data agent is available on the network, but records the amount of time that the request was in transit on the network. This allows you to get an idea of the responsiveness of the various storage servers on the network as well. Notice that when you generate a ping report the **Report Filter Specification** section on the **Selection** page enables you to limit the time range to use in generating the report. If you would like to limit the time range, you can check the check box next to **Limit To This Time Range** and then use the date selection boxes to specify the range for a report A ping report shows: • the name of a computer • percentage of time the computer is available on the network • number of pings that have been transmitted and the number that have been received • number of damaged requests and the average transit time for the request to go between the Server and the monitored agent machine that is being monitored • maximum transit time • minimum transit time • report start time and report end time for each computer in the report **Note:** You can sort on any of these columns by simply clicking on the column name at the top of the report. You can also generate charts and drill down on each computer for more detail.
Computer Uptime	Use the computer uptime report to track the availability of the storage based upon what percentage of the time the server that hosts that storage is down or not running. This is accomplished by the Data agent component that runs on the monitored machine. The agent keeps track of this information and it is periodically sent to the server. A computer uptime shows the: • name of a computer • percentage of time the computer is up and running • number of reboots that have been recorded • total amount of time the computer has been down and unavailable as well as the average down time that has been recorded • report start time and report end time for each computer in the report **Note:** You can sort on any of these columns by simply clicking on the column name at the top of the report. You can also generate charts and drill down on each computer for more detail.

TPC-wide Storage Space Reports

Use these reports to view storage capacity of your resources at a file system, computer, or entire network level.

For a specific computer, group of computers, or all the computers in your enterprise, you can view the following:

- How much storage you have
- How much storage is being used
- What percentage of storage is free for expansion

Reporting Type	Description
Disk Space	Use these reports to view and chart disk capacity, per disk, per computer, per computer group, per domain, and for the whole network. These reports include one row for each computer that shows the total amount of storage capacity on that computer and the associated computer and disk storage information. For detail on any computer in the report, click on the magnifying glass on the left of the report row. You can also click on the chart button within a report to view an easy-to-interpret bar chart, showing the total storage capacity of each computer that you have chosen to report on and chart. To get the exact amount of storage for a computer, simply position the cursor over the associated bar and a box will appear showing the total amount for the selected computer. **Note:** To ensure the accuracy and consistency of the data that appears in Disk Capacity reports, run regularly scheduled probes against the hosts that import volumes and the storage subsystems upon which you want to report. You should run probes against storage subsystems after running them against any hosts. You can use the following By clauses to view disk space for the following storage entities: By Disk, By Storage Subsystem, By Cluster, By Computer, By Computer Group, By Domain, and Network-wide.
File System Space	Use these reports to view information for monitored file systems, including the total amount of storage capacity on those file systems and statistics for used space, free space, and total number of files. You can use the following By clauses to view file system space for the following storage entities: By Filesystem, By Filesystem Group, By Cluster, By Computer, By Computer Group, By Domain, Network-wide.
Consumed File System Space	Use these reports to view information about the amount of space consumed on a filesystem. You can use the following By clauses to view consumed file system space for the following storage entities: By Filesystem, By Filesystem Group, By Cluster, By Computer, By Computer Group, By Domain, Network-wide.
Available File System Space	Use Filesystem Used Space reports to view information about the amount of free space on a filesystem. You can use the following By clauses to available file system space for the following storage entities: By Filesystem, By Filesystem Group, By Cluster, By Computer, By Computer Group, By Domain, Network-wide.

TPC-Wide Storage Space - Disk Space report:

Use this report to view and chart disk storage capacity at the following levels: by disk, by computer, by cluster, by storage subsystem, by computer monitoring or reporting group, by storage subsystem monitoring or reporting group, by domain, and for the whole network.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

By Disk

Chart icon 
Click this button to generate charts about the capacity of disks that are displayed in the report.

Magnifying glass icon 
Click this button to view a more detailed report about the capacity of a disk.

Device
Name of the device on which a disk resides. The following disks are displayed in this report:

- Computer disks that are internal to the computers monitored by IBM Tivoli Storage Productivity Center.
- Computer disks on virtual machines that are monitored by Tivoli Storage Productivity Center.
- Storage subsystem disks on storage subsystems that are monitored by Tivoli Storage Productivity Center.
- Storage subsystem disks that no longer exist after a probe of a storage subsystem because information about that subsystem was removed based on settings in the Removed Resource Retention panel.
- Back-end storage volumes that are not identified by Tivoli Storage Productivity Center as being assigned to a monitored storage subsystem virtualizer.
- Computer disks that are not identified by Tivoli Storage Productivity Center as being assigned to a monitored storage subsystem.

The following disks are not displayed in this report:

- Computer disks that are assigned to storage subsystems monitored by Tivoli Storage Productivity Center.
- Back-end storage volumes that are assigned to storage subsystem virtualizers monitored by Tivoli Storage Productivity Center.

Path Name of a disk.

Disk Space (Capacity)
Total storage capacity for a disk.

Note: The value that appears in the TOTAL row for this column will not include computer disk space that is on virtual machines monitored by Tivoli Storage Productivity Center.

Available Disk Space (Unallocated Space)
Amount of space on the disk that is not allocated to any logical volume (or volumes for disks within storage subsystems). This value does not include RAID overhead on storage subsystems that are identified by Tivoli Storage Productivity Center.

Note: The value that appears in the TOTAL row for this column includes the unallocated capacity of all disks within the storage subsystem, including the capacity of spare disks.

In Storage Subsystem
Indicates **Yes** when a disk is either of the following:

- An actual physical disk within a storage subsystem that is monitored by Tivoli Storage Productivity Center.
- A back-end storage volume within a storage subsystem virtualizer that is monitored by Tivoli Storage Productivity Center.

Manufacturer
Name of the manufacturer of a disk. Name of the manufacturer of a disk.

Model
Model number of a disk.

Serial Number
Serial number of a disk.

Primary Defects
Number of defects on the disk when it was new.

Note: Defect data is not collected for IBM Tivoli Storage Enterprise Storage Server, IBM FAStT, HP StorageWorks, Hitachi Data Systems storage subsystems.

Grown Defects
Number of defects detected on the disk since it was new.

Note: Defect data is not collected for IBM Tivoli Storage Enterprise Storage Server, IBM FAStT, HP StorageWorks, Hitachi Data Systems storage subsystems.

Fibre Attached
Indicates **Yes** when a disk is attached to a port that Tivoli Storage Productivity Center identifies as fibre.

Mapped to Storage Subsystem
Indicates **Yes** when a computer disk or a back-end storage volume can be mapped to a storage subsystem.

Detectable
Indicates **Yes** when a disk is identified as being available to the server or the storage subsystem following a successful probe.

Is Virtual
Indicates **Yes** when a disk resides on a machine that Tivoli Storage Productivity Center identifies as a virtual machine.

By Storage Subsystem

Chart icon 
Click this button to generate charts about the disk capacity of storage subsystems that are displayed in the report.

Magnifying glass icon 
Click this button to view a more detailed report about the capacity of disks associated with a storage subsystem.

Name
Name or alias of a storage subsystem that contains the disks or the name or alias of a storage subsystem virtualizer (for example, SAN Volume Controller) that manages the disks. This report does not display storage subsystems that have been discovered but not probed.

Disk Space (Disk Capacity)
Total raw (unformatted) disk capacity of a storage subsystem. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Capacity of spare disks identified on IBM Tivoli Storage Enterprise Storage Server, IBM DS8000, and IBM DS6000 storage subsystems.
- Capacity of storage subsystem disks that become missing after a Tivoli Storage Productivity Center storage subsystem probe.

Available Disk Space (Disk Unallocated Space)
Total unformatted disk freespace of a storage subsystem. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- RAID overhead on storage subsystems that are identified by Tivoli Storage Productivity Center.
- External back-end storage for TagmaStore subsystems.
- Storage space that is consumed by volumes that reside on external disks when Hitachi Data Systems (HDS) storage subsystems are configured as storage virtualizers.

Number of Disks
Total number of disks in a storage subsystem. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- External back-end storage for TagmaStore subsystems.

Note: If you click the magnifying glass icon next to storage system, the resulting drill-down report generates a list of detectable storage subsystem disks or back-end storage volumes only. Therefore, the total number of disks for each storage subsystem drill-down report will equal this value.

Number of Volumes (LUNs)
Total number of storage volumes in a storage subsystem or the total number of virtual volumes that are virtual storage. This value does not include storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.

Volume Space (LUN Capacity)
Total capacity of all storage subsystem volumes or virtual storage volumes within a storage subsystem. This value does not include capacity information for storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.

Note: For space-efficient FlashCopy volumes, the volume capacity is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage capacity reported for a storage subsystem using space-efficient FlashCopy volumes. This also applies to the Total Usable Volume Space and Total FlashCopy Target Capacity calculations.

Consumable Volume Space (Usable LUN Capacity)
Total amount of unique storage subsystem volume space and virtual

storage volume space on monitored storage subsystem arrays that you can assign to servers within the network or are already assigned to servers within the network.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a probe.
- Storage subsystem volumes or virtual storage volumes that are used as flash copy targets.
- Storage subsystem volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers. For example, SAN Volume Controller.

Note: For IBM FAStT, HP StorageWorks, and Hitachi Data Systems storage subsystems, the value in this column is the same as the value in the Volume Capacity column.

FlashCopy Target Volume Space (Capacity)
Total amount of unique storage subsystem volume space and virtual storage volume space on monitored storage subsystem arrays that is identified as flash copy target space. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are not used as flash copy targets.

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Total amount of storage subsystem volume space or virtual storage volume space on systems where the systems disks are identified by Tivoli Storage Productivity Center as residing within a monitored storage subsystem. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Disks that are assigned to a monitored storage subsystem but are not identified by Tivoli Storage Productivity Center as a disk residing within the storage subsystem. Some reasons why these disks are not identified as residing within the storage subsystem include:
  - Tivoli Storage Productivity Center does not support the identification of disks for some storage subsystem types.
  - The connection of the disk to the storage subsystem is lost and is no longer detectable during a probe.
- Storage subsystem volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.

- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers. For example, SAN Volume Controller.

Volume Space Not Correlated (LUN Capacity not Visible to Data Manager)
Total amount of unique storage subsystem volume space and virtual storage volume space discovered on monitored storage subsystem arrays that can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Disks that reside within a monitored storage subsystem.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers.

Unavailable Disk Space (Overhead)
Amount of storage subsystem volume or virtual storage volume space that is dedicated to redundancy. This value is dependent on the storage subsystem SMI-S provider returning a valid RAID value that is used to determine the overhead. If the value is zero for a storage subsystem, the overhead cannot be calculated.

This amount does not include storage space information from storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.

Formatted Space (Capacity)
Total amount of formatted storage space associated with the storage subsystem that is used or can be used for subsystem storage volumes or virtual storage volumes. This value is dependent on the storage subsystem SMI-S provider returning a formatted value.

For some storage subsystems, this is the total managed space which includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Note: This value does not apply to IBM Storage Volume Controller.

Formatted Space with No Volumes
Total amount of formatted storage space associated with the storage subsystem that can be used for subsystem storage volumes or virtual storage volumes. This value is dependent on the storage subsystem SMI-S provider returning a formatted value.

For some subsystems, this value represents the remaining managed space that includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Note: This value does not apply to SAN Volume Controller.

Overall Unavailable Disk Space (Total Overhead)
(IBM Tivoli Storage Enterprise Storage Server, IBM DS6000, and IBM DS8000 storage subsystems only) Total amount of unformatted space within disk groups or array sites.

Available Disk Group or Array Site Space
This value does not apply to all storage subsystems. For IBM FAStT, IBM Tivoli Storage Enterprise Storage Server, IBM DS6000, and IBM DS8000 storage subsystems, this is the total raw disk space of any unformatted disk groups or array sites. This column displays **N/A** for the storage subsystems where this value is not applicable. For IBM XIV Storage System this value shows the physical (hard) space of the XIV system available for pool creation or expansion.

Cache
Amount of cache memory that is internal to the storage subsystem.

Note: This column is blank for SAN Volume Controllers.

Manufacturer
Name of the manufacturer of a disk.

Model
Model number of a disk.

Serial Number
Serial number of a disk.

Firmware Revision
Microcode (OS) version.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage subsystem volume space within the storage subsystem that is mapped or assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are not mapped or are not assigned to host systems.

Unassigned Volume Space (Subsystem Unassigned LUN Capacity)
Total storage subsystem volume space within a storage subsystem that is not mapped or not assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are mapped or assigned to host systems.

Unassigned and Available Formatted Disk Space (Subsystem Total Formatted Free Space)
Total storage subsystem space that is not mapped or not assigned to host systems and the total amount of formatted storage space associated with the storage subsystem that can be used for subsystem storage volumes.

This total amount of formatted storage space is dependent on the storage subsystem SMI-S provider returning a formatted value. For some subsystems, this value represents the remaining managed space that might include space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are mapped or assigned to host systems.

zOS Volume Space (zOS LUN Capacity)
Total storage subsystem space that is reserved for use on mainframe storage such as used by zOS. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are reserved for use on open system server or non-mainframe storage.

Open System Volume Space (Open System LUN Capacity)
Total storage subsystem space that is reserved for use on open system storage such as used by operating systems like LINUX, AIX, HP/UX, SUN Solaris, and Windows. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are reserved for use on mainframe storage.

Physical Disk Space (Capacity)
Total amount of physical disk space discovered on the monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller) that are attached from a storage subsystem.

Backend Volume Space (Backend LUN Capacity)
Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage subsystem virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are back-end disks on a storage subsystem virtualizer that is not monitored.

By Storage Subsystem Group

Chart icon 
Click this button to generate charts about the disk capacity of storage subsystems in the groups that are displayed in the report.

Magnifying glass icon 
Click this button to view a more detailed report about the disk capacity of storage subsystems in a group.

Creator
ID of the user who created the storage-subsystem montoring group or storage subsystem reporting group.

Group Name of a group.

Disk Space (Disk Capacity)
Total raw (unformatted) disk capacity of the storage subsystems that are in a group. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Capacity of spare disks identified on IBM Tivoli Storage Enterprise Storage Server, IBM DS8000, and IBM DS6000 storage subsystems.
- Capacity of storage subsystem disks that become missing after a Tivoli Storage Productivity Center storage subsystem probe.

Available Disk Space (Disk Unallocated Space)
Total unformatted disk free space of the storage subsystems that are in a group. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- RAID overhead on storage subsystems that are identified by Tivoli Storage Productivity Center.
- External back-end storage for TagmaStore subsystems.
- Storage space that is consumed by volumes that reside on external disks when Hitachi Data Systems (HDS) storage subsystems are configured as storage virtualizers.

Number of Disks
Total number of disks in the storage subsystems that are in a group. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- External back-end storage for TagmaStore subsystems.

Note: If you click the magnifying glass icon next to storage system, the resulting drill-down report generates a list of detectable storage subsystem disks or back-end storage volumes only. Therefore, the total number of disks for each storage subsystem drill-down report will equal this value.

Number of Volumes (LUNs)
Total number of storage volumes in a storage subsystem or the total number of virtual volumes that are virtual storage. This value does not include storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.

Volume Space (LUN Capacity)
Total capacity of all storage subsystem volumes or virtual storage volumes within the storage subsystem. This value does not include capacity information of storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.

Note: For space-efficient FlashCopy volumes, the volume capacity is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage capacity reported for a storage subsystem using space-efficient FlashCopy volumes. This also applies to the Total Usable Volume Space and Total FlashCopy Target Capacity calculations.

Consumable Volume Space (Usable LUN Capacity)
Total amount of unique storage subsystem volume space and virtual storage volume space on monitored storage subsystem arrays in a group that you can assign to servers within the network or are already assigned to servers within the network.

Tivoli Storage Productivity Center does not include the following volums in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a probe.
- Storage subsystem volumes or virtual storage volumes that are used as flash copy targets.
- Storage subsystem volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers. For example, SAN Volume Controller.

Note: For IBM FAStT, IBM XIV Storage System, HP StorageWorks, and Hitachi Data Systems storage subsystems, the value in this column is the same as the value in the Volume Space column.

FlashCopy Target Volume Space (Capacity)
Total amount of unique storage subsystem volume space and virtual storage volume space on monitored storage subsystem arrays in a group that is identified as flash copy target space. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are not used as flash copy targets

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Total amount of storage subsystem volume space or virtual storage volume space on systems where the systems disks are identified by Tivoli Storage Productivity Center as residing within the monitored storage subsystems in a group. Tivoli Storage Productivity Center does not include the following in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Disks that are assigned to a monitored storage subsystem but are not identified by Tivoli Storage Productivity Center as a disk residing within the storage subsystem. Some reasons why these disks are not identified as residing within the storage subsystem include:

- Tivoli Storage Productivity Center does not support the identification of disks for some storage subsystem types.
- The connection of the disk to the storage subsystem is lost and is no longer detectable during a probe.

• Storage subsystem volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
• Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
• Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers. For example, SAN Volume Controller.

Volume Space Not Correlated (LUN Capacity not Visible to Data Manager)
Total amount of unique storage subsystem volume space and virtual storage volume space discovered on monitored storage subsystem arrays in a group that can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following in its calculation of the value for this column:

• Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
• Disks that reside within a monitored storage subsystem.
• Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers.

Unavailable Disk Space
Amount of volume or virtual storage volume space of storage subsystems in a group that is dedicated to redundancy. This value is dependent on the storage subsystem SMI-S provider returning a valid RAID value that is used to determine the overhead. If the value is zero for a storage subsystem, the overhead cannot be calculated.

This amount does not include storage space information from storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.

Formatted Space
Total amount of formatted storage space associated with the storage subsystems in a group that is used or can be used for subsystem storage volumes or virtual storage volumes. This value is dependent on the storage subsystem SMI-S provider returning a formatted value.

For some storage subsystems, this is the total managed space which includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Note: This value does not apply to IBM Storage Volume Controller.

Formatted Space with No Volumes
Total amount of formatted storage space associated with the storage subsystems in a group that can be used for subsystem storage volumes or virtual storage volumes. This value is dependent on the storage subsystem SMI-S provider returning a formatted value.

For some subsystems, this value represents the remaining managed space that includes space that cannot be used for storage volumes or virtual

storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Available Disk Group or Array Site Space
This value does not apply to all storage subsystems. For IBM FAStT, IBM Tivoli Storage Enterprise Storage Server, IBM DS6000, and IBM DS8000 storage subsystems, this is the total raw disk space of any unformatted disk groups or array sites. For XIV Storage System, the total available physical (hard) space that is left for pool creation. This column displays **N/A** for the storage subsystems where this value is not applicable.

Cache
Amount of cache memory that is internal to the storage subsystems in the group.

Note: This column does not include values for SAN Volume Controllers.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total volume space of the storage subsystems within a group that is mapped or assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are not mapped or are not assigned to host systems.

Unassigned Volume Space (Subsystem Unassigned LUN Capacity)
Total volume space of the storage subsystems in a group that is not mapped or not assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are mapped or assigned to host systems.

zOS Volume Space (zOS LUN Capacity)
Total space of storage subsystems in a group that is reserved for use on mainframe storage such as used by zOS. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are reserved for use on open system server or non-mainframe storage.

Open System Volume Space (Open System LUN Capacity)
Total space of the storage subsystems in a group that is reserved for use on open system storage such as used by operating systems such as LINUX, AIX, HP/UX, SUN Solaris, and Windows. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are reserved for use on mainframe storage.

Physical Disk Space (Capacity)
Total amount of physical disk space on the storage subsystem arrays that are in a group. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller) that are attached from a storage subsystem.

Backend Volume Space (Backend LUN Capacity)
Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on the monitored storage subsystem virtualizers (for example, SAN Volume Controller) that are in a group. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are back-end disks on a storage subsystem virtualizer that is not monitored.

By Cluster

Chart icon 
Click this button to generate charts about the disk capacity of the clusters that are displayed in the report.

Magnifying glass icon 
Click this button to view a more detailed report about the capacity of disks in a cluster.

Cluster
Name of a cluster.

Disk Space (Total Capacity)
Total amount of unique disk storage space for the computers in a cluster. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value in this column:

- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note:

- If the computer is an IBM SAN File System client, then this value will also include the capacity of any storage subsystem volumes visible to this computers that are owned by the SAN File System.
- The drill-down reports for each computer generates a list of detectable computer disks only. Therefore, the total capacity of disks for each computer drill-down report will equal this value.

Available Disk Space (Total Unallocated Space)
Total amount of unique disk storage space for a computer that is not

allocated to any logical volume on the computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note: If the computer is an IBM SAN File System client, then this value does not include the unallocated space on any storage subsystem volumes visible to it that are owned by the SAN File System.

Owned Disk Space (Capacity)
Total amount of unique disk storage space for computers in a cluster that is owned by those computers. If Data Manager determines that a disk is configured for a different host, that disk's capacity is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note:

- If the computer is an IBM SAN File System client, this value does not include the capacity of any volumes that are visible to it that are owned by the SAN File System.
- If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads from the host ID. If this does not agree with the local host ID, the disk is classified as belonging to a different (Solaris) host.

Owned Available Disk Space (Owned Unallocated Space)
Total amount of unique disk storage space for computers in a cluster that is not allocated to any logical volume within the computers and is owned by the computers. If Data Manager determines that a disk is configured for a different computer, that disk's unallocated space is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

If the computer is an IBM SAN File System client, then this value excludes the unallocated space on any storage subsystem volumes visible to it that are owned by the SAN File System.

If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product will examine the Veritas disk-label and read off the host-ID. If this does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)
Total amount of unique disk space on the computers in a cluster that is not

attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Non-Fibre Channel Attached Available Disk Space (Total Other Attached Unallocated Space)

Total amount of disk storage space on the computers in a cluster that is not allocated to any logical volume within the computers and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Disk Space (Total Fibre Channel Attached Capacity)

Total amount of unique disk space on the computers in a cluster that is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Data agent.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)

Total amount of disk storage space on the computers in a cluster that is not allocated to any logical volume within the computers and is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Data agent.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

By Computer

Chart icon 

Click this button to generate charts about the disk capacity of the computers that are displayed in the report.

Magnifying glass icon 
Click this button to view a more detailed report about the disks on a computer.

Computer
Name of a computer.

Disk Space (Total Capacity)
Total amount of unique disk storage space for a computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note:

- The value that appears in the **TOTAL** row for this column does not include disks on virtual machines where the Data agent is installed.
- If the computer is an IBM SAN File System client, then this value will also include the capacity of any storage subsystem volumes visible to this computers that are owned by the SAN File System.
- The drill-down reports for each computer generates a list of detectable computer disks only. Therefore, the total capacity of disks for each computer drill-down report will equal this value.

Available Disk Space (Total Unallocated Space)
Total amount of unique disk storage space for a computer that is not allocated to any logical volume on the computer. Tivoli Storage Productivity Center does not include the following disks from in its calculation of the value in this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note: If the computer is an IBM SAN File System client, then this value does not include the unallocated space on any storage subsystem volumes visible to it that are owned by the SAN File System.

Owned Disk Space (Capacity)
Total amount of unique disk storage space for a computer that is owned by that computer. If Data Manager determines that a disk is configured for a different host, that disk's capacity is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

If the computer is an IBM SAN File System client, this value does not include the capacity of any volumes that are visible to it that are owned by the SAN File System.

If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads from the host ID. If this does not agree with the local host ID, the disk is classified as belonging to a different (Solaris) host.

Owned Available Disk Space (Owned Unallocated Space)
Total amount of unique disk storage space for a computer that is not allocated to any logical volume within the computer and is owned by the computer. If Data Manager determines that a disk is configured for a different computer, that disk's unallocated space is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note:

- If the computer is an IBM SAN File System client, then this value excludes the unallocated space on any storage subsystem volumes visible to it that are owned by the SAN File System.
- If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product will examine the Veritas disk-label and read off the host-ID. If this does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)
Total amount of unique disk space on a computer that is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Non-Fibre Channel Attached Available Disk Space (Total Other Attached Unallocated Space)
Total amount of disk storage space for a computer that is not allocated to any logical volume within the computer and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Disk Space (Total Fibre Channel Attached Capacity)
Total amount of unique disk space on a computer that is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Data agent.

- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)
Total amount of disk storage space on a computer that is not allocated to any logical volume within the computer and is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Data agent.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

OS Type
Operating system running on a computer.

Network Address
Network address of a computer.

IP Address
IP address of a computer.

Time Zone
Time zone in which a computer is running.

CPU Architecture
Architecture of the processor (for example, IA32).

By Computer Group

Chart icon 
Click this button to generate charts about the disk capacity of computers in the groups that are displayed in the report.

Magnifying glass icon 
Click this button to view a more detailed report about the disk capacity of computers in a group.

Creator
Creator of a computer monitoring group or computer reporting group.

Group
Name of a group.

Disk Space (Total Capacity)
Total amount of unique disk storage space for the computers in a group. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that are on a virtual machine monitored by Tivoli Storage Productivity Center.
- Computer disks on a computer that has not been probed.

- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Available Disk Space (Total Unallocated Space)
Total amount of unique disk storage space for computers in a group that is not allocated to any logical volume on the computers. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note: If the computer is an IBM SAN File System client, then this value does not include the unallocated space on any storage subsystem volumes visible to it that are owned by the SAN File System.

Owned Disk Space (Capacity)
Total amount of unique disk storage space for computers in a group that is owned by those computers. If Data Manager determines that a disk is configured for a different host, that disk's capacity is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on a cluster of virtual machines that is monitored by Tivoli Storage Productivity Center.
- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note:

- If the computer is an IBM SAN File System client, this value does not include the capacity of any volumes that are visible to it that are owned by the SAN File System.
- If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads from the host ID. If this does not agree with the local host ID, the disk is classified as belonging to a different (Solaris) host.

Owned Available Disk Space (Owned Unallocated Space)
Total amount of unique disk storage space for the computers in a group that is not allocated to any logical volume within the computer group and is owned by a computer within the group. If Data Manager determines that a disk is configured for a different computer, that disk's unallocated space is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Note:

- If the computer is an IBM SAN File System client, then this value excludes the unallocated space on any storage subsystem volumes visible to it that are owned by the SAN File System.
- If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product will examine the Veritas disk-label and read off the host-ID. If this does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)
Total amount of unique disk space on the computers in a group and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks discovered on virtual machine servers where the Data agent is installed.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Non-Fibre Channel Attached Available Disk Space (Total Other Attached Unallocated Space)
Total amount of disk storage space on the computers in a group that is not allocated to any logical volume within the computers and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Disk Space (Total Fibre Channel Attached Capacity)
Total amount of unique disk space on the computers in a group that is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Data agent.
- Computer disks discovered on virtual machine servers where the Data agent is installed.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)
Total amount of disk storage space on the computers in a group that is not allocated to any logical volume within the computers and is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Data agent.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

By Domain

Chart icon 
Click this button to generate charts about the disk capacity of the computers in a domain.

Magnifying glass icon 
Click this button to view a more detailed report about the disk capacity of computers in a domain.

Domain
Name of a domain.

Disk Space (Capacity)
Total amount of unique disk storage space for the computers in a domain. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Available Disk Space (Unallocated Space)
Total amount of unique disk storage space for computers in a domain that is not allocated to any logical volume on the computers. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

Network-wide

Chart icon 
Click this button to generate charts about the disk capacity across a network.

Magnifying glass icon 
Click this button to view a more detailed report about the disk capacity of computers in a network.

Disk Space (Capacity)
Total amount of internal physical disk space on monitored storage subsystem arrays and disk space on monitored systems that are not storage subsystem volumes. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage subsystem disks or computer disks that become missing after a probe.
- Back-end storage volumes that are not identified by Tivoli Storage Productivity Center as being assigned to a monitored storage subsystem virtualizer.
- Computer disks that are identified by Tivoli Storage Productivity Center as being assigned to a monitored storage subsystem.
- Computer disks that reside within a virtual machine server where the Data agent is installed.

Consumed File System Space (Filesystem Used Space)
Total amount of unique file system storage space that is used or consumed by the systems where the Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems on discovered hypervisors that are assigned to discovered virtual machines where the Data agent is installed.
- File systems on discovered virtual machines where the Data agent is installed, but the hypervisor for the virtual machine has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Available File System Space (Filesystem Free Space)
Total amount of unique file system storage space that is not used or available to the systems where the Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Raw Volume Space
Total amount of storage subsystem volume space that is the following:

- Assigned or mapped to hosts.
- Storage subsystem volume space on a system that it is monitored by Tivoli Storage Productivity Center.
- Logical volume space on hosts that do not reside within a storage subsystem monitored by Tivoli Storage Productivity Center.

Unavailable Disk Space (Overhead)
Total amount of storage subsystem volume space that is identified as the following:

- For RAID or mirror redundancy on a storage subsystem.
- As storage subsystem volume space used for replication targets.
- As logical volume space on hosts that is for RAID or mirror redundancy on those hosts.

Available Disk Space (Unallocated Space)
Total amount of disk space that does not include the following from the network capacity:

- Network Raw Volume Space
- Network Overhead
- Network Unknown Volume Space

Unknown Storage Subsystem Volume Space (Unknown LUN Capacity)
Total amount of storage subsystem volume space that is not mapped or assigned to a host and is not identified as residing on a host that has been probed by a Data agent.

TPC-Wide Storage Space - File System Space report:

Use these reports to view detailed information about the storage capacity of file systems on machines in your environment. You can view this information in different ways using the following By clauses: by filesystem, by filesystem group, by cluster, by computer, by computer group, by domain, and for the entire network.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

By Filesystem

Chart icon 
Click this button to generate charts about the storage space for a file system.

Computer
Name of the computer on which a file system resides.

Mount Point
Name or mount point (UNIX/Linux) of the file system (for example, c:\, d;\, /opt, /export/home, etc.). The file systems displayed in this report include the following:

- File systems on virtual machines that are monitored by IBM Tivoli Storage Productivity Center.
- File systems on hypervisors that are monitored by Tivoli Storage Productivity Center.
- File systems that become missing after a probe.

File System Space (Capacity)
Amount of file system storage space for a file system.

Note:

- The value that appears in the **TOTAL** row for this column does not include file systems on virtual machines where the Data agent is installed.
- The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater

than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

Detectable
Indicates **Yes** when a file system is identified as being available to the server following a successful probe.

Percent Used Space
Percentage of file system storage space that is used within a file system.

Consumed File System Space (Used Space)
Amount of used storage space within a file system. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

The value that appears in the **TOTAL** row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Data agent has been installed.
- File systems on discovered virtual machine servers where the Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.

Available File System Space (Free Space)
Amount of unused storage space within a file system.

File Count
Total number of files within a file system. This value does not include files on file systems that have not been scanned.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a file system. This value does not include directories on file systems that have not been scanned.

Percent Used Inodes
Percentage of used Inodes within a file system.

Used Inodes
Number of used Inodes within a file system.

Free Inodes
Number of free Inodes within a file system.

Volume Group
Volume group to which a logical volume belongs. If the volume is not part of a volume group this column displays **N/A**.

Logical Volume Space (Logical Volume Allocation)
Logical volume space used by the logical volume containing the file system.

Disk Detectable
Indicates **Yes** when a file system's logical volume is identified as having an underlying host disk.

By Filesystem Group

Chart icon 
Click this button to generate charts about the capacity of file systems in file system monitoring groups or file system reporting groups.

Magnifying glass icon 
Click this button to view more detailed information about the capacity of file systems in a file system monitoring group or file system reporting group.

Creator
ID of the user who created a group.

Group Name of a file system monitoring group or file system reporting group.

File System Space (Capacity)
Amount of file system storage space for the file systems in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Data agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

Percent Used Space
Percentage of total file system storage space that is used within a group.

Consumed File System Space (Used Space)
Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Data agent is installed if both the virtual machine file systems and hypervisor file systems reside within the group.
- File systems on virtual machine servers where the Data agents installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is

available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are not mounted.

File Count
Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a group.

Used Inodes
Number of used Inodes within a group.

Free Inodes
Number of free Inodes within a group.

By Cluster

Chart icon 
Click this button to generate charts about the used file system space for a cluster.

Magnifying glass icon 
Click this button to view more detailed information about the used storage space on file systems that are in a cluster.

Cluster
Name of cluster.

File System Space (Capacity)
Total amount of unique file system storage space in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.

- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

Percent Used Space
Percentage of total file system storage space that is used within a cluster.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the file systems in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

File Count
Total number of files within a cluster. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Directory Count
Total number of directories within a cluster. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a cluster.

Used Inodes
Number of used Inodes within a cluster.

Free Inodes
Number of free Inodes within a cluster.

By Computer

Chart icon 
Click this button to generate a chart about the used file system storage space of the computers that are displayed in the report.

Magnifying glass icon 
Click this button to view a more detailed report about the used file system storage space on a computer.

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computer where the Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note:

- The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.
- The value that appears in the **TOTAL** row for this column does not include file systems on virtual machines where the Data agent is installed.
- The drill-down reports for each computer will generate a list of detectable file systems only. Therefore, the total capacity of file systems for each computer drill-down report will equal this value.

Percent Used Space
Percentage of total file system storage space that is used within a computer.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computer where the Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to a computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

File Count
Total number of files within a computer. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the **File Count** column is 200.

Directory Count
Total number of directories within a computer. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a computer.

Used Inodes
Number of used Inodes within a computer.

Free Inodes
Number of free Inodes within a computer.

By Computer Group

Chart icon 
Click this button to generate charts about the used file system storage space for a computer monitoring group or computer reporting group.

Magnifying glass icon 
Click this button to view more detailed information about the file systems in a computer monitoring group or computer reporting group.

Creator
ID of the user who created a group.

Group Name of a computer monitoring group or computer reporting group.

File System Space (Capacity)
Amount of file system storage space for the computers in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Data agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

Percent Used Space
Percentage of total file system storage space that is used within a group.

Consumed File System Space (Used Space)
Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Data agent is installed if both the virtual machine file systems and hypervisor file systems reside within the group.
- File systems on virtual machine servers where the Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are not mounted.

File Count
Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.

- Files on file systems that have not been scanned by a Data agent.

Note: For groups that contain Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the **File Count** column is 200.

Directory Count
Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a group.

Used Inodes
Number of used Inodes within a group.

Free Inodes
Number of free Inodes within a group.

By Domain

Chart icon 
Click this button to generate charts about the file system used space for a domain.

Magnifying glass icon 
Click this button to view more detailed information about the file system used space of computers that are in a domain.

Domain
Name of a domain.

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

Percent Used Space
Percentage of total file system storage space that is used within a domain.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server
- File systems that are not mounted.

File Count
Total number of files within a domain. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a domain. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a domain.

Used Inodes
Number of used Inodes within a domain.

Free Inodes
Number of free Inodes within a domain.

Network-wide

Chart icon 
Click this button to generate a chart about the file system used space in a network.

Magnifying glass icon 
Click this button to view more detailed information about the file system used space in a network.

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

Percent Used Space
Percentage of total file system storage space that is used within a network.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server

- File systems that are not mounted.

File Count
Total number of files within a network. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a network. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a network.

Used Inodes
Number of used Inodes within a network.

Free Inodes
Number of free Inodes within a network.

TPC-Wide Storage Space - Consumed File System Space report:

Use these reports to view detailed information about the used storage space in file systems on computers in your environment. You can view this information in different ways using the following By clauses: by filesystem, by filesystem group, by cluster, by computer, by computer group, by domain, and for the entire network.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

By Filesystem

Chart icon 
Click this button to generate charts about the used space information for a file system.

Computer
Name of the computer on which a file system resides.

Mount Point
Name or mount point (UNIX/Linux) of the file system (for example, c:\, d;\, /opt, /export/home, etc.). The file systems displayed in this report include the following:

- File systems on virtual machines that are monitored by IBM Tivoli Storage Productivity Center.
- File systems on hypervisors that are monitored by Tivoli Storage Productivity Center.
- File systems that become missing after a probe.

Consumed File System Space (Used Space)
Amount of used storage space within a file system. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

The value that appears in the **TOTAL** row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Data agent has been installed.
- File systems on discovered virtual machine servers where the Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.

Detectable
Indicates **Yes** when a file system is identified as being available to the server following a successful probe.

Percent Used Space
Percentage of file system storage space that is used within a file system.

Available File System Space (Free Space)
Amount of unused storage space within a file system.

File System Space (Capacity)
Amount of file system storage space for a file system.

Note:

- The value that appears in the **TOTAL** row for this column does not include file systems on virtual machines where the Data agent is installed.
- The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

File Count
Total number of files within a file system. This value does not include files on file systems that have not been scanned.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the

number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a file system. This value does not include directories on file systems that have not been scanned.

Percent Used Inodes
Percentage of used Inodes within a file system.

Used Inodes
Number of used Inodes within a file system.

Free Inodes
Number of free Inodes within a file system.

Volume Group
Volume group to which a logical volume belongs. If the volume is not part of a volume group this column displays **N/A**.

Logical Volume Space (Allocation)
Logical volume space used by the logical volume containing the file system.

Disk Detectable
Indicates **Yes** when a file system's logical volume is identified as having an underlying host disk.

By Filesystem Group

Chart icon 
Click this button to generate charts about the used file system storage space in a file system monitoring group or file system reporting group.

Magnifying glass icon 
Click this button to view more detailed information about the used file system storage space in a file system monitoring group or file system reporting group.

Creator
ID of the user who created a group.

Group Name of a file system monitoring group or file system reporting group.

Consumed File System Space (Used Space)
Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Data agent is installed if both the virtual machine file systems and hypervisor file systems reside within the group.
- File systems on virtual machine servers where the Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.

- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Percent Used Space
Percentage of total file system storage space that is used within a group.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are not mounted.

File System Space (Capacity)
Amount of file system storage space for the file systems in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Data agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

File Count
Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a group.

Used Inodes
Number of used Inodes within a group.

Free Inodes
Number of free Inodes within a group.

By Cluster

Chart icon 
Click this button to generate charts about the used file system space for a cluster.

Magnifying glass icon 
Click this button to view more detailed information about the used storage space on file systems that are in a cluster.

Cluster
Name of cluster.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the file systems in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Percent Used Space
Percentage of total file system storage space that is used within a cluster.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

File System Space (Capacity)
Total amount of unique file system storage space in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

File Count
Total number of files within a cluster. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Directory Count
Total number of directories within a cluster. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a cluster.

Used Inodes
Number of used Inodes within a cluster.

Free Inodes
Number of free Inodes within a cluster.

By Computer

Chart icon 
Click this button to generate charts about the used file system storage space of the computers that are displayed in the report.

Magnifying glass icon 
Click this button to view a more detailed report about the used file system storage space on a computer.

Computer
Name of the computer on which a file system resides.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computer where the Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Percent Used Space
Percentage of total file system storage space that is used within a computer.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to a computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computer where the Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note:

- The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.
- The value that appears in the **TOTAL** row for this column does not include file systems on virtual machines where the Data agent is installed.
- The drill-down reports for each computer will generate a list of detectable file systems only. Therefore, the total capacity of file systems for each computer drill-down report will equal this value.

File Count
Total number of files within a computer. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the

number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a computer. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a computer.

Used Inodes
Number of used Inodes within a computer.

Free Inodes
Number of free Inodes within a computer.

By Computer Group

Chart icon 
Click this button to generate charts about the used file system storage space for a computer monitoring group or computer reporting group.

Magnifying glass icon 
Click this button to view more detailed information about the file systems in a computer monitoring group or computer reporting group.

Creator
ID of the user who created a group.

Group Name of a computer monitoring group or computer reporting group.

Consumed File System Space (Used Space)
Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Data agent is installed if both the virtual machine file systems and hypervisor file systems reside within the group.
- File systems on virtual machine servers where the Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Percent Used Space
Percentage of total file system storage space that is used within a group.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are not mounted.

File System Space (Capacity)
Amount of file system storage space for the computers in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Data agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

File Count
Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: For groups that contain Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a group.

Used Inodes
Number of used Inodes within a group.

Free Inodes
Number of free Inodes within a group.

By Domain

Chart icon 
Click this button to generate charts about the file system used space for a domain.

Magnifying glass icon 
Click this button to view more detailed information about the file system used space of computers that are in a domain.

Consumed File System Space (Domain)
Name of a domain.

Used Space
Total amount of unique file system storage space that is used or consumed by the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Percent Used Space
Percentage of total file system storage space that is used within a domain.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server
- File systems that are not mounted.

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

File Count
Total number of files within a domain. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a domain. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a domain.

Used Inodes
Number of used Inodes within a domain.

Free Inodes
Number of free Inodes within a domain.

Network-wide

Chart icon 
Click this button to generate charts about the file system used space in a network.

Magnifying glass icon 
Click this button to view more detailed information about the file system used space in a network.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Percent Used Space
Percentage of total file system storage space that is used within a network.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server
- File systems that are not mounted.

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

File Count
Total number of files within a network. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a network. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a network.

Used Inodes
Number of used Inodes within a network.

Free Inodes
Number of free Inodes within a network.

TPC-Wide Storage Space - Available File System Space report:

Use these reports to view detailed information about the available storage space in file systems on systems in your environment. You can view this information in different ways using the following By clauses: by filesystem, by filesystem group, by cluster, by computer, by computer group, by domain, and for the entire network.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

By Filesystem

Chart icon 
Click this button to generate a chart about the used space information for a file system.

Computer
Name of the computer on which a file system resides.

Mount Point
Name or mount point (UNIX/Linux) of a file system (for example, c:\, d;\, /opt, /export/home, etc.). The file systems displayed in this report include the following:

- File systems on virtual machines that are monitored by IBM Tivoli Storage Productivity Center.
- File systems on hypervisors that are monitored by Tivoli Storage Productivity Center.
- File systems that become missing after a probe.

Available File System Space (Free Space)
Amount of unused storage space within a file system.

Detectable
Indicates **Yes** when a file system is identified as being available to the server following a successful probe.

Percent Available File System Space (Percent Free Space)
Percentage of file system storage space that is unused within a file system.

Consumed File System Space (Used Space)
Amount of used storage space within a file system. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

The value that appears in the **TOTAL** row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Data agent has been installed.

- File systems on discovered virtual machine servers where the Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.

File System Space (Capacity)
Amount of file system storage space for a file system.

Note:

- The value that appears in the **TOTAL** row for this column does not include file systems on virtual machines where the Data agent is installed.
- The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

File Count
Total number of files within a file system. This value does not include files on file systems that have not been scanned.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a file system. This value does not include directories on file systems that have not been scanned.

Percent Used Inodes
Percentage of used Inodes within a file system.

Used Inodes
Number of used Inodes within a file system.

Free Inodes
Number of free Inodes within a file system.

Volume Group
Volume group to which a logical volume belongs. If the volume is not part of a volume group this column displays **N/A**.

Logical Volume Space (Allocation)
Logical volume space used by the logical volume containing the file system.

Disk Detectable
Indicates **Yes** when a file system's logical volume is identified as having an underlying host disk.

By Filesystem Group

Chart icon 
Click this button to generate charts about the used file system storage space in a file system monitoring group or file system reporting group.

Magnifying glass icon 
Click this button to view more detailed information about the used file system storage space in a file system monitoring group or file system reporting group.

Creator
ID of the user who created a group.

Group Name of a file system monitoring group or file system reporting group.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are not mounted.

Percent Available File System Space (Percent Free Space)
Percentage of total file system storage space that is not used or available within a group.

Consumed File System Space (Used Space)
Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Data agent is installed if both the virtual machine file systems and hypervisor file systems reside within the group.
- File systems on virtual machine servers where the Data agents installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

File System Space (Capacity)
Amount of file system storage space for the file systems in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Data agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

File Count
Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a group.

Used Inodes
Number of used Inodes within a group.

Free Inodes
Number of free Inodes within a group.

By Cluster

Chart icon 
Click this button to generate a chart about the used file system space for a cluster.

Magnifying glass icon 
Click this button to view more detailed information about the used storage space on file systems that are in a cluster.

Cluster
Name of cluster.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Percent Available File System Space (Percent Free Space)
Percentage of total file system storage space that is available within a cluster.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the file systems in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

File System Space (Capacity)
Total amount of unique file system storage space in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

File Count
Total number of files within a cluster. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Directory Count
Total number of directories within a cluster. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a cluster.

Used Inodes
Number of used Inodes within a cluster.

Free Inodes
Number of free Inodes within a cluster.

By Computer

Chart icon 
Click this button to generate a chart about the used file system storage space of the computers that are displayed in the report.

Magnifying glass icon 
Click this button to view a more detailed report about the used file system storage space on a computer.

Computer
Name of the computer on which a file system resides.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to a computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Percent Available File System Space (Percent Free Space)
Percentage of total file system storage space that is available within a computer.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computer where the Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computer where the Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.

- File systems that are not mounted.

Note:

- The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.
- The value that appears in the **TOTAL** row for this column does not include file systems on virtual machines where the Data agent is installed.
- The drill-down reports for each computer will generate a list of detectable file systems only. Therefore, the total capacity of file systems for each computer drill-down report will equal this value.

File Count

Total number of files within a computer. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count

Total number of directories within a computer. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes

Percentage of used Inodes within a computer.

Used Inodes

Number of used Inodes within a computer.

Free Inodes

Number of free Inodes within a computer.

By Computer Group

Chart icon 

Click this button to generate charts about the used file system storage space for a computer monitoring group or computer reporting group.

Magnifying glass icon 
Click this button to view more detailed information about the file systems in a computer monitoring group or computer reporting group.

Creator
ID of the user who created a group.

Group Name of a computer monitoring group or computer reporting group.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are not mounted.

Percent Available File System Space (Percent Free Space)
Percentage of total file system storage space that is available within a group.

Consumed File System Space (Used Space)
Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Data agent is installed if both the virtual machine file systems and hypervisor file systems reside within the group.
- File systems on virtual machine servers where the Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

File System Space (Capacity)
Amount of file system storage space for the computers in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Data agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

File Count
Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: For groups that contain Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a group.

Used Inodes
Number of used Inodes within a group.

Free Inodes
Number of free Inodes within a group.

By Domain

Chart icon 
Click this button to generate charts about the file system used space for a domain.

Magnifying glass icon 
Click this button to view more detailed information about the file system used space of computers that are in a domain.

Domain
Name of a domain.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server
- File systems that are not mounted.

Percent Available File System Space (Percent Free Space)
Percentage of total file system storage space that is available within a domain.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

File Count
Total number of files within a domain. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a domain. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.

- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a domain.

Used Inodes
Number of used Inodes within a domain.

Free Inodes
Number of free Inodes within a domain.

Network-wide

Chart icon 
Click this button to generate a chart about the file system used space in a network.

Magnifying glass icon 
Click this button to view more detailed information about the file system used space in a network.

Available File System Space (Free Space)
Total amount of unique file system storage space that is not used or is available to the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server
- File systems that are not mounted.

Percent Available File System Space (Percent Free Space)
Percentage of total file system storage space that is available within a network.

Consumed File System Space (Used Space)
Total amount of unique file system storage space that is used or consumed by the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Note: On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

File System Space (Capacity)
Total amount of unique file system storage space that was discovered on the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.

- File systems that are not mounted.

Note: The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

File Count
Total number of files within a network. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Data agent.

Note: On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

Directory Count
Total number of directories within a network. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Data agent.

Percent Used Inodes
Percentage of used Inodes within a network.

Used Inodes
Number of used Inodes within a network.

Free Inodes
Number of free Inodes within a network.

Usage Reports

Use Usage reports to view the usage and growth of the consumption of your storage. You can view these reports based on specific filesystems and computers, groups of filesystems and computers, or throughout your entire enterprise. These reports help you project storage consumption for the future and maximize the current storage assets that you have in place by eliminating wasted space and making the most of the space you have.

Use these reports to do the following:

- View which servers and filesystems are experiencing the heaviest (or lightest) load of data and storage access.
- Identify wasted space by pinpointing files that are no longer needed or have not been accessed for the longest time.

- Associate users with the amount of storage they are consuming on specific filesystems, groups of filesystems, specific computers, groups of computers, and throughout the entire network.
- Select the files against which to run an IBM Tivoli Storage Manager archive/backup job. Tivoli Storage Manager protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage.

Access Load reports: Use Access Load reports to view which machines and file systems within your environment are experiencing the heaviest (or lightest) load of data and storage access.

Access Time reports

Use Access Time reports to view and chart the amount of data and the number of files that have been accessed during the last day, the last week, the last month, the last year, and over. You can view this information at the directory level, the file system level, the computer level, at the domain level, and for the entire network.

The following report shows a single pie chart that depicts all of the data across the entire enterprise, giving you an easy way to see how much of your corporate data is accessed every day, every week, and even data that has not been accessed in over a year.

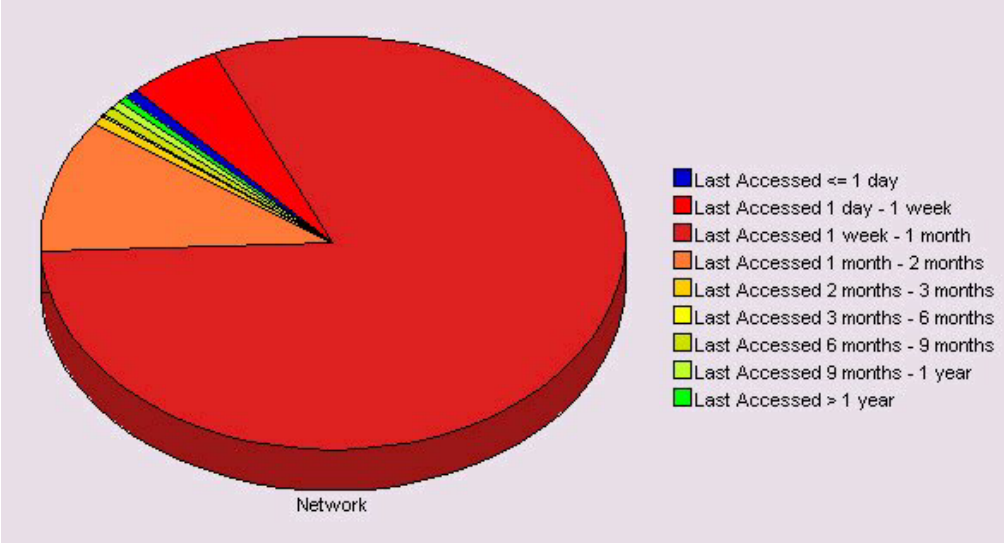

Figure 53. Usage Reporting—Access Load—Access Time—Chart Network Wide

Modification Time reports

Use Modification Time reports to view and chart the amount of data and the number of files that have been modified during the last day, the last week, the last month, the last year, and over. You can view this information at the directory level, the file system level, the computer level, at the domain level, and for the entire network.

The following report shows a single pie chart that depicts all of the data across the entire enterprise, allowing you to see how much of your corporate data is modified every day, every week, and even data that has not been modified in over a year.

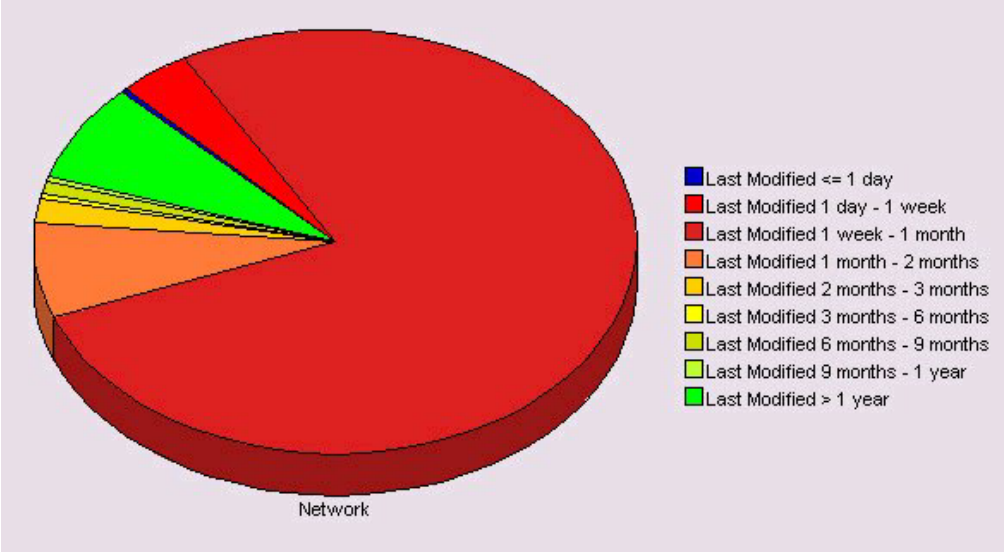

Figure 54. Usage Reporting—Access Load—Modification Time—Chart Network Wide

Creation Time reports

Use Creation Time reports to view and chart the amount of data and the number of files that have been created during the last day, the last week, the last month, the last year, and over. You can view this information at the directory level, the file system level, the computer level, at the domain level, and for the entire network.

Note: Because scans do not collect creation time information for the files stored on UNIX machines, creation time reports do not contain any data for the non-Windows machines in your environment.

Files reports: Use Files Reports to report on the files found on the storage in your environment during theData Manager scan process. During the scan process, Data Manager gathers statistics on your files and attributes about the files. Select the files against which to run an IBM Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 355 for information on how to create archive and backup jobs from the files shown in a Files report. The information gathered can be reported on in a variety of ways using this facility:

Largest Files reports

Use these reports to view information about the largest files in your environment. Largest Files reports enable you to do the following:

- Generate reports that contain detailed information about the largest files found in your environment. You can view reports on the largest files in the following ways: by directory, by directory group, by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

  Note: The number of files that are gathered and stored in the repository is dependent upon the number that you specify in the profile that you use during scanning. The default profile gathers information on the 20 largest files. You can select a profile for the report by clicking the **Profile** list box.

- Select files against which to run anTivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 355 for information on how to create archive/backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive/backup job. See "Archive and backup functions" on page 355 for information on how to add files from reports to an archive/backup job.

Duplicate Files reports

Use these reports to view information about the files found in your environment during a scan that contain duplicate file names. Duplicate Files reports enable you to do the following:

- Generate reports on the files found in your environment during a scan that contain duplicate file names. Use this report to view a list of files that might no longer be needed and could be wasting storage space. You can use the data generated in a duplicate files report to view a list of files that might no longer be needed and could be wasting storage space. You can view reports on duplicate files in the following ways: by directory, by directory group, by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

  Note: The number of files that are gathered and stored in the repository is dependent upon the number that you specify in the profile that you use during scanning. The default profile gathers information on the 20 largest files. You can select a profile for the report by clicking the **Profile** list box.
- Select files against which to run an Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 355 for information on how to create archive and backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 355 for information on how to add files from reports to an archive or backup job.

Before generating a Duplicate Files report for environments that contain a large amount of files (i.e., environments that might contain over 100,000 duplicate files), we recommend you tune theData Manager repository. To do this, run the appropriate RDBMS utility to collect statistics on the repository tables. Run the following utilities depending on the RDBMS where the repository is stored:

- UDB: runstats
- Oracle: analyze
- Microsoft SQL Server: update stats
- Sybase: update stats

Run the appropriate utility against all the tables in the repository schema you specified during the installation process.

To generate the data for a duplicate files report, you can either create a constraint that will save filenames for *n* number of violation files, or run a scan that includes one or all of the following existing profiles:

- TPCUser.Largest Files
- TPCUser.Largest Orphans
- TPCUser.Most at Risk

- TPCUser.Most Obsolete

There is not a specific Duplicate Files profile that collects information on all duplicate files in a file system. The filenames that are triggered by the previously listed profiles or constraints will be saved to the repository and the duplicate files report is generated from the data gathered from the scan or constraint.

For example, create a scan using the Largest Files profile to collect information on the 100 largest files. The 100 largest files are saved in the t_stat_file table in the repository. If five duplicate files were saved as a result of being flagged by the Largest File Profile, then those files will be reported as duplicates in the Duplicate file Report.

Most Obsolete Files reports

Use these reports to view information about the files found in your environment that have not been accessed in the longest period of time. Most Obsolete Files reports enable you to do the following:

- Generate reports on the files found in your environment that have not been accessed in the longest period of time. It is very likely that many of these files are no longer needed, at least online, and could be wasting space. You can view reports on the largest files in the following ways: by directory, by directory group, by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

  Note: The number of files that are gathered and stored in the repository is dependent upon the number that you specify in the profile that you use during scanning. The default profile gathers information on the 20 largest files. You can select a profile for the report by clicking the **Profile** list box.
- Select files against which to run an Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 355 for information on how to create archive/backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 355 for information on how to add files from reports to an archive or backup job.

Orphan Files reports

Use these reports to view information about files that are owned by users that are no longer active in the environment. Orphan files reporting enables you to do the following:

- Generate reports on files that are owned by users that are no longer active in the environment. It is very likely that many of these files are no longer needed and could be wasting space. You can view reports on orphan files in the following ways: by directory, by directory group, by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

  You can also generate reports on the oldest orphan files that have not been accessed for the longest periods of time, as well as largest orphan files.

  Note: The number of files that are gathered and stored in the repository is dependent upon the number that you specify in the profile that you

use during scanning. The default profile gathers information on the 20 largest files. You can select a profile for the report by clicking the **Profile** list box.

- Select files against which to run an Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 355 for information on how to create archive and backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive or backup job. See "Archive and backup functions" on page 355 for information on how to add files from reports to an archive or backup job.

File Size Distribution reports

Use these reports to view information about the distribution of file sizes across storage resources. Use File Size Distribution reports to generate reports that show the distribution of file sizes on a specific directory, a file system, a computer, the domain, or across your entire network.

The following report shows a single pie chart for each computer in the report.

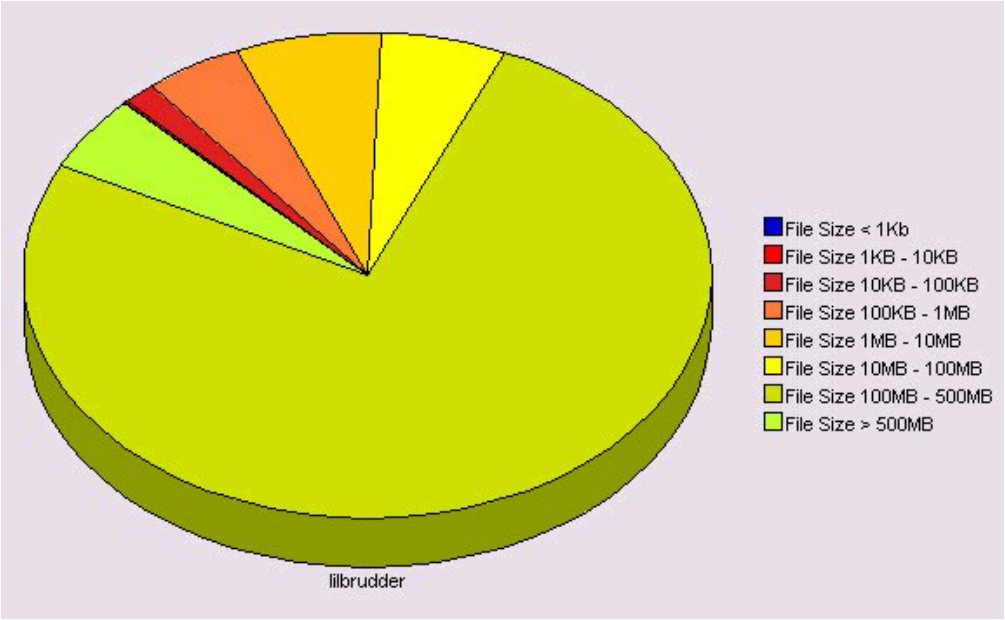

Figure 55. Usage Reporting—Files—File Size Distribution—File Count Chart Tab

Each pie chart depicts the number of files that fall into each file size bucket that are described in the legend. For example, the blue portion of the pie represents the number of files on each computer that are smaller than 1024 bytes long. This gives you an easy way to get a quick grasp of what are the size distribution of your data files.

File Summary reports

Use these reports to view general information about the files in your environment. Use File Summary reports generate overview information for files, also at the directory level, the file system level, the computer level, the domain level or for the whole network.

This report shows you the following (for each entity at the level you are reporting by):

- The total size of the files
- The number of files
- The number of directories
- The average file size

File Types reports

Use these reports to view information about the storage usage of different file types in your environment. File Types reports enable you to generate reports on the file types found in your environment during a scan. The scan that collects data for these reports must use the TPCUser.Summary By File Type profile. For example, when you run a scan using the TPCUser.Summary By File Type profile,Data Manager will collect storage usage information about such file types as .exe, .zip, .sys, .pdf, .doc, .dll, .wav, .mp3, .avi, etc. Each of these file types is represented by its own row in the generated reports.

Use these reports to do the following:

- Relate the space used by applications to the total capacity and used space. For example, you can use these reports to view the total amount of storage consumed by Adobe Acrobat files and Lotus Notes mail databases.
- View which applications are consuming the most space across a given set of storage resources

You can view these reports in the following ways:

- **Group by Resource**. These reports organize data by storage resource (for example, file system, computer, cluster, etc.) and space used.
- **Group by File Type**. These reports organizes data by file type and space used. They show where the most spaced used by each type of file resides.

Archive/Backup jobs from reports:

Use the Archive/Backup function from Reporting to define an archive or backup job to run against the files that you select from specified reports.

Tivoli Storage Manager protects your organization's data from hardware failures and other let you back up specified files. These files are not excluded in the include-exclude list and meet the requirement for serialization in the backup copy group of the management class assigned to each file. An incremental backup enables you to back up new or changed files or directories from a client domain or from specified directories or files. These directories or files are not excluded in the include-exclude list and meet the requirements for frequency, mode, and serialization as defined by a backup copy group of the management class assigned to each file. See *IBM Tivoli Storage Manager Client User's Guide* for more information about Tivoli Storage Manager archives and backups.

This function enables you to select a specific file or group of files fromData Manager reports that you want to archive or back up using the Tivoli Storage Manager tool.

Archive/backup is run from**Data Manager -> Reporting**. You cannot create archive and backup jobs from the Archive/Backup node - rather you create these jobs by selecting files from the previously listed reports or the reports that you create. Each

archive and backup that you create from reports appears as a node under the Archive/Backup node. The reports that contain files against which you can run an archive or backup job include:

- **Data Manager -> Reporting -> Usage -> Files -> Largest Files**
- **Data Manager -> Reporting -> Usage -> Files -> Most Obsolete Files**
- **Data Manager -> Reporting -> Usage -> Files -> Orphan Files**
- **Data Manager -> Reporting -> Usage -> Files -> Duplicate Files**
- **Data Manager -> Reporting -> Backup -> Most at Risk Files ->**
- **Data Manager -> Reporting -> Usage Violations -> Constraint Violations**

You can also run an archive or backup job from a report you create.

Note:

See *Creating Archive/Backup jobs* in *Reporting* for information about how to create an archive or backup job from the files shown in reports.

Directories reports:

Use Directories reports to report upon the directories found on the storage in your environment during the scan process. During the scan process, Data Manager gathers a vast number of statistics on your directories and a number of attributes about the directories. Before you can view information about directories in these reports, you must do the following:

1. Use the **Monitoring → Groups → Directory** function to create directory groups that contain the directories upon which you want to report.
2. Include the directory groups you created in scans.
3. Schedule and run the scans that contain the directory groups (and directories) upon which you want to report.

Largest Directories reports

Use Largest Directories reports to view reports on the *n* largest directories found in your environment:

- for specific directories
- on specific computers or groups of computers
- on specific file systems or groups of file systems
- in specific domains
- on the entire network

The number of largest (or the value of *n* in finding the *n* largest directories), is determined by the profile that you use for doing the reporting. When you define a profile in the **Monitoring → Profiles** node, you specify the number of largest directories upon which to collect information. When you generate a Largest Directories report, you can select what profile to use.

Monitored Directory Space reports

Use Monitored Directory reports to report on the monitored directories that you have included in a group via the **Monitoring → Groups → Directory** facility and included in a scan. During the scan process, Data Manager gathers a number of statistics on your directories and a number of attributes about the directories.

Users reports: Use Users Reports to view information about the users that are found as file owners in your environment. These reports provide an easy and convenient way to view and aggregate storage resource consumption by owner across directories, file systems, computers, domains, and across your entire network.

OS User Groups reports: OS User Groups reports give you the exact same capabilities as the reports under Users reports. The only difference is that you are doing your analysis of OS User Groups as opposed to actual single users.

Usage Violation reports

Use Usage Violation reports to view information about quota and constraint violations.

Usage Violation reports enable you to:

- Create and enforce corporate storage usage policies based on reports that show quota and constraint violations.

  Use quotas to control how much storage a user, or a group of users, can consume on a filesystem or group of filesystems, a computer or group of computers, and throughout the entire network. Use Constraints to set policies about which types of files, owners, and file sizes are allowed on a filesystem or group of file systems, a computer or group of computer, and throughout the entire network.
- Select files in Constraint Violation reports against which to run an IBM Tivoli Storage Manager archive/backup job. See "Archive and backup functions" on page 355 for information on how to create archive/backup jobs from the files shown in a Files report.
- Add files from Constraint Violation reports to an existing Tivoli Storage Manager archive or backup job. See "Create an archive/backup job" on page 356 for information on how to add files from reports to an archive/backup job.

Constraint Violations reports: Constraint Violations reporting enables you to do the following:

- View the files that have violated a constraint. The report shows one row for each Constraint. The row shows the detail of the constraint and the violating space totals for the files that violate the constraint. You can view reports on the files that violate constraints in the following ways: by file system, by file system group, by computer, by computer group, by domain, and by constraint.
- Select files against which to run an Tivoli Storage Manager archive or backup job. See the "Archive and backup functions" on page 355 section for information on how to create archive and backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive or backup job. See the "Archive and backup functions" on page 355 section for information on how to add files from reports to an archive or backup job.

Quota Violations reports: Use Quota Violations reports to report on the quotas that you have defined and see how much space is being used by the various entities that have quotas defined upon them.

Quotas can be defined on either of two categories of users: on a specific user or group of users that you have defined within either Users or on OS User Groups.

Quota Category	Description
User	This allows you to define a quota on a user or group of users. The users are the owners of the files that Data Manager finds as it scans your systems.
OS User Group	This allows you to define a quota on an OS user group or group of OS user groups. OS user groups are defined outside the tool and are found by Data Manager as it scans your systems.

Within each of those categories, you can define quotas at any one of following three levels:

Quota Level	Description
Network	This allows you to define quotas on users or groups of users at the entire network level. This enables you to control how much any single user or group of users is using on all of the storage throughout your enterprise.
Computer	This allows you to define quotas on users or groups of users at the computer level. This enables you to control how much any single user or group of users is using on a single computer, or on any given group of computers.
Filesystem	This allows you to define quotas on users or groups of users at the file system level. This enables you to control how much any single user or group of users is using on a single file system, or on any given group of file systems.

All Quotas reports

Use All Quotas reports to report on all quotas that have been violated, including those at the network level, computer level, and file system level. The **Selection** tab allows you to select upon what users you wish to do quota reporting.

The **Objects** tab shows you for each user and computer and file system:

- how much space is being used
- what the Quota threshold is
- what the percentage of the Quota that is being consumed

This gives you a simple way to analyze how much space is being used relative to the Quotas that you have defined on your systems and users.

Network-wide Quotas reports

Use Network-wide Quota reports to view all the network-wide quotas that have been violated by users and OS user groups in your environment. Information about the quotas that appear on these reports includes the amount of space being used, the threshold value, and the percentage above the threshold value. You can view these reports sorted by user, by OS user group, and by the creator of the quota.

Computer Quotas reports

Use Computer Quota reports to view all the Filesystem Quotas reports computer quotas that have been violated by users and OS user groups in your environment. Information about the quotas that appear on these reports includes the amount of space being used, the threshold value, and the percentage above the threshold value. You can view these reports sorted by user, by OS user group, and by the creator of the quota.

Filesystem Quotas reports

Use Filesystem Quota reports to view all the computer quotas that have

been violated by users and OS user groups in your environment. Information about the quotas that appear on these reports includes the amount of space being used, the threshold value, and the percentage above the threshold value. You can view these reports sorted by user, by OS user group, and by the creator of the quota.

Backup reports

Use Backup reports to identify the most at risk files in your enterprise that are not backed up properly. Critical data that is not backed up might expose you to costly losses of information.

Backup reports also enable you to:

- Investigate the sizes of data that would be required to backup to achieve full backups of any subset of your enterprise (including the whole enterprise).
- Determine the amount of space required to perform an incremental backup of any or all of your data at any given point in time.
- Select files in Most at Risk Files reports against which to run an IBM Tivoli Storage Manager archive/backup job.
- Add files from Most at Risk Files reports to an existing Tivoli Storage Manager archive/backup job.

Most At Risk Files reports: Most At Risk Files reporting enables you to do the following:

- View the files that have the earliest modification time, but have not been backed up since they were modified. This can alert you to a dangerous situation where you may have vulnerable data that could be lost and not recovered due to lack of backups.
- Select files against which to run an IBM Tivoli Storage Manager archive or backup job. See the "Archive and backup functions" on page 355 section for information on how to create archive or backup jobs from the files shown in a Files report.
- Add files to an existingTivoli Storage Manager archive or backup job. See the "Archive and backup functions" on page 355 section for information on how to add files from reports to an archive or backup job.

Modified Files not Backed Up reports: Use Modified Files not Backed Up reports to view the numbers of files that have the earliest modification time, but have not been backed up since they were modified. This can alert you to a dangerous situation where you may have vulnerable data that could be lost and not recovered due to lack of backups.

Backup Storage Requirements reports: Use Backup Storage Requirements reports to determine, at a network, domain computer, file system, or directory level, how much storage will be required to do backups. You can determine how much storage will be required to do both full backups and to do incremental backups.

Reporting Type	Description
Full Backup Size	Use Full Backup Size reports to view how much space it will take to do full backups of specific computers, file systems, directories, groups of those entities, domains, and the entire network.

Reporting Type	Description
Incremental Backup Size	(Windows only) Use this report to view the total size of all the files that have not been backed up since the last scan. You can generate this report by Directory, by Directory Group, by Filesystem, by Filesystem Group, by Computer, by Computer Group, by Domain and Network Wide. This allows you to see for any of these levels, how much data would need to be backed up for any of these desired time intervals. This is useful for backup sizing and can be used to determine appropriate backup groups and sets for timing your backups.
Incremental Range Sizes	Use this report to determine how much storage space you would need to back up data that has been modified in the last day, week, month, two months, three months, six months, nine months, or year. For example, if you want to back up the files that have been modified in the last 3 months only, generate this report and view the values in the ″Last Modified <= 3 months″ column to determine 1) how many files changed and 2) how much storage those changed files are consuming. You can generate this report by Directory, by Directory Group, by Filesystem, by Filesystem Group, by Cluster, by Computer, by Computer Group, by Domain and Network Wide. This allows you to see for any of these levels, how much data would need to be backed up for any of these desired time intervals. This is useful for backup sizing and can be used to determine appropriate backup groups and sets for timing your backups.

Monitored Computer Storage Space reports

Use these host-based reports to view the data that is collected by probes for computers, network attached storage, and hypervisors. These reports contain detailed information about the consumption of file system space and enable you to relate that space back to external or internal storage without counting shared space multiple times.

Monitored computer storage space reports are organized into the following categories:

Disk Storage
Use these reports to view information about disks that reside on computers monitored by IBM Tivoli Storage Productivity Center. This includes information about storage subsystem disks, known disks with serial numbers, or unknown disks without serial numbers.

Non-Disk Storage
Use these reports to view information about storage space that does not reside on computer disks that are monitored by Tivoli Storage Productivity Center. This type of storage includes the following:

- Storage that resides on a remote mount or a network attached server.
- Storage that is temporarily allocated to processor memory.

You must perform the following actions before viewing these reports:

- Run a probe against the storage entities upon which you want to report.
- If you want to generate By Computer Group, By Filesystem Group, and By Storage Subsystem Group reports, perform the following tasks:

- Create computer monitoring group and file system monitoring groups using the **Data Manager** > **Monitoring** > **Groups** node and then probe them.
- Create storage subsystem monitoring group using the **Disk Manager** > **Monitoring** > **Groups** node and probe them.
- Create a computer reporting group or file system reporting group using the **Data Manager** > **Reporting** > **Groups** and **Disk Manager** > **Reporting** > **Groups** node. Probe the entities in the group.
- Create a subsystem reporting group using the **Disk Manager** > **Reporting** > **Groups** node. Probe the storage subsystems in the group.

Related Topics

- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Monitored Computer Storage Space reports - Storage entities:

Monitored Computer Storage Space reports display information about storage entities that you probe. Use this topic to view descriptions of those entities.

You can use Monitored Computer Storage Space reports to view information about the following storage entities:

Computer
Computer that is monitored by a Data agent. This computer can be a single host, a virtual machine, or a hypervisor. A hypervisor is a machine that hosts one or more virtual machines.

Disk on a storage subsystem
Computer disk that originates from a storage subsystem. If the disk's storage subsystem is not monitored by IBM Tivoli Storage Productivity Center, that storage subsystem is represented by a portion of the serial number of the disk. If the storage subsystem is monitored by Tivoli Storage Productivity Center, that storage subsystem is represented by the name collected by a probe.

Disk not on a storage subsystem with a serial number
Computer disk that Tivoli Storage Productivity Center is unable to determine if it originates from a storage subsystem. Tivoli Storage Productivity Center is able to extract the serial number of that disk.

Disk not on a storage subsystem without a serial number
A valid computer disk for which Tivoli Storage Productivity Center is unable to do the following:

- Determine if the disk originates from a storage subsystem.
- Extract the serial number of the disk.

Disk Group
A collection of disk drives on a computer.

File System
Storage entity that creates, removes, and modifies files on a computer. File systems can be shared among monitored computers and span portions of different disks. If a file system spans disk entities, then there will be a single logical volume beneath it that describes how the file system is spanned.

(Parent) Logical Volume
Logical representation of a volume that spans many logical volumes. If there is no file system created on this logical volume, information about the parent logical volume is displayed in the reports.

(Child) Logical Volume
Logical representation of a volume that spans one or more portions of a single or multiple disks. This entity helps determine how the parent logical volumes and file systems are allocated on the disks. If there is no file system created on this logical volume, or there is no parent logical volume for this logical volume, the child logical volume is displayed in the reports.

Non-Disk Storage
Storage space that does not reside on a computer disk. This includes storage residing on a remote mount or a network attached server and storage that is temporarily allocated to processor memory.

Volume Group
A collection of logical volumes on a computer.

Virtual Machine (VM) Disk File
Storage entity that represents a computer disk on a computer that is a virtual machine.

Virtual Machine File System (VMFS)
Storage entity that creates, removes, and modifies VM Disk files on a computer that is a hypervisor hosting one or more virtual machines.

Monitored Computer Storage Space reports - Common storage space values:

Monitored Computer Storage Space reports provide detailed storage space information about the computers probed by IBM Tivoli Storage Productivity Center. Use this topic to view descriptions of the common storage values that appear in these reports.

You can view the following storage space information in Monitored Computer Storage Space reports:

Space Total amount of storage space for the entity that a row represents:

- For file system rows, this is the total logical volume space where a file system resides.
- For logical volume rows, this is the total logical volume space.
- For computer disk rows, this is the total disk space of a computer disk.
- For computer space rows, this is the total storage subsystem volume space on a computer.
- For storage subsystem space rows, this is the total storage subsystem volume space on all monitored computers.

Consumed Space
Amount of consumed storage space for the entity that a row represents:

- For file system rows, this is the space that is consumed by data that can be accessed by the end user of a computer.
- For logical volume rows, this is the space that can be consumed by a file system, but is not yet consumed by a file system.
- For computer disk rows, this is the space that is consumed by the file systems and logical volumes on a disk.

- For computer space rows, this is the storage subsystem volume space that is consumed on a monitored computer.
- For storage subsystem space rows, this is the storage subsystem volume space that is consumed on all monitored computers.

Available Space
Amount of available storage space for the entity that a row represents:
- For file system rows, this is the space that is available for consumption.
- For logical volume rows, this value is zero (0) because no file system has been created on the logical volume that would help determine the amount of available space.
- For computer disk rows, this is the space that is available on the file systems on a computer disk and the space on a computer disk not yet allocated to a logical volume.
- For computer space rows, this is the storage subsystem volume space that is available for consumption on a computer.
- For storage subsystem space rows, this is the storage subsystem volume space that is available for consumption on all monitored computers.

Unavailable Space
Amount of unavailable storage space for the entity that a row represents:
- For file system rows, this space is unavailable for consumption because it is reserved by the software RAID level on the underlying logical volume and is required for data that is used by the operating system to maintain a file system.
- For logical volume rows, this space is unavailable because it is reserved for the software RAID level of a logical volume.
- For computer disk rows, this is the space that is unavailable on the file systems and the logical volumes on a computer disk.
- For computer space rows, this is the storage subsystem volume space that is unavailable on a computer.
- For storage subsystem space rows, this is the storage subsystem volume space that is unavailable on all monitored computers.

Monitored Computer Storage Space - Disk Storage reports:

Use these reports to view information about disks residing within computers that are monitored by IBM Tivoli Storage Productivity Center. This includes information about storage subsystem disks, known disks with serial numbers, and unknown disks without serial numbers.

Disk Storage reports are organized into the following categories:

All Disk Storage
Use these reports to view information about all the disks residing on computers that are monitored by Tivoli Storage Productivity Center.

Disk Storage On Storage Subsystems
Use these reports to view information about disks that Tivoli Storage Productivity Center has identified as storage subsystem volumes.

Disk Storage Not On Storage Subsystems
Use these reports to view information about disks that Tivoli Storage Productivity Center has not identified as storage subsystem volumes.

Related Topics

- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Disk Storage - All Disk Storage reports:

Use these reports to view information about all the disks residing on computers that are monitored by IBM Tivoli Storage Productivity Center.

See All Disk Storage reports - Column Descriptions for descriptions of all the columns that appear in All Disk Storage reports.

Use By clauses to filter the entities that appear in a report. For example, if you want to view an All Disk Storage report that displays specific computers, select **By Computer** and use the **Selection** and **Filter** options to determine which computers appear in that report.

You can customize the information in each report by relating the monitored entities to other storage. The following options for relating entities to storage in these reports includes the following:

- Use the **Relate To:** options in the By Computer, By Computer Group, and By Disk/Volume Group reports to relate storage entities to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space
  - Computer Disk Space
  - Distributed Computer Disk Space
- Use the **Relate To:** options in the By File System/Logical Volume and By File System Groups reports to relate storage entities to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space

Related Topics

- All Disk Storage reports - Column Descriptions
- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

All Disk Storage reports - Column Descriptions:

Use these reports to view information about all the disks residing on computers that are monitored by IBM Tivoli Storage Productivity Center. Use this topic to view descriptions of the columns that appear in these reports.

The following buttons and columns appear in the By clauses of All Disk Storage reports:

Note: Not all of the following columns appear in each By clause.

Selection...
Click this button to select the specific entities that you want to include in the report. For example, in the **By Filesystem/Logical Volume** report, click this button to select the specific file systems or logical volumes upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict the entities that you want to include in the report.

Relate <entity> To:
Select how you want to relate monitored storage entities in the report to other storage.

Magnifying glass icon 
This button appears on all **By <entity> Group** reports. Click this button to view a more detailed report about the entities in a group.

Available Disk Space
Amount of space on a disk that is available for use. For computer groups, this is the amount of space on disks that is available for use by the computers in a group. For disk and volume groups, this is the external disk space available for use within those groups.

This space includes the file system space that is available on the disk and the space that has not been allocated to a logical volume.

Available Disk/Volume Group Space
Amount of available space in a disk group or volume group.

Available File System/Logical Volume Space
Amount of space on the file system or logical volume that is available for use. For computers groups, this is the amount of space on file systems or logical volumes that is available for use by the computers within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that can be consumed. For disk groups and volume groups, this is the external logical volume or file system space available for use within those groups.

Computer
Name of a computer where a disk, file system, logical volume, or disk group is located.

Computer OS Type
The operating system running on a computer.

Consumed Disk Space
Amount of space on a disk that is consumed. For computer groups, this is the amount of disk consumed by the computers within those groups. For disk and volume groups, this is the external disk space consumed within those groups.

This space includes the file system space that is consumed on a disk. If a logical volume on this disk space has no file system, then this space includes the consumed portion of the logical volume that has been allocated to the disk.

Consumed Disk/Volume Group Space
Amount of consumed disk or volume group space.

Consumed File System or Logical Volume Space
Amount of space on the file system or logical volume consumed by the computer on a disk. For computer groups, this is the amount of space on

file systems or logical volumes that is consumed by the computers within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that is consumed. For disk groups and volume groups, this is the external logical volume or file system space consumed within those groups.

Correlated
Indicates **Yes** if Tivoli Storage Productivity Center is able to relate the host storage directly to a storage subsystem. You can relate this correlated storage back to the actual storage subsystem storage for a more detailed view of how the host storage is allocated.

If a computer disk, file system, or logical volume appears in the Disk Storage On Storage Subsystem reports but is not correlated, it might be the result of the following reasons:

- The storage subsystem was not probed. Run a probe to collect data about the storage subsystem then generate this report.
- The storage subsystem storage was unassigned from the host storage, but the host has not been probed. Run a probe to collect data about the host then generate this report.
- The storage subsystem storage was not correctly identified by a probe of the host. If this is the problem, check the probe log of the host for messages relating to SCSI command errors and the probe results of the storage subsystem.

Creator
ID of the user who created a group.

Disk Path
Path name of a disk as seen by the computer.

Disk Space
Total amount of space allocated to a computer disk. For computer groups, this is the total amount of space allocated to the disks on the computers within those groups. For disk and volume groups, this is the external disk space within the disk group or volume group.

Note: This space is distributed across the hypervisor disk space from where the virtual machine disk's file resides. The portion of space identified to a particular external hypervisor disk is reported only. This is true for all space values in this report.

Disk/Volume Group Name
Name of a disk group or volume group where a logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Disk/Volume Group Space
Total amount of disk group or volume group space. This is space inclusive of all storage space and applies to all space values related to disk groups or volume groups.

Disk/Volume Group Type
Type of a disk group.

Fibre Attached
Indicates **Yes** when Tivoli Storage Productivity Center determines that a disk is behind a fibre channel adapter port.

File System Type
Type of a file system, such as NTFS.

File System/Logical Volume Path
Name of a file system or logical volume path.

Group Name of a group.

Grown Defects
Number of defects detected on the disk since it was new.

Hypervisor Name
Name of the hypervisor where a virtual machine was created. This column displays **N/A** if not applicable to a disk.

Logical Volume Space
Total amount of space allocated to the logical volume on a disk. For computer groups, this is the total amount of space allocated to the logical volumes on disks within those groups. For file system groups, this is the total amount of space allocated to the logical volumes within those groups. For disk groups and volume groups, this is the external logical volume space within those groups.

This space includes the file system space on the disk reported if the entity displayed is a file system. The actual size of the file system on the disk reported is in addition to the consumed and available space.

Note: When relating a monitored entity to distributed file system/logical volume space, the space is distributed across the disks where a logical volume or file system resides. The portion identified to a disk is reported only.

(Disk) Manufacturer
The manufacturer of a disk.

Mapped to Storage Subsystem
Indicates **Yes** if Tivoli Storage Productivity Center has identified a disk as being located on a storage volume from a storage subsystem. The storage subsystem might no be monitored by Tivoli Storage Productivity Center.

(Disk) Model
Model number of a disk.

Number of Disks
Number of disks in a computer group. For disk groups and volume groups, this is the number of external disks within those groups.

Number of File Systems/Logical Volumes
Number of logical file systems or logical volumes within a group. For disk groups and volume groups, this is the number of external logical file systems or logical volumes within those groups.

The logical volumes underlying the file systems are not counted for the value in the column.

Overallocated
Indicates **Yes** if Tivoli Storage Productivity Center is unable to uniquely identify disk storage across more than one host. Check the probe log of the host for messages relating to SCSI command errors.

Path Type
Path type of the entity being reported: file system or a logical volume.

Primary Defects
Number of defects on the disk when it was new.

Shared
Indicates **Yes** if a file system or logical volume is shared by more than one computer.

Serial Number
Serial number of a disk.

Unavailable Disk Space
Amount of space on the disk not available to be consumed. For the computer groups, this is the amount of space on disks that is not available to be consumed by the computers in a group. For disk and volume groups, this is the external disk space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Unavailable Logical Volume Space
Amount of space on a logical volume that is not available for use by a computer. For computer groups, this is the amount of space on logical volumes that is not available for use by the computers within those groups. For file system groups, this is the amount of space on the logical volumes within those groups that is not available for use. For disk groups and volume groups, this is the external logical volume space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

VM Disk File
Virtual machine disk's file on a monitored hypervisor that represents the disk on the virtual machine. This column displays **N/A** if not applicable to a disk.

VMFS Disk
Hypervisor disk where the virtual machine disk resides. This column displays **N/A** if not applicable to a disk.

Note: If a virtual machine disk's file spans multiple hypervisor disks, duplicate entries for the virtual machine disk are generated. If the report is distributed, the space is distributed across the hypervisor disks that are not external.

VFMS Mount Point
Hypervisor file system's mount point where the virtual machine disk's file resides. This column displays **N/A** if not applicable to a disk.

Related Topics

- All Disk Storage reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Disk Storage - Disk Storage on Storage Subsystems reports:

Use these reports to view information about computer disks that IBM Tivoli Storage Productivity Center has identified as storage subsystem volumes.

See Disks On Storage Subsystems - Column Descriptions for descriptions of all the columns that appear in Disk Storage on Storage Subsystem reports.

Use By clauses to filter the entities that appear in a report. For example, if you want to view a report that displays specific computers, select **By Storage Subsystem** and use the **Selection** and **Filter** options to determine which storage subsystems appear in the generated report.

You can customize the information in each By clause by relating the monitored entities to other storage. The following options for relating storage entities (storage subsystems, computers, groups, file systems, logical volumes, disk groups, volume groups) to other storage in these reports includes the following:

- Use the **Relate Storage Subsystems To:** options in the By Storage Subsystem and By Storage Subsystem Group reports to relate storage subsystems to the following storage:
  - Computer Space
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space
  - Computer Disk Space
  - Distributed Computer Disk Space
- Use the **Relate Computers To:** and **Relate Disk/Volume Group To:** options in the By Computer, By Computer Group, and By Disk/Volume Group reports to relate storage computers, computer groups, disk groups, or volume groups to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space
  - Computer Disk Space
  - Distributed Computer Disk Space
  - Storage Subsystem Space (not available for disk groups and volume groups)
- Use the **Relate FileSystems/Logical Volumes To:** options in the By File System/Logical Volume and By Filesystem Group reports to relate file systems or logical volumes to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space

Considerations

Keep in mind the following considerations when generating Disk Storage on Storage Subsystems reports:

By Storage Subsystem (Relate Storage Subsystems To: File System/Logical Volume Space, Distributed File System/Logical Volume Space)

Duplicate entries appear in By Storage Subsystem reports for file systems or logical volumes that span multiple storage subsystems. The storage space shown in the report is adjusted in the totals. The storage space values include space that is not storage subsystem storage if a file system or logical volume spans disks that are both storage subsystem disks and other storage disks.

When relating storage subsystems or computers to distributed file system/logical volume space, the reports displays duplicate entries for file systems or logical volumes that span multiple disks, but the displayed storage space unique. The values for storage space exclude space that is not storage subsystem storage if a file system or logical volume spans disks that are both storage subsystem disks and other storage disks.

By Storage Subsystem (Relate Storage Subsystems To: Computer Disk Space, Distributed Computer Disk Space)

Duplicate entries appear in these reports for virtual machines whose disks spans multiple hypervisor disks. The storage space shown in the report is adjusted in the totals. The storage space values include space that is not storage subsystem storage if the virtual machine disk's file spans hypervisor disks that are external or the virtual machine disk's file spans other storage disks.

When relating storage subsystems or computers to distributed computer disk space, this report displays duplicate entries for virtual machines whose disks spans multiple hypervisor disks, but the storage space shown in the report is unique. The storage space values exclude space that is not storage subsystem storage if the virtual machine disk's file spans hypervisor disks that are external or the virtual machine disk's file spans other storage disks.

Related Topics

Keep in mind the following considerations when generating Disk Storage on Storage Subsystems reports:

- Disks On Storage Subsystems - Column Descriptions
- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Disks On Storage Subsystems - Column Descriptions:

Use these reports to view information about disks on monitored computers that IBM Tivoli Storage Productivity Center has identified as storage subsystem volumes.

The following buttons and columns appear in the By clauses of Disks On Storage Subsystems reports:

Note: Not all of the following columns appear in each By clause.

Selection...

Click this button to select the specific entities that you want to include in the report. For example, in the **By Storage Subsystem** report, click this button to select the specific storage subsystems upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict the entities that you want to include in the report.

Relate <entity> To:
Select how you want to relate monitored storage entities in the report to other storage. See Disk Storage - Disk Storage Not On Storage Subsystems reports for more information.

Magnifying glass icon 
This button appears on all **By <entity> Group** reports. Click this button to view a more detailed report about the entities in a group.

Available Disk Space
Amount of disk space on a storage subsystem that is available for a computer to consume. For computer groups, this is the amount of space on disks that is available for use by the computers within those groups. For storage subsystem groups, this is the amount of space on disks that is available for use by the storage subsystems within those groups. For disk and volume groups, this is the external disk space available for use within those groups.

This space includes the file system space that is available on the disk and the space that has not been allocated to a logical volume.

Available Disk/Volume Group Space
Amount of available space in a disk group or volume group.

Available File System/Logical Volume Space
Amount of space on the file system or logical volume that is available for use by the computer on a storage subsystem. For computers groups and storage subsystem groups, this is the amount of space on file systems or logical volumes that is available for use on storage subsystems by the entities within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that can be consumed. For disk groups and volume groups, this is the external logical volume or file system space available for use within those groups.

Computer
Name of a computer where a disk, file system, logical volume, or disk group is located.

Computer OS Type
The operating system running on a computer.

Consumed Disk Space
Amount of disk space on a storage subsystem that is consumed by a computer. For storage subsystem groups, this is the amount of consumed disk space within those groups. For computer groups, this is the amount of disk space consumed by the computers those groups. For disk and volume groups, this is the external disk space consumed within those groups.

This space includes the file system space that is consumed on a disk. If a logical volume on this disk space has no file system, then this space includes the consumed portion of the logical volume that has been allocated to the disk.

Consumed Disk/Volume Group Space
Amount of consumed disk or volume group space.

Consumed File System/Logical Volume Space
Amount of space on the file system or logical volume that is consumed by the computer on a storage subsystem. For computer groups and storage subsystem groups, this is the amount of space on file systems or logical volumes that is consumed by the entities within those groups. For file

system groups, this is the amount of storage subsystem space on file systems or logical volumes within those groups that is consumed. For disk groups and volume groups, this is the external logical volume or file system space consumed within those groups.

Correlated
Indicates **Yes** if Tivoli Storage Productivity Center is able to relate the host storage directly to a storage subsystem. You can relate this correlated storage back to the actual storage subsystem storage for a more detailed view of how the host storage is allocated.

If a computer disk, file system, or logical volume appears in the Disk Storage On Storage Subsystem reports but is not correlated, it might be the result of the following reasons:

- The storage subsystem was not probed. Run a probe to collect data about the storage subsystem then generate this report.
- The storage subsystem storage was unassigned from the host storage, but the host has not been probed. Run a probe to collect data about the host then generate this report.
- The storage subsystem storage was not correctly identified by a probe of the host. If this is the problem, check the probe log of the host for messages relating to SCSI command errors and the probe results of the storage subsystem.

Creator
ID of the user who created a group.

Disk Path
Path name of a disk as seen by the computer.

Disk Space
Total disk space on a storage subsystem seen by a computer. For computer groups, this is the total disk space within those groups that is associated with storage subsystems. For storage subsystem groups, this is the total disk space on the storage subsystems within those groups. For disk and volume groups, this is the external disk space within the disk group or volume group.

Note: This space is distributed across the hypervisor disk space from where the virtual machine disk's file resides. The portion of space identified to a particular external hypervisor disk is reported only. This is true for all space values in this report.

Disk/Volume Group Name
Name of a disk group or volume group where a logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Disk/Volume Group Space
Total amount of disk group or volume group space. This is space inclusive of all storage space and applies to all space values related to disk groups or volume groups.

Disk/Volume Group Type
Type of a disk group. For example, a volume group might be LVM or Veritas. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Fibre Attached
Indicates **Yes** when Tivoli Storage Productivity Center determines that a disk is behind a fibre channel adapter port.

File System Type
Type of a file system, such as NTFS.

File System/Logical Volume Path
Name of a file system or logical volume path.

Group Name of a group.

Grown Defects
Number of defects detected on the disk since it was new.

Hypervisor Name
Name of the hypervisor where a virtual machine was created. This column displays **N/A** if not applicable to a disk.

Logical Volume Space
Total amount of space allocated to the logical volume on a disk. For computer groups, this is the total logical volume space in computer groups that is associated with storage subsystems. For storage subsystem groups, this is the total logical volume space on the storage subsystems within those groups. For file system groups, this is the total amount of space allocated to the logical volumes within those groups. For disk groups and volume groups, this is the external logical volume space within those groups.

This space includes the file system space on the disk reported if the entity displayed is a file system. The actual size of the file system on the disk reported is in addition to the consumed and available space.

Note: When relating a monitored entity to distributed file system/logical volume space, the space is distributed across the disks where a logical volume or file system resides. The portion identified to a disk is reported only.

(Disk) Manufacturer
The manufacturer of a disk.

Mapped to Storage Subsystem
Indicates **Yes** if Tivoli Storage Productivity Center has identified a disk as being located on a storage volume from a storage subsystem. The storage subsystem might not be monitored by Tivoli Storage Productivity Center.

(Disk) Model
Model number of a disk.

Number of Disks
Number of disks assigned to a storage subsystem or group. For disk groups and volume groups, this is the number of external disks within those groups.

Number of File Systems/Logical Volumes
Number of logical file systems or logical volumes within a group. For disk groups and volume groups, this is the number of external logical file systems or logical volumes within those groups.

The logical volumes underlying the file systems are not counted for the value in the column.

Overallocated
Indicates **Yes** if Tivoli Storage Productivity Center is unable to uniquely identify disk storage across more than one host. Check the probe log of the host for messages relating to SCSI command errors.

Path Type
Path type of the entity being reported: file system or a logical volume.

Primary Defects
Number of defects on the disk when it was new.

Shared
Indicates **Yes** if a file system or logical volume is shared by more than one computer.

Serial Number
Serial number of a disk.

Storage Subsystem Manufacturer
The manufacturer of a storage subsystem.

Storage Subsystem Type
The type of storage subsystem.

Unavailable Disk Space
Amount of disk space on the storage subsystem that is not available for a computer to consume. For the computer groups, this is the amount of space on disks that is not available to be consumed by the computers within those groups. For the storage subsystem groups, this is the amount of space on disks that is not available to be consumed within those groups. For disk and volume groups, this is the external disk space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Unavailable Logical Volume Space
Amount of space on the logical volume that is not available for use by the computer on a storage subsystem. For computer groups and storage subsystem groups, this is the amount of space on logical volumes that is not available for use by the entities within those groups. For file system groups, this is the amount of space on the logical volumes within those groups that is not available for use. For disk groups and volume groups, this is the external logical volume space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

VM Disk File
Virtual machine disk's file on a monitored hypervisor that represents the disk on the virtual machine. This column displays **N/A** if not applicable to a disk.

VMFS Disk
Hypervisor disk where the virtual machine disk resides. This column displays **N/A** if not applicable to a disk.

Note: If a virtual machine disk's file spans multiple hypervisor disks, duplicate entries for the virtual machine disk are generated. If the report is distributed, the space is distributed across the hypervisor disks that are not external.

VFMS Mount Point
Hypervisor file system's mount point where the virtual machine disk's file resides. This column displays **N/A** if not applicable to a disk.

Related Topics

- Disks On Storage Subsystems reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Disk Storage - Disk Storage Not On Storage Subsystems reports:

Use these reports to view information about computer disks that IBM Tivoli Storage Productivity Center has not identified as storage subsystem volumes.

These reports are organized into the following categories:

Disks With Serial Numbers
These reports display information about non-storage subsystem disks that Tivoli Storage Productivity Center has identified with a valid serial number.

Disks Without Serial Numbers
These reports display information about non-storage subsystem disks that Tivoli Storage Productivity Center has not identified with serial numbers that uniquely identify those disks. This might occur on internal disks within a hypervisor and virtual machines that have storage created from internal disks within a hypervisor. If this problem occurs for other types of storage, check the probe log of the host for messages relating to SCSI command errors.

See Disks Not On Storage Subsystems - Column Descriptions for descriptions of all the columns that appear in Disk Storage Not On Storage Subsystems reports.

Use By clauses to filter the entities that appear in a report. For example, if you want to view a Disk Storage Not On Storage Subsystems report that displays specific computers, select **By Computer** and use the **Selection** and **Filter** options to determine which computers appear in that report.

You can customize the information in each By clause by relating the monitored entities to other storage. The following options for relating storage entities (computers, computer groups, file systems, file system groups, logical volumes, disk groups, volume groups) to other storage in these reports includes the following:

- Use the **Relate Computers To:** and **Relate Disk/Volume Group To:** options in the By Computer, By Computer Group, and By Disk/Volume Group reports to relate computers, computer groups, disk groups, and volume groups to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space
  - Computer Disk Space
  - Distributed Computer Disk Space

- Use the **Relate FileSystems/Logical Volumes To:** options in the By File System/Logical Volume and By Filesystem Group reports to relate file systems or logical volumes to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space

Related Topics

- Disks Not On Storage Subsystems - Column Descriptions
- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Disks Not On Storage Subsystems - Column Descriptions:

Use Disks Not On Storage Subsystems reports to view information about disk storage that is not associated with storage subsystems. IBM Tivoli Storage Productivity Center has not identified the disks in this report as storage subsystem volumes. Use this topic to view descriptions of the columns that appear in these reports.

The following buttons and columns appear in the By clauses of Disks Not On Storage Subsystems reports:

Note: Not all of the following columns appear in each By clause.

Selection...
Click this button to select the specific entities that you want to include in the report. For example, in the **By Filesystem/Logical Volume** report, click this button to select the specific file systems or logical volumes upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict the entities that you want to include in the report.

Relate <entity> To:
Select how you want to relate monitored storage entities in the report to other storage. See Disk Storage - Disk Storage Not On Storage Subsystems reports for more information.

Magnifying glass icon 
This button appears on all **By <entity> Group** reports. Click this button to view a more detailed report about the entities in a group.

Available Disk Space
Amount of space on a disk that is available for use by a computer. For computer groups, this is the amount of space on disks that is available for use by the computers in a group. For disk and volume groups, this is the external disk space available for use within those groups.

This space includes the file system space that is available on the disk and the space that has not been allocated to a logical volume.

Available Disk/Volume Group Space
Amount of available space in a disk group or volume group.

Available File System/Logical Volume Space
Amount of space on the file system or logical volume that is available for use on a disk. For computers groups, this is the amount of space on file systems or logical volumes that is available for use by the computers within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that can be consumed. For disk groups and volume groups, this is the external logical volume or file system space available for use within those groups.

Computer
Name of a computer where a disk, file system, logical volume, or disk group is located.

Computer OS Type
The operating system running on a computer.

Consumed Disk Space
Amount of space on a disk that is consumed. For computer groups, this is the amount of disk consumed by the computers in those groups. For disk and volume groups, this is the external disk space consumed within those groups.

This space includes the file system space that is consumed on a disk. If a logical volume on this disk space has no file system, then this space includes the consumed portion of the logical volume that has been allocated to the disk.

Consumed Disk/Volume Group Space
Amount of consumed disk or volume group space.

Consumed File System/Logical Volume Space
Amount of space on the file system or logical volume consumed by the computer on a disk. For computer groups, this is the amount of space on file systems or logical volumes that is consumed by the computers within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that is consumed. For disk groups and volume groups, this is the external logical volume or file system space consumed within those groups.

Creator
ID of the user who created a group.

Disk Path
Path name of a disk as seen by the computer.

Disk Space
Total amount of space allocated to a computer disk. For computer groups, this is the total amount of space allocated to the disks on the computers within those groups. For disk and volume groups, this is the external disk space within the disk group or volume group.

Note: This space is distributed across the hypervisor disk space from where the virtual machine disk's file resides. The portion of space identified to a particular external hypervisor disk is reported only. This is true for all space values in this report.

Disk/Volume Group Name
Name of a disk group or volume group where a logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Disk/Volume Group Space
Total amount of disk group or volume group space. This is space inclusive of all storage space and applies to all space values related to disk groups or volume groups.

Disk/Volume Group Type
Type of a disk group. For example, a volume group might be LVM or Veritas. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Fibre Attached
Indicates **Yes** when Tivoli Storage Productivity Center determines that a disk is behind a fibre channel adapter port.

File System Type
Type of a file system, such as NTFS.

File System/Logical Volume Path
Name of a file system or logical volume path.

Group Name of a group.

Grown Defects
Number of defects detected on the disk since it was new.

Hypervisor Name
Name of the hypervisor where a virtual machine was created. This column displays **N/A** if not applicable to a disk.

Logical Volume Space
Total amount of space allocated to the logical volume on a disk. For computer groups, this is the total amount of space allocated to the logical volumes on disks within those groups. For file system groups, this is the total amount of space allocated to the logical volumes within those groups. For disk groups and volume groups, this is the external logical volume space within those groups.

This space includes the file system space on the disk reported if the entity displayed is a file system. The actual size of the file system on the disk reported is in addition to the consumed and available space.

Note: When relating a monitored entity to distributed file system/logical volume space, the space is distributed across the disks where a logical volume or file system resides. The portion identified to a disk is reported only.

(Disk) Manufacturer
The manufacturer of a disk.

(Disk) Model
Model number of a disk.

Number of Disks
Number of disks in a computer group. For disk groups and volume groups, this is the number of external disks within those groups.

Number of File Systems/Logical Volumes
Number of logical file systems or logical volumes within a group. For disk groups and volume groups, this is the number of external logical file systems or logical volumes within those groups.

The logical volumes underlying the file systems are not counted for the value in the column.

Overallocated
Indicates **Yes** if Tivoli Storage Productivity Center is unable to uniquely identify disk storage across more than one host. Check the probe log of the host for messages relating to SCSI command errors.

Path Type
Path type of the entity being reported: file system or a logical volume.

Primary Defects
Number of defects on the disk when it was new.

Shared
Indicates **Yes** if a file system or logical volume is shared by more than one computer.

Serial Number
Serial number of a disk.

Unavailable Disk Space
Amount of space on the disk not available to be consumed. For the computer groups, this is the amount of space on disks that is not available to be consumed by the computers in a group. For disk and volume groups, this is the external disk space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Unavailable Logical Volume Space
Amount of space on a logical volume that is not available for use by a computer. For computer groups, this is the amount of space on logical volumes that is not available for use by the computers within those groups. For file system groups, this is the amount of space on the logical volumes within those groups that is not available for use. For disk groups and volume groups, this is the external logical volume space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

VM Disk File
Virtual machine disk's file on a monitored hypervisor that represents the disk on the virtual machine. This column displays **N/A** if not applicable to a disk.

VMFS Disk
Hypervisor disk where the virtual machine disk resides. This column displays **N/A** if not applicable to a disk.

Note: If a virtual machine disk's file spans multiple hypervisor disks, duplicate entries for the virtual machine disk are generated. If the report is distributed, the space is distributed across the hypervisor disks that are not external.

VFMS Mount Point
Hypervisor file system's mount point where the virtual machine disk's file resides. This column displays **N/A** if not applicable to a disk.

Related Topics

- Disks Not On Storage Subsystems reports overview

- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Disk Storage - Non-Disk Storage reports:

Use these reports to view information about storage space that does not reside on disks that are monitored by IBM Tivoli Storage Productivity Center. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory.

Use By clauses to filter the entities that appear in a report. For example, if you want to view a report that displays specific computers, select **By Computer** and use the **Selection** and **Filter** options to determine which computers appear in the generated report.

The following By clauses are available for Non-Disk Storage reports:

- By Computer
- By Computer Group
- By Filesystem/Logical Volume
- By Filesystem Group
- By Disk/Volume Group

You can customize the information in each By clause by relating the monitored entities to other storage. The following options for relating storage entities (computers, computer groups, file systems, file system groups, logical volumes, disk groups, volume groups) to other storage in these reports includes the following:

- File System or Logical Volume Space
- Distributed File System or Logical Volume Space

Related Topics

- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

Non-Disk Storage - By Computer report, By Filesystem/Logical Volume report:

Use these reports to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. Select **By Computer** if you want to determine which computers are displayed in the report. Select **By Filesystem/Logical Volume** if you want to determine which file systems or logical volumes appear in the report.

Selection...

For the **By Computer** reports, click this button to select the specific computers upon which you want to report. For the **By Filesystem/Logical Volume** reports, click this button to select the specific file systems or logical volumes upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict computers or file system/logical volumes upon which you want to report.

Relate <entity> To:
Select how you want to relate computers or file systems in the report to other storage entities. See]Disk Storage - Non-Disk Storage reports for more information.

Computer
Name of a computer.

File System/Logical Volume Path
Name of a file system or logical volume path.

Note: Unique file system/logical volume path entries are generated in this report. This occurs because if the storage space is not distributed, there are no underlying disks no need to report distribution. If storage space is distributed, space is distributed on logical volumes that do not have underlying disks only.

Logical Volume Space
Total amount of space allocated to the logical volume on a disk. This space includes the file system space if the entity displayed is a file system. The actual size of the file system is in addition to the consumed and available space.

Note: When relating computers to distributed file system/logical volume space, the space is distributed across the logical volumes where the file system or parent logical volume resides. The portion identified to a particular file system or parent logical volume is reported only. If a file system does not have an underlying logical volume, the logical volume space is the size of that file system.

Consumed File System/Logical Volume Space
Amount of space on a file system or logical volume that is consumed by a computer.

Available File System/Logical Volume Space
Amount of space on a file system or logical volume that is available for use by a computer.

Unavailable Logical Volume Space
Amount of space on the logical volume that is not available for use by a computer. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Computer OS Type
Operating system running on a computer.

Disk/Volume Group Name
Name of a disk or volume group where the logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Disk/Volume Group Type
Type of a disk or volume group where the logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

File System Type
Type of a file system, such as NTFS.

Path Type
Path type of an entity: file system or a logical volume.

Shared
Indicates **Yes** if a file system or logical volume is shared by more than one computer.

Non-Disk Storage - By Computer Group report, By Filesystem Group report:

Use these reports to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. Select **By Computer Group** if you want to determine which computer groups are displayed in the report. Select **By Filesystem Group** if you want to determine which file system groups appear in the report.

Selection...
For the **By Computer Group** reports, click this button to select the specific computer groups upon which you want to report. For the **By Filesystem Group** reports, click this button to select the specific file system group upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict groups upon which you want to report.

Relate <entity> To:
Select how you want to relate groups in the report to other storage entities. See Disk Storage - Non-Disk Storage reports for more information.

A unique row appears for each group in this report.

Magnifying glass icon 
Click this button to view a more detailed report about the computers or file systems in a group.

Creator
ID of the user who created the group.

Group Name of a group.

Logical Volume Space
Total amount of space allocated to the computers or file systems in a group. This space includes the file system space if the entity displayed is a file system. The actual size of the file system is in addition to the consumed and available space.

Note: When relating groups to distributed file system/logical volume space, this space is distributed across the disks from where the logical volume or file system reside. This is true for all space values in the Distributed File System/Logical Volume Space report.

Consumed File System/Logical Volume Space
Amount of space on file systems or logical volumes within a group that is consumed.

Available File System/Logical Volume Space
Amount of space on file systems or logical volumes within a group that can be consumed.

Unavailable Logical Volume Space
Amount of space on the logical volumes in a group that is not available for use. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Number of File Systems/Logical Volumes
Number of file systems or logical volumes in a group.

Non-Disk Storage - By Computer report, By Filesystem/Logical Volume report:

Use these reports to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. Select **By Computer** if you want to determine which computers are displayed in the report. Select **By Filesystem/Logical Volume** if you want to determine which file systems or logical volumes appear in the report.

Selection...
For the **By Computer** reports, click this button to select the specific computers upon which you want to report. For the **By Filesystem/Logical Volume** reports, click this button to select the specific file systems or logical volumes upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict computers or file system/logical volumes upon which you want to report.

Relate <entity> To:
Select how you want to relate computers or file systems in the report to other storage entities. See]Disk Storage - Non-Disk Storage reports for more information.

Computer
Name of a computer.

File System/Logical Volume Path
Name of a file system or logical volume path.

Note: Unique file system/logical volume path entries are generated in this report. This occurs because if the storage space is not distributed, there are no underlying disks no need to report distribution. If storage space is distributed, space is distributed on logical volumes that do not have underlying disks only.

Logical Volume Space
Total amount of space allocated to the logical volume on a disk. This space includes the file system space if the entity displayed is a file system. The actual size of the file system is in addition to the consumed and available space.

Note: When relating computers to distributed file system/logical volume space, the space is distributed across the logical volumes where the file system or parent logical volume resides. The portion identified to a particular file system or parent logical volume is reported only. If a file system does not have an underlying logical volume, the logical volume space is the size of that file system.

Consumed File System/Logical Volume Space
Amount of space on a file system or logical volume that is consumed by a computer.

Available File System/Logical Volume Space
Amount of space on a file system or logical volume that is available for use by a computer.

Unavailable Logical Volume Space
Amount of space on the logical volume that is not available for use by a computer. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Computer OS Type
Operating system running on a computer.

Disk/Volume Group Name
Name of a disk or volume group where the logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

Disk/Volume Group Type
Type of a disk or volume group where the logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

File System Type
Type of a file system, such as NTFS.

Path Type
Path type of an entity: file system or a logical volume.

Shared
Indicates **Yes** if a file system or logical volume is shared by more than one computer.

Non-Disk Storage - By Computer Group report, By Filesystem Group report:

Use these reports to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. Select **By Computer Group** if you want to determine which computer groups are displayed in the report. Select **By Filesystem Group** if you want to determine which file system groups appear in the report.

Selection...
For the **By Computer Group** reports, click this button to select the specific computer groups upon which you want to report. For the **By Filesystem Group** reports, click this button to select the specific file system group upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict groups upon which you want to report.

Relate <entity> To:
Select how you want to relate groups in the report to other storage entities. See Disk Storage - Non-Disk Storage reports for more information.

A unique row appears for each group in this report.

Magnifying glass icon 
Click this button to view a more detailed report about the computers or file systems in a group.

Creator
ID of the user who created the group.

Group Name of a group.

Logical Volume Space
Total amount of space allocated to the computers or file systems in a group. This space includes the file system space if the entity displayed is a file system. The actual size of the file system is in addition to the consumed and available space.

Note: When relating groups to distributed file system/logical volume space, this space is distributed across the disks from where the logical volume or file system reside. This is true for all space values in the Distributed File System/Logical Volume Space report.

Consumed File System/Logical Volume Space
Amount of space on file systems or logical volumes within a group that is consumed.

Available File System/Logical Volume Space
Amount of space on file systems or logical volumes within a group that can be consumed.

Unavailable Logical Volume Space
Amount of space on the logical volumes in a group that is not available for use. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Number of File Systems/Logical Volumes
Number of file systems or logical volumes in a group.

Disk Storage on Storage Subsystems reports - By Disk/Volume Group:

Use this report to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. This report enables you to determine which disk groups or volume groups are displayed.

Selection...
Click this button to select the specific disk groups or volume groups upon which you want to report.

Filter Click this button to apply conditions to the columns within a report to further select or restrict disk groups or volume groups upon which you want to report.

Relate Disk/Volume Group To:
Select how you want to relate the non-disk storage in disk groups or volume groups to other storage entities. See Disk Storage - Non-Disk Storage reports for more information.

A unique row appears for each disk group or volume group.

Magnifying glass icon 
Click this button to view a more detailed report about the disk or file system within the disk group or volume group.

Disk/Volume Group Name
Name of a group.

Disk/Volume Group Type
Type of a group. For example, a volume group might be LVM or Veritas.

Computer
Name of the computer where a disk group or volume group resides.

Computer OS Type
Operating system running on a computer.

Disk/Volume Group Space
Total amount of disk group or volume group space. This is space inclusive of all storage space and applies to all space values related to disk groups or volume groups.

Consumed Disk/Volume Group Space
Amount of consumed disk group or volume group space.

Available Disk/Volume Group Space
Amount of available disk group or volume group space.

Logical Volume Space
External logical volume space within a disk group or volume group.

Consumed File System/Logical Volume Space
External logical volume or file system space that is consumed within a disk group or volume group.

Available File System/Logical Volume Space
External logical volume or file system space that is available for use within a disk group or volume group.

Unavailable Logical Volume Space
External logical volume space that is not available for use within a disk group or volume group. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

Number of File Systems/Logical Volumes
Number of external logical file systems or logical volumes within a disk or volume group. The logical volumes underlying the file systems are not counted.

Additional reporting support and information

Learn about the additional storage devices upon which you can report within Data Manager.

IBM FAStT Subsystem support: The IBM FAStT subsystem reporting feature further extends the subsystem reporting capabilities of Data Manager. You can use the Reporting facility to view information about the capacity, allocation, and volumes of an IBM FAStT subsystem. You can view IBM FAStT subsystem information by using the following menu options to generate reports:

- **Reporting → Asset → By Storage Subsystem →** <IBM FAStT subsystem name>
- **Reporting → Asset → System-wide → Storage Subsystems**
- **Reporting → Asset → System-wide → Volumes**
- **Reporting → Asset → System-wide → Disks**

Note: Data Manager shows an incorrect value for the **Storage Pool Free Space** column in IBM FAStT subsystem reports (for example, **Asset → By Storage Subsystem → <FAStT_subsystem_name>**, and **Asset → System-wide → Storage**

Subsystems). The value for **Storage Pool Free Space** in these reports reflects the space from one Free Capacity node rather than all the free space within the defined arrays on a FAStT subsystem.

Hitachi Data Systems Subsystem support: Use the Reporting facility to report on Hitachi Data Systems (HDS) storage subsystems. You can view reports that contain information about the capacity, allocation, and volumes of an HDS subsystem.

Any storage pools or volumes created using the HDS **HiCommand** software will not be reported through the HDS CIMOM provider. **HiCommand** is the native interface for HDS subsystems. If you configure an HDS subsystem through **HiCommand** , IBM Tivoli Storage Productivity Center will not be able to gather information about storage pools and volumes from that subsystem. The HDS CIM agent provider will return all storage pools created through the **HiCommand** interface as a single pool containing all of the **HiCommand** created pools. This single pool is always reported as Raid 5.

Note: Data Manager does not report on the raw capacity of disks internal to the HDS subsystem. Information on the internal disks of the HDS subsystem is not available through the CIMOM interface provided by the Hitachi HiCommand Device Manager.

You can view HDS storage subsystem information by using the following menu options to generate reports:

- **Reporting → Asset → By Storage Subsystem →** <HDS subsystem name>
- **Reporting → Asset → System-wide → Storage Subsystems**
- **Reporting → Asset → System-wide → Volumes**
- **Reporting → Asset → System-wide → Disks**
- **Reporting → TPC-wide Storage Space → Disk Space**

HP StorageWorks Subsystem support: Use the Reporting facility to report on HP StorageWorks storage subsystems. You can view reports that contain information about the capacity, allocation, and volumes of an HP StorageWorks subsystem.

You can view HP StorageWorks storage subsystem information by using the following menu options to generate reports:

- **Reporting → Asset → By Storage Subsystem →** <HP StorageWorks subsystem name>
- **Reporting → Asset → System-wide → Storage Subsystems**
- **Reporting → Asset → System-wide → Volumes**
- **Reporting → Asset → System-wide → Disks**

IBM Tivoli Storage SAN File Systems support: You can also perform scans and probes to gather information about IBM Tivoli Storage SAN File Systems and generate reports based on that information using the Reporting facility. You can view information gathered by probes of the SAN File System by using the following menu options to generate reports:

- **Reporting → Asset → By OS Type → IBM SAN File System**
- **Reporting → Asset → By Computer**
- **Reporting → TPC-wide Storage Space → File System Space**
- **Reporting → TPC-wide Storage Space → Consumed File System Space**
- **Reporting → TPC-wide Storage Space → Available File System Space**

Note: Due to the complexity of the SAN File System, the values displayed for total capacity, used space and free space in the Capacity and Asset reports might not be accurate. The values displayed are the values returned by the and they may not represent actual values for your SAN File System. We recommend using the SAN File System admin console to obtain the actual values.
You can view information gathered by scans of the SAN File System by using the following menu options to generate reports:

- **Reporting → Usage**
- **Reporting → Usage Violations**
- **Reporting → Backup**

Refer to the *Creating Directory Groups* section in the *Monitoring* chapter for more information about how to configure SAN File System support within the Reporting facility.

Report Array, Loop, and Adapter Pair for IBM Tivoli Storage Enterprise Storage Server Disk Groups support: Disk Manager provides report information on Array, Loop, and Adapter Pairs for Tivoli Storage Enterprise Storage Server Disk Groups. This information is available when managing an Tivoli Storage Enterprise Storage Server through IBM CIM Agent (ICAT) for Tivoli Storage Enterprise Storage Server version 1.2.0.13. Contact IBM Customer Support to obtain IBM CIM Agent for Tivoli Storage Enterprise Storage Server version 1.2.0.13.

You can view Array, Loop, and Adapter Pair information by using the menu option **Reporting → Asset Reporting → By Storage Subsystem** to generate reports.

Calculating the Size for HP/Compaq volumes:

HP HSG Element Manager calculates the size for HP or Compaq volumes by using 1,000 bytes to a kilobyte. Data Manager does not use 1,000 bytes to the kilobyte to determine the size of a HP or Compaq volume. Instead, it uses the following method:

1. Retrieve the value for the size of the volume from HSG Element Manager.
2. Multiply the value by 1,000,000 to determine the value in bytes.
3. Divide the value in bytes by 1,024 (for megabytes and gigabytes accordingly) to calculate the true bytes value.

For example, if HSG Element Manager reports the size of a volume to be 62,000 megabytes, Data Manager takes that value and multiplies it by 1,000,000 to calculate the true value of 62,000,000,000. To determine the true value in megabytes, it divides that number by 1,024 to the second power, which is 59, 128 megabytes. To determine the true value in gigabytes, it divides that number by 1,024 to the third power, which is 57.7 gigabytes.

VMware capacity reports:

You must probe both the ESX Server and the Data agent in 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 (this is 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 will be shown as follows:

Column heading	Capacity	Used Space (this is calculated as capacity minus free space, any negative values will be 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 will be shown as follows:

Column heading	Capacity	Used Space (this is 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

Data Manager for Databases reports

Use Data Manager for Databases for reports to view both overview and detailed information about your RDBMS storage resources, including instances, tablespaces, databases, devices, containers, data files, fragments, tables, control files, redo logs, archive log directories, users. These reports are constructed from the statistical information gathered by data collection jobs and accumulated in the database repository.

Using Data Manager for Databases reports

Learn how to use Data Manager for Databases reports to view information about the instances, tablespaces, databases, devices, containers, data files, fragments, tables, control files, redo logs, archive log directories, and users within your environment.

Before you begin:

Learn about information you should know before generating and viewing Data Manager for Databases reports.

Before you can use reports to view detailed information about tables in an RDBMS, you must have 1) defined a table group that contains the tables about which you want to view storage information and 2) scheduled and run a scan against that table group. See "Working with groups of resources" on page 236 for information about using groups. See "Collecting storage usage and trending information (scans)" on page 210 for information about scans and how to create them.

Data Manager for Databases - Report Categories:

Use the report categories within Data Manager for Databases reports to collect information about RDBMS storage assets, capacity, usage, and usage violations.

The following table dsecribes the categories of reports available in Data Manager for Databases.

Report Category	Use these reports to do the following	What job collects information for these reports?
Asset Reports	View the information gathered by probes. While running probes, Data Manager for Databases itemizes the information for each type of database. • Oracle: instances, tablespaces, tables, data files, control files, redo logs, and archive log directories • SQL Server: instances, databases, tables, and data files • Sybase: instances, devices, databases, tables, and fragments. • UDB: instances, databases, tablespaces, tables, containers, and users	Probes

Report Category	Use these reports to do the following	What job collects information for these reports?
Capacity Reports	View the storage capacity of different entities within a RDBMS. Specifically, use Capacity reports to view: • **Instance storage capacity within all DBMSs.** If you are monitoring the instances for different DBMSs, you can use the **All DBMSs** node to view a summary of the storage capacity for all those instances in a single view. • **Oracle databases** – How much storage you have for instances – How much storage is being used for data files – What percentage of a data file's storage is free for expansion • **SQL Server databases** – How much storage you have for instances – How much storage is being used for data files – What percentage of a data file's storage is free for expansion • **Sybase databases** – How much storage you have for instances – How much storage is being used for devices – What percentage of a device's storage is free for expansion – How much storage is being used for fragments – What percentage of a fragment's storage is free for expansion • **UDB databases** – How much allocated storage you have for a UDB Instance and the objects it contains – How much storage is being used for a UDB Instance and the objects it contains – What percentage of UDB storage is free for expansion The Capacity reporting function provides you with the ability to view the storage capacity, used space, and free space for your RDBMS objects in a single view.	Probes

Report Category	Use these reports to do the following	What job collects information for these reports?
Usage Reports	Monitor and report on the usage and growth of the consumption of your DBMS storage: • **Storage usage within all DBMSs.** If you are monitoring the multiple DBMSs, use this node to view common storage usage information for databases, tablespaces, tables, and users across all monitored DBMSs. • **Oracle databases** – View detailed information for tablespaces including their total size, used space, empty used space, # of tables and indexes, etc. – Identify the tables consuming the most space within your environment. – Identify the segments with the most extents and most empty used space. – Associate users with the amount of storage they are consuming on specific databases, tablespaces, groups of tablespaces, specific computers, groups of computers, and throughout the entire network. • **SQL Server databases** – View detailed information for databases including their total size, used space, empty used space, # of tables and indexes, etc. – Identify the tables consuming the most space within your environment. – Associate users with the amount of storage they are consuming on specific databases, groups of databases, specific computers, groups of computers, and throughout the entire network. • **Sybase databases** – View detailed information for databases including their total size, used space, empty used space, # of tables and indexes, etc. – Identify the tables consuming the most space within your environment. – Associate users with the amount of storage they are consuming on specific databases, groups of databases, specific computers, groups of computers, and throughout the entire network. • **UDB databases** – View detailed information for tablespaces including their total size, used space, empty used space, # of tables and indexes, etc. – Identify the tables consuming the most space within your environment. You can view this information at the tablespace, database, and computer level, as well as at a network-wide level. – Associate users with the amount of storage they are consuming on specific databases, tablespaces, groups of tablespaces, specific computers, groups of computers, and throughout the entire network.	Scans
Usage Violation Reports	Create and enforce corporate storage usage policies. You can report on violations of quotas that are set at the network-wide, instance, tablespace, and database levels.	Scans, Quotas

Reporting on RDBMS Objects:

Learn about the RDBMS objects upon which you can report using Data Manager for Databases reports.

The following table lists the RDBMS objects whose storage usage is monitored by Data Manager for Databases.

Oracle	SQL Server	Sybase	DB2/UDB
Instances	**Instances**	**Instances**	**Instances**
Databases	**Databases**	**Devices**	**Databases**
Tablespaces	**Datafiles**	**Databases**	**Tablespaces**
Datafiles	**Tables**	**Fragments**	**Containers**
Tables	**Indexes**	**Tables**	**Tables**
Indexes	**Log Files**	**Indexes**	**Indexes**
Extents			
Segments			
Redo Log files			
Control Files			
Archive Log Directories			

Note:

1. Each of the major reporting categories in the Reporting facility can be expanded in the navigation tree to reveal additional sub-categories of reporting. Some of the nodes within the reporting tree will appear only if you set up access for the RDBMS to which that node applies. For example, if you set up for monitoring Oracle and SQL Server instances only, nodes representing the other RDBMSs supported by Data Manager for Databases (such as Sybase and UDB) will not appear within the navigation tree. If you have set up access for only one RDBMS, the **All DBMSs** node will also not be available in some nodes.
2. Data Manager for Databases does not currently support the monitoring of clustered database applications.

Using By clauses:

Use the By clauses associated with Data Manager for Databases reports as a method for viewing the data in those reports from different perspectives.

As you drill down through Data Manager for Databases's reports, you will notice a set of options that repeats for each report type. These *By Clauses* allow you to generate variations of the reports with different perspectives and groupings of the data. Some of the available By Clauses are described in the following table:

By Clause	Description
By DBMS Type	Use this clause to select the DBMS that you want to report upon.
By Instance	Use this clause to select the instances that you want to report upon and the report that will be generated based upon those instances.
By Device (Sybase only)	Use this clause to select the devices that you want to report upon and the report that will be generated based upon those devices.
By Database	Use this clause to select the database that you want to report upon and the report that will be generated based upon those databases.
By Tablespace	Use this clause to select the tablespaces that you want to report upon and the report that will be generated based upon those tablespaces.

By Clause	Description
By Tablespace Group	Use this clause to select the tablespace groups that you want to report upon and the report that will be generated based upon those tablespace groups.
By Fragments (Sybase only)	Use this clause to select the fragments that you want to report upon and the report that will be generated based upon those fragments.
By Container (UDB only)	Use this clause to select the containers that you want to report upon and the report that will be generated based upon those containers.
By Computer	Use this clause to select the computers that you want to report upon and the report that will be generated based upon those computers.
By Computer Group	Use this clause to select the computer groups that you want to report upon and the report that will be generated based upon those computer groups.
By User	Use this clause to select the users that you want to report upon and the report that will be generated based upon those users.
Network Wide	The reports and charts that you see will be generated on all data, network wide. This allows you to get a global view of the storage and storage usage for your whole enterprise.

Asset reports

Use Asset reports to view information about the instances, devices, databases, tablespaces, tables, containers, data files, fragments, control files, redo logs, and archive log directories, and general hardware inventory of RDBMS storage assets within your environment.

While running probes and scans, Data Manager for Databases itemizes your RDBMS storage assets and provides a hierarchical view of that information so you can drill down to view your those assets in greater detail. You can view information about RDBMS storage resources through the following views:

Asset Reporting View	Use these reports to...
Asset > By Computer	View the instances and their related objects based upon the computers on which they reside.
Asset > By DBMS Type	This node appears if you are using Data Manager for Databases to monitor more than one RDBMS. Expand this node to view asset information about instances organized according to their DBMS.
Asset > System-wide	Investigate all of your instances, databases, tablespaces, tables, container, data files, users and other assets for your *entire* system. **Note:** (UDB only) For the **All DBMSs** reports available under the **Asset > System-wide** reporting nodes, the Database-Tablespaces view shows information about UDB tablespaces only.

By Computer, By DBMS Type:

Learn about the reports available within the **By Computer** and **By DBMS Type** nodes within Asset reporting.

Asset Report Type	Use these reports to...	What's in these reports?
By Computer By Computer > *machine_name* > *instance/database_name* **By DBMS Type** By DBMS Type > *dbms_type* > *instance/database_name*	Drill down through your DBMS storage assets in a hierarchical manner according to the computers on which they reside or the DBMS to which they belong.	Depending on the DBMS of the Instance you want to view, you can drill down and expose its: • Tablespaces • Databases • Devices • Monitored tables • Containers • Data files • Fragments • Control files • Redo logs • Archive log directories

Asset Report Type	Use these reports to...	What's in these reports?
Instance Information **By Computer** By Computer > *machine_name* > *instance_name* **By DBMS Type** By DBMS Type > *dbms_type* > *instance_name*	View information about all of your monitored instances. Specifically, you can view information about the following information based upon DBMS type: **SQL Server** • Databases • Monitored Tables • Data Files **Oracle** • Tablespaces • Monitored Tables • Data Files • Control Files • Redo Logs • Archive Log Directories **Sybase** • Devices • Databases • Monitored Tables • Fragments **UDB** • Instances • Databases • Tablespaces • Tables • Containers • Users	**SQL Server** • RDBMS Type, Name, Version, Loginid • Number of Databases, Number of Datafiles, Number of Log Files • Total Size of all Storage, Capacity of Data Files, Free Space of Data Files, Capacity of Log Files, Free Space of Log Files • Probe Status, Last Probe Time **Oracle** *Database Information tab* • Name, Owner, Block Size, Log Mode, Status • Maximum Number of Instances, Number of Tablespaces, Maximum Number of Data Files, Number of Data Files, Maximum Number of Log Files, Number of Log Files • Total Size of all Storage, Capacity Of Data Files, Free Space of Data Files • Create Time, Discovered Time, Dropped Time *Instance Information tab* • RDBMS Type, Name, Version, Home Directory, Loginid, Port, Probe Status • Last Probe Time, Last Start Time, Discovered Time **Sybase** • RDBMS Type, Name, Version, Home Directory, Loginid, Port • Number of Databases, Number of Devices, Number of Fragments, Number of Log Fragments • Total Size of all Storage, Mirror Overhead, Capacity of Data Fragments, Free Space in Data Fragments, Capacity of Log Fragments, Free Space in Log Fragments • Probe Status, Last Probe Time **UDB** • RDBMS Type, Name, Version, Home Directory, Loginid, Port • Number of Tablespaces, Number of Containers, Number of Log Files • Total Size of all Storage, Container Capacity, DMS Container Free Space, Total Size of all Log Files • Probe Status, Last Probe Time

Asset Report Type	Use these reports to...	What's in these reports?
Oracle Detailed Instance Information **By Computer** By Computer > *machine_name* > *instance_name*	View information about the objects associated with a selected Oracle Instance.	• **Tablespaces**: When you expand the **Tablespaces** node, click the name of a tablespace to view information about that tablespace and its related data file. • **Monitored Tables**: When you expand the **Monitored Tables** node, click on the name of a table to view its information. Only the tables that belong to Table Groups included in scans will be available to view through this node. • **Data Files**: Click the **Data Files** node, click the name of a data file to view information about the data files associated with the Instance, including: Data File Path, Used Space, Percent Used, Free Space, Total Size, Coalesced Extents, Minimum Free Extent Size, Maximum Free Extent Size, and Create Time. • **Control Files**: Click this node to view information on the control files associated with the Instance, including the File Name (full path) and File Size. • **Redo Logs**: Click this node to view information on the redo logs associated with the Instance. – for Log Group: Sequence Number, Status, Archived, Total Size – for Log File: File Name (full path) • **Archive Log Directories**: Click this node to view information on the archive log directory associated with the Instance, including Directory Name, Size of Files, File Count, Status, Last Failure Time.
SQL Server Detailed Instance Information **By Computer** By Computer > *machine_name* > *server_name* **By DBMS Type** By DBMS Type > SQL/Server > *server_name*	View information about the objects associated with a selected SQL Server Instance.	• **Databases**: When you expand the **Databases** node, click the name of a Database to view information about that database, its related data file, and its related log file. • **Monitored Tables**: When you expand the **Monitored Tables** node, click on the name of a table to view its information including the computer, instances, and database on which is resides, as well as its creator, total size, empty used space, index size, number of indexes, row count, etc. Only the tables that belong to table groups included in scans will be available to view through this node. • **Data Files**: Click the **Data Files** node, click the name of a data file to view information about the data files associated with the Instance, including: Data File Path, Used Space, Percent Used, Free Space, Total Size, Coalesced Extents, Minimum Free Extent Size, Maximum Free Extent Size, and Create Time.

Asset Report Type	Use these reports to...	What's in these reports?
Sybase Detailed Instance Information **By Computer** By Computer > *machine_name* > *instance_name* **By DBMS Type** By DBMS Type > Sybase > *instance_name*	view information about the objects associated with a selected Sybase Instance.	• **Devices**: Click this node to view information about the Device associated with the Instance, including its name, computer on which it resides, path, free space, percent free, used space, total size, fragments, mirror path, etc. • **Databases**: When you expand the **Databases** node, click the name of a Database to view information about that database, its device information, and its related fragment information. • **Monitored Tables**: When you expand the **Monitored Tables** node, click on the name of a table to view its information including the computer, instances, and database on which is resides, as well as its creator, total size, empty used space, index size, number of indexes, row count, etc. Only the tables that belong to table groups included in scans will be available to view through this node. • **Fragments**: Click this node to view about the Fragments associated with the Instance including the computer, instances, and database on which they resides as well as their path, used space, percent used, free space, total size, data type, etc.
UDB Detailed Instance Information **By Computer** By Computer > *machine_name* > *instance_name* **By DBMS Type** >	view information about the objects associated with a selected UDB Instance.	• **Databases**: When you expand the **Databases** node, click the name of a databases to view information about that database. • **Monitored Tables**: When you expand the **Monitored Tables** node, click on the name of a table to view its information. Only the tables that belong to Table Groups included in scans will be available to view through this node. • **Containers**: Click the **Containers** node, click the name of a container to view information about the containers associated with the Instance, including: Database, Tablespace, Container Node, Container Name, Used Space, Percent Used, Free Space, Total Size, Container Type, Status, Discovered Time.

System-wide:

Use System-wide reports to view your RDBMS storage assets as a whole, across the storage environment

Asset Report Type	Use these reports to...	What's in these reports?
System-Wide Asset > System-wide	drill down through your DBMS storage assets in a hierarchical manner.	Depending on the DBMS of the Instance you want to view, you can drill down and expose its: • Tablespaces • Databases • Devices • Monitored tables • Containers • Data files • Fragments • Control files • Redo logs • Archive log directories
All DBMSs Asset > System-wide >All DBMSs	view information about all the DBMSs within your environment based on the objects related to those instances.	• **Instances**: Expand the **Instances** node and click the sorting method by which you want to view instances. • **Databases-Tablespaces**: Expand **Databases-Tablespaces** node and click the sorting method by which you want to view information about the databases and/or tablespaces within your environment. • **Tables**: Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to table groups included in scans will appear within the reports viewed through this node. • **Files**: Expand the **Files** node and click the sorting method by which you want to view information about data files/containers, which includes: File Path, Free Space, Percent Free, Used Space, Total Size, Create Time, and Discovered Time. • **Users**: Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored DBMS within your environment.

Asset Report Type	What's in these reports?
All DBMSs (continued)	• **Instances**: Expand the **Instances** node and click the sorting method by which you want to view instances. – By Instance: sort based on Instance name – By Version: sort based on the version of the DBMS under which an Instance is running – By Start Time: sort based upon when an Instance was last started – By Probe Time: sort based upon when an Instance was last probed – By RDBMS Type: sort based upon the DBMS of an Instance • **Databases-Tablespaces**: Expand the **Databases-Tablespaces** node and click the sorting method by which you want to view information about the databases or tablespaces within your environment. – By Database-Tablespace: sort based on DB/TB name – By Freespace: sort based on the amount of free space on a DB/TB – By Total Size: sort based on the total size of a DB/TB – By Table Count: sort based upon the number of tables on a DB/TB – By Scan Time: sort based upon when a DB/TB was last scanned – By Discovered Time: sort based upon when a DB/TB was discovered – Dropped Tablespaces: sort based upon dropped DBs or TBs. – By RDBMS Type: sort based upon the DBMS of a DB/TB • **Tables**: Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans will appear within the reports viewed through this node. – By Table: sort based on table name – By Number of Indexes: sort based upon the number of indexes associated with a table – By Total Size: sort based on the total size of a table – By Row Count: sort based upon the number of rows in a table – By Chained Row Count: sort based upon the number of chained rows in a table – By Discovered Time: sort based upon when a table was discovered – Dropped Tables: sort based upon dropped tables – By RDBMS Type: sort based upon the DBMS of a table • **Files**: Expand the **Files** node and click the sorting method by which you want to view information about DBMS files. – By File: sort based on file name – By Free Space: sort based upon file free space – By Total Size: sort based on the total size of a file – By Discovered Time: sort based upon when a file was discovered – By RDBMS Type: sort based upon the DBMS of a file • **Users**: Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored DBMS within your environment.

Asset Report Type	What's in these reports?
Oracle Asset > System-wide > Oracle	• **Instances**: Expand the **Instances** node and click the sorting method by which you want to view instances. – By Instance: sort based on Instance name – By Version: sort based on the version of the DBMS under which an Instance is running – By Start Time: sort based upon when an Instance was last started – By Probe Time: sort based upon when an Instance was last probed • **Tablespaces**: Expand this node and click the sorting method by which you want to view information about tablespaces within your environment. – By Tablespace: sort based on tablespace name – By Max Free Extent: sort based on the maximum free extent size – By Freespace: sort based on the amount of free space on a tablespace – By Total Size: sort based on the total size of a tablespace – By Coalesced Extents: sort based on the number of coalesced extents – By Table Count: sort based upon the number of tables – By Scan Time: sort based upon when a tablespace was last scanned – By Discovered Time: sort based upon when a tablespace was discovered – Dropped Tablespaces: sort based upon dropped tablespaces • **Tables**: Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans will appear within the reports viewed through this node. – By Table: sort based on table name – By Number of Indexes: sort based upon the number of indexes associated with a table – By Total Size: sort based on the total size of a table – By Row Count: sort based upon the number of rows in a table – By Chained Row Count: sort based upon the number of chained rows in a table – By Discovered Time: sort based upon when a table was discovered – Dropped Tables: sort based upon dropped tables • **Data Files**: Expand the **Data Files** node and click the sorting method by which you want to view information about DBMS files. – By Data File: sort based on file name – By Max Free Extent: sort data files sorted by their max free extent sizes – By Free Space: sort based upon file free space – By Total Size: sort based on the total size of a file – With Auto-extend: sort data files that have auto-extend specified – Without Auto-extend: sort data files that do not have auto-extend specified – By Discovered Time: sort based upon when a file was discovered • **Users**: Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored Oracle instances within your environment.

Asset Report Type	What's in these reports?
SQL Server Asset > System-wide > SQL/Server	• **Instances**: Expand the **Instances** node and click the sorting method by which you want to view instances. – By Instance: sort based on Instance name – By Version: sort based on the version of the DBMS under which an Instance is running – By Probe Time: sort based upon when an Instance was last probed • **Databases**: Expand this node and click the sorting method by which you want to view information about databases within your environment. – By Database: sort based on database name – By Freespace: sort based on the amount of free space on a database – By Total Size: sort based on the total size of a database – By Table Count: sort based upon the number of tables in the database – By Scan Time: sort based upon when a database was last scanned – By Discovered Time: sort based upon when a database was discovered – Dropped Databases: sort based upon dropped databases • **Tables**: Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans will appear within the reports viewed through this node. – By Table: sort based on table name – By Number of Indexes: sort based upon the number of indexes associated with a table – By Total Size: sort based on the total size of a table – By Row Count: sort based upon the number of rows in a table – By Discovered Time: sort based upon when a table was discovered – Dropped Tables: sort based upon dropped tables • **Data Files**: Expand the **Data Files** node and click the sorting method by which you want to view information about data files. – By Data File: sort based on file name – By Free Space: sort based upon file free space – By Total Size: sort based on the total size of a file – With Auto-extend: sort data files that have auto-extend specified – Without Auto-extend: sort data files that do not have auto-extend specified – By Discovered Time: sort based upon when a file was discovered • **Users**: Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored SQL Server instances. – By Users: sort based on user names – By Discovered Time: sort based when a user was discovered

Asset Report Type	What's in these reports?
Sybase Asset > System-wide > Sybase	• **Instances**: Expand the **Instances** node and click the sorting method by which you want to view instances. – By Instance: sort based on Instance name – By Version: sort based on the version of the DBMS under which an Instance is running – By Probe Time: sort based upon when an Instance was last probed • **Devices**: Expand this node and click the sorting method by which you want to view information about devices within your environment. – By Device: sort based on device name – By Freespace: sort based on the amount of free space on a device – By Total Size: sort based on the total size of a device – By Fragments: sort based upon the number of fragments for a device – By Discovered Time: sort based upon when a device was discovered – Dropped Devices: sort based upon dropped devices • **Databases**: Expand this node and click the sorting method by which you want to view information about databases within your environment. – By Database: sort based on database name – By Freespace: sort based on the amount of free space on a database – By Total Size: sort based on the total size of a database – By Table Count: sort based upon the number of tables in the database – By Scan Time: sort based upon when a database was last scanned – By Discovered Time: sort based upon when a database was discovered – By Last Dump Time: sort based on the date/time when devices were last dumped – Dropped Databases: sort based upon dropped databases • **Fragments**: Expand this node and click the sorting method by which you want to view information about fragments. – By Fragments: sort based on fragment name – By Free Space: sort based upon the free space in fragments – By Total Size: sort based on the total size of fragments – By Discovered Time: sort based upon when fragments were discovered • **Tables**: Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans will appear within the reports viewed through this node. – By Table: sort based on table name – By Number of Indexes: sort based upon the number of indexes associated with a table – By Total Size: sort based on the total size of a table – By Row Count: sort based upon the number of rows in a table – By Chained Row Count: sort based upon the number of chained rows in a table – By Discovered Time: sort based upon when a table was discovered – Dropped Tables: sort based upon dropped tables • **Users**: Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored Sybase instances.

Asset Report Type	What's in these reports?
UDB Asset > System-wide > UDB	• **Instances**: Expand the **Instances** node and click the sorting method by which you want to view instances. – By Instance: sort based on Instance name – By Version: sort based on the version of the DBMS under which an Instance is running – By Probe Time: sort based upon when an Instance was last probed • **Databases**: Expand this node and click the sorting method by which you want to view information about databases within your environment. – By Database: sort based on database name – By Freespace: sort based on the amount of free space on a database – By Total Size: sort based on the total size of a database – By Table Count: sort based upon the number of tables in the database – By Scan Time: sort based upon when a database was last scanned – By Discovered Time: sort based upon when a database was discovered – Dropped Databases: sort based upon dropped databases • **Tablespaces**: Expand this node and click the sorting method by which you want to view information about tablespaces. – By Tablespace: sort based on tablespace name – By FreeSpace: sort based upon the freespace in tablespaces – By Total Size: sort based on the total size of tablespaces – By Table Count: sort based upon the number of tables associated with tablespaces – By Scan Time: sort based upon the date/time when tablespaces were scanned – By Discovered Time: sort based upon the date/time when tablespaces were discovered – Dropped Tablespaces: sort based upon the names of dropped tablespaces • **Tables**: Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans will appear within the reports viewed through this node. – By Table: sort based on table name – By Number of Indexes: sort based upon the number of indexes associated with a table – By Total Size: sort based on the total size of a table – By Row Count: sort based upon the number of rows in a table – By Chained Row Count: sort based upon the number of chained rows in a table – By Discovered Time: sort based upon the date/time when a table was discovered – Dropped Tables: sort based upon the names of dropped tables

Asset Report Type	What's in these reports?
	• **Containers**: Expand the **Containers** node and click the sorting method by which you want to view containers. – By Container: sort based on container name – By Freespace: sort based upon the freespace in containers – By Total Size: sort based on the total size of containers – With Auto–extend: sort based upon whether containers with Auto-extend on – Without Auto–extend: sort based upon whether containers with Auto-extend off – By Discovered Time: sort based upon the date/time when container were discovered • **Users**: Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored UDB instances.

Capacity reports

Use Capacity reports to view the storage capacity, used space, and free space for your RDBMS objects.

Capacity enable you to view storage capacity at the computer, instance, database, device, fragment, container, and data file level. For a specific Instance, computer, group of computers, database, group of databases, tablespace, group of tablespaces, datafiles, containers, fragments, or all the computers in your enterprise, you can view the following:

- How much storage you have for a DBMS and the objects it contains
- How much storage is being used for a DBMS and the objects it contains
- What percentage of DBMS storage is free for expansion

All DBMSs Capacity reports:

Use these reports to view storage capacity information about the different RDBMS instances that you are monitoring within your environment.

Capacity Reporting View	Use these reports to...
Capacity > All DBMSs > Total Instance Storage	If you have multiple RDBMSs that you monitor using Data Manager, use this node to view detailed storage information about the different RDBMS instances within your environment. Depending on the RDBMS of the instance you want to view, use this report to view information about its: • RDBMS instance • RDBMS type • total size • file capacity • file free space • DB-TS count • file count • log file count

Oracle Capacity reports:

Use these reports to view storage capacity information about the Oracle instances that you are monitoring within your environment.

Capacity Reporting View	Use these reports to...
Oracle Capacity > Oracle **By Clauses** for these reports: • By Instance: sort based on Instance name • By Database: sort a report by database name • By Tablespace: sort a report by tablespace name • By Tablespace Group: view the desired storage information organized by tablespace group • By Datafile: sort based on file name • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor Oracle instances. Expand this node to view the following information: • **Total Instance Storage.** Generate a report under this node to view detailed information about Oracle instances, including: – total size – file capacity – file free space – DB-TS count – file count – log file count • **Data File Used Space.** Generate a report under this node to view detailed information about the storage usage of datafiles associated with the tablespaces in an Instance, including: – tablespace – data file path – used space – percent used – free space – total size – free extents – coalesced extents – maximum free extent size • **Data File Free Space**. Generate a report under this node to view detailed information about the storage free space of datafiles associated with the tablespaces in an Instance, including: – tablespace – data free space – percent free – data used space – data size – free extents – coalesced extents – maximum free extent size

SQL Server Capacity reports:

Use these reports to view storage capacity information about the SQL Server instances that you are monitoring within your environment.

Capacity Reporting View	Use these reports to...
SQL Server Capacity > SQL/Server **By Clauses** for these reports: • By Instance: sort based on Instance name • By Database: sort a report by database name • By Database Group: view the desired storage information organized by database group • By Datafile: sort based on file name • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor SQL Server instances. Expand this node to view the following information: • **Total Instance Storage.** Generate a report under this node to view detailed information about SQL Server instances, including: – total size – data file capacity – data file free space – log file capacity – log file free space – database count – data file count – log file count • **Data File Used Space.** Generate a report under this node to view detailed information about the storage usage of datafiles associated with the databases in an Instance, including: – database – data used space – percent used – data free space – data size – log used space – log percent used – log free space – log size – number data files • **Data File Free Space**. Generate a report under this node to view detailed information about the storage free space of datafiles associated with the databases in an Instance, including: – database – data free space – percent free – data used space – data size – log free space – log percent used – log used space – log size – number data files

Sybase Capacity reports:

Use these reports to view storage capacity information about the Sybase instances that you are monitoring within your environment.

Capacity Reporting View	Use these reports to...
Sybase Capacity > Sybase **By Clauses** for these reports: • By Device: sort based on Device name • By Instance: sort based on Instance name • By Fragments: sort based upon fragments • By Database: sort a report by database name • By Database Group: view the desired storage information organized by database group • By Datafile: sort based on file name • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor Sybase instances. Expand this node to view the following information: • **Total Instance Storage.** Generate a report under this node to view detailed information about Sybase instances, including: – total device size – mirrored device size – device free space – data fragment capacity – data fragment free space – log fragment capacity – log fragment free space – device count – database count – fragment count – log fragment count • **Device Used Space.** Generate a report under this node to view detailed information about the storage usage of devices associated with instances, including: – Instance – device path – used space – percent used – free space – total size – number fragments – mirror path – number devices • **Device Free Space**. Generate a report under this node to view detailed information about the storage free space of devices associated instances, including: – Instance – device path – free space – percent free – used space – total size – number fragments – low, high – mirror path – number devices

Capacity Reporting View	Use these reports to...
Sybase (continued)	• **Fragment Used Space.** Generate a report under this node to view detailed information about the storage usage of fragments associated with databases, including: – Instance – fragment path – database – data used space – percent used – data free space – data size – data type – log used space – log percent used – log free space – log size – number fragments • **Fragment Free Space**. Generate a report under this node to view detailed information about the storage free space of fragments associated databases, including: – Instance – fragment path – database – data free space – percent free – data used space – data size – log free space – log used space – log percent used – log used space – log size – number fragments

UDB Capacity reports:

Use these reports to view storage capacity information about the UDB instances that you are monitoring within your environment.

Capacity Reporting View	Use these reports to...
UDB Capacity > UDB **By Clauses** for these reports: • By Container: sort based on Container name • By Tablespace: sort based on Tablespace name • By Database: sort a report by database name • By Database Group: view the desired storage information organized by database group • By Instance: sort based on Instance name • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor UDB instances. Expand this node to view the following information: • **Total Instance Storage.** Generate a report under this node to view detailed information about UDB instances, including: – total size – container capacity – container free space – log file capacity – tablespace count – container count – log file count – log mode • **Container Used Space.** Generate a report under this node to view detailed information about the storage usage of containers associated with a tablespace, including: – Instance – database – tablespace – container node – container name – used space – percent used – free space – total size – container type – status – discovered time • **Container Free Space**. Generate a report under this node to view detailed information about the storage free space of containers associated with a tablespace, including: – Instance – database – tablespace – container node – container name – free space – percent free – used space – total size – container type – status – discovered time

Usage reports

Use Usage reports to monitor and report on the usage and growth of the storage consumption for your databases, tablespaces, tables, segments, and users.

Use these reports to do the following:

- View detailed information for tablespaces and databases including their total size, used space, empty used space, # of tables and indexes, etc.
- Identify the tables consuming the most space within your environment. You can view this information at the tablespace, database, and computer level, as well as at a network-wide level.
- (Oracle) Identify the segments with the most extents and most empty used space. You can view this information at the tablespace, database, and computer level, as well as at a network-wide level.
- Associate users with the amount of storage they are consuming on specific databases, groups of databases, tablespaces, groups of tablespaces, specific computers, groups of computers, and throughout the entire network.

All DBMS Usage reports:

Use these reports to view storage usage information about the different RDBMS instances that you are monitoring within your environment.

Usage Reporting View	**Use these reports to...**
Usage > All DBMSs	If you have multiple RDBMSs that you monitor using Data Manager, use this node to view detailed storage information about the different RDBMS instances within your environment.
Usage > All DBMSs > Databases-Tablespaces	view storage usage statistics about the databases and tablespaces with your monitored DBMSs. Some of the information you can view through these reports includes: • database name, tablespace name • computer • Instance • Instance type • total size • used space • empty used space • table count • index count • other count
Usage > All DBMSs > Tables	view storage usage statistics about the tables with your monitored DBMSs. Some of the information you can view through these reports includes the following: • Largest tables • Total size • Creator • Table name • Empty used space • Index size • Row count • Chained row count

Usage Reporting View	Use these reports to...
Usage > All DBMSs > Users	associate users with the amount of storage they are consuming on specific DBMS objects, specific computers, groups of computers, and throughout the entire network. • **Group by User Name**: select the users and then the resources on which you want to report storage consumption. For example, when setting up a report to view user storage consumption on an Oracle or UDB tablespace, you must do the following: – Select the users upon which you want to report – Select the tablespaces whose storage consumption by the selected users you want to view • **Group by Resource**: select the resources and then the users on which you want to report storage consumption. For example, when setting up a report to view storage consumption on Oracle or UDB tablespaces for users, you must do the following: – Select the tablespaces upon which you want to report – Select the users whose storage consumption on the selected tablespaces you want to view

Oracle Usage reports:

Use these reports to view storage usage information about the Oracle tablespaces, tables, segments, and users that you are monitoring within your environment.

Usage Reporting View	Use these reports to...
Oracle Usage > Oracle **By Clauses** for these reports: • By Database—sort a report by database name • By Tablespace—sort a report by tablespace name • By Tablespace Group—view the desired storage information organized by tablespace group • By Table—sort a report by table name • By Table Group—view the desired storage information organized by table group • By Computer—sort based on the computer where the object's Instance resides • By Computer Group—view the desired storage information organized by computer group • Network-wide—view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor Oracle instances. Expand this node to view the following information: • **Tablespaces.** View detailed information for tablespaces including the following: – Instance – Database – Total size – Used space – Empty used space – Freelist blocks – Table count – Index count • **Tables.** Identify the tables consuming the most space within your environment. You can view this information at the tablespace, database, and computer level, as well as at a network-wide level. Some of the information in these reports includes the following: – Table name – Total size – Empty used space – Index size – Freelist blocks – Number partitions – Number indexes – Number segments – Number extents – Percent free, used

Usage Reporting View	Use these reports to...
Oracle Usage → Oracle **By Clauses** for these reports: • By Database—sort a report by database name • By Tablespace—sort a report by tablespace name • By Tablespace Group—view the desired storage information organized by tablespace group • By Table—sort a report by table name • By Table Group—view the desired storage information organized by table group • By Computer—sort based on the computer where the object's Instance resides • By Computer Group—view the desired storage information organized by computer group • Network-wide—view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor Oracle instances. Expand this node to view the following information: • **Segments**. Identify the segments with the most extents and most empty used space. Some of the information in these reports includes the following: – Number extents – Segment name, creator – Tablespace – Parent name, creator – Total size – Empty used space – Number extents – Freelist blocks, groups • **Users**. Associate users with the amount of storage they are consuming on specific databases, tablespaces, groups of tablespaces, specific computers, groups of computers, throughout the entire network, and applied Quotas. Some of the information included in these reports includes the following: – Computer, instance, database, tablespace – Total size – Empty used space – Table count – Index count – Largest segment, 2nd largest segment – Segment sizes: <50Kb, 500Kb-100Mb, 1Mb-10Mb, 10Mb-100Mb, 100Mb-500MB, 500Mb-1Gb, >1Gb – Quota name – Percent of Quota – Quota threshold

SQL Server Usage reports:

Use these reports to view storage usage information about the SQL Server databases, tables, and users that you are monitoring within your environment.

Usage Reporting View	Use these reports to...
SQL Server Usage > SQL/Server **By Clauses** for these reports: • By Instance: sort based on Instance name • By Database: sort a report by database name • By Database Group: view the desired storage information organized by database group • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network • By User Quota: sort based upon the users to which Quotas are applied. A user who has not violated a quota will not appear in this report.	This node appears if you are using Data Manager for Databases to monitor storage usage on SQL Server instances. Expand this node to view the following information: • **Databases.** Generate a report under this node to view detailed information about the storage on SQL Server databases, including the following: – Total size – Used space – Empty used space – Table count – Index count • **Tables.** Identify the tables consuming the most space within your environment. You can view this information at the table, database, and computer level, as well as at a network-wide level. Some of the information in these reports includes the following: – Table name – Total size – Empty used space – Index count – Row count • **Users**. Associate users with the amount of storage they are consuming on specific databases, groups of databases, specific computers, groups of computers, throughout the entire network, and applied Quotas . Some of the information in these reports includes the following: – Computer, instance, database – Total size – Empty used space – Table count – Index count – Largest object, 2nd largest object – Object sizes: <50Kb, 500Kb-100Mb, 1Mb-10Mb, 10Mb-100Mb, 100Mb-500MB, 500Mb-1Gb, >1Gb – Quota name – Percent of Quota – Quota threshold

Sybase Usage reports:

Use these reports to view storage usage information about the Sybase databases, tables, and users that you are monitoring within your environment.

Usage Reporting View	Use these reports to...
Sybase Usage > Sybase • By Database: sort a report by database name • By Database Group: view the desired storage information organized by database group • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network • By User Quota: sort based upon the users to which Quotas are applied. A user who has not violated a quota will not appear in this report.	This node appears if you are using Data Manager for Databases to monitor storage usage on Sybase instances. Expand this node to view the following information: • **Databases.** Generate a report under this node to view detailed information about the storage on Sybase databases, including the following: – Database name, computer, instance – Total size – Used space – Empty used space – Table count – Index count • **Tables.** Identify the tables consuming the most space within your environment. You can view this information at the table, database, and computer level, as well as at a network-wide level. Some of the information in these reports includes the following: – Table name, database, computer, Instance, creator – Total size – Empty used space – Index size – Maximum rows per page – Expected row size – Fill factor – Page gap – Average row length – Row count • **Users**. Associate users with the amount of storage they are consuming on specific databases, groups of databases, specific computers, groups of computers, throughout the entire network, and applied Quotas. Some of the information included in these reports includes the following: – Computer, instance, database – Total size – Empty used space – Table, index count – Largest object, 2nd largest object – Object sizes: <50Kb, 500Kb-100Mb, 1Mb-10Mb, 10Mb-100Mb, 100Mb-500MB, 500Mb-1Gb, >1Gb – Quota name – Percent of Quota – Quota threshold

UDB reports:

Use these reports to view storage usage information about the UDB tablespaces, tables, and users that you are monitoring within your environment.

Usage Reporting View	Use these reports to...
UDB Usage > UDB • By Tablespace: sort a report by tablespace name • By Tablespace Group: view the desired storage information organized by tablespace group • By Database: sort a report by database name • By Instance: sort a report by Instance name • By Computer: sort based on the computer where the object's Instance resides • By Computer Group: view the desired storage information organized by computer group • Network-wide: view the desired storage information for your entire network	This node appears if you are using Data Manager for Databases to monitor storage usage on UDB instances. Expand this node to view the following information: • **Tablespaces.** Generate a report under this node to view detailed information about the storage on UDB tablespaces, including the following: – Database name, computer, instance – Total size – Used space – Empty used space – Table count – Index count • **Tables.** Identify the tables consuming the most space within your environment. You can view this information at the table, database, and computer level, as well as at a network-wide level. Some of the information in these reports includes the following: – Table name, database, computer, Instance, creator – Total size – Empty used space – Index size – Maximum rows per page – Expected row size – Fill factor – Page gap – Average row length – Row count **Note:** When viewing reports under the **Monitored Tables** node, keep in mind that a table will appear in this report if: 1) it is included in the definition of a table group and 2) the table group in which it is included is defined as part of a scan. • **Users**. Associate users with the amount of storage they are consuming on specific UDB instances, tablespaces, groups of tablespaces, computers, groups of computers, and throughout the entire network. Some of the information included in these reports includes the following: – Computer, instance – Total size – Empty used space – Table, index count – Largest object, 2nd largest object – Object sizes: <50Kb, 500Kb-100Mb, 1Mb-10Mb, 10Mb-100Mb, 100Mb-500MB, 500Mb-1Gb, >1Gb – Quota name – Percent of Quota – Quota threshold

Usage Violation reports

Use Usage Violation reports to view information about quota violations and help enforce corporate storage usage policies.

With quotas, you can control how much storage a user, or a group of users, can consume on the following:

- Database or group of databases
- Tablespace or group of tablespaces
- Instance

- Throughout the entire network

The following table contains information about Usage Violation reports:

Usage Violation Reporting Type	Use these reports to...	What's in these reports?
Quota Violations Usage Violations > Quota Violations	identify the quota violations within the environment	
All Quotas Usage Violations > Quota Violations > All Quotas	view all the quotas and quota violations in the environment • **By User**. View information about the users who have violated quotas. • **By User Quota**. View information about the quotas that have been violated.	• User name • RDBMS type • Instance • Database-tablespace • Space used • Threshold • % of threshold • Number of violations • Quota creator • Quota name • Quota type
Network Wide Quotas Usage Violations > Quota Violations > Network-wide Quotas	view quota violations at the network-wide level • **By User**. View information about the users who have violated quotas. • **By User Quota**. View information about the quotas that have been violated and which users have violated them.	• User name • Space used • Threshold • % of threshold • Number of violations • Quota creator • Quota name • Quota type
Instance Quotas Usage Violations > Quota Violations > Instance Quotas	view quota violations at the Instance level • **By Instance**. View Instance quota violations sorted by Instance • **By Computer**. View Instance quota violations sorted by computer • **By Computer Group**. View Instance quota violations sorted by computer group • **By User**. View Instance quota violations sorted by users • **By User Quota**. View Instance quota violations sorted by quota	• Computer • RDBMS type • Group creator (by group clauses only) • Group name (by group clauses only) • Instance • Database-tablespace • User name • Space used • Threshold • % of threshold • Number of violations • Quota creator • Quota name • Quota type

Usage Violation Reporting Type	Use these reports to...	What's in these reports?
Database-Tablespace Quotas Usage Violations > Quota Violations > Database-Tablespace Quotas	view quota violations at the tablespace level • **By Database-Tablespace**. View quota violations sorted by the database or tablespace name on which the violation occurred • **By Database-Tablespace Group**. View quota violations sorted by the database or tablespace group in which the violation occurred • **By Instance**. View database or tablespace quota violations sorted by Instance • **By Computer**. View database or tablespace quota violations sorted by computer • **By Computer Group**. View database or tablespace quota violations sorted by computer group • **By User**. View database or tablespace quota violations sorted by users • **By User Quota**. View database or tablespace quota violations sorted by quota	• Computer • Group creator (by group clauses only) • Group name (by group clauses only) • Instance • Database-tablespace • User name • Space used • Threshold • % of threshold • Number of violations • Quota creator • Quota name • Quota type

Data Manager for Chargeback reports

Use invoices generated by Data Manager for Chargeback to charge users or groups of users for storage usage within a storage environment.

See "Working with invoices" on page 454 for information about how to use Data Manager for Chargeback to generate invoices.

Disk Manager reports

Use Disk Manager reports to view the relationships among the components of a storage subsystem and the storage capacity for a computer, filesystem, storage subsystem, volume, Array Sites, Ranks, and disk.

Reporting on storage subsystems

Learn how to use Disk Manager to report on storage subsystems.

These reports also enable you to view the relationships among the components of a storage subsystem. Storage-subsystems reports use data gathered by probes.

To ensure the accuracy and consistency of the data that appears in storage-subsystem reports, run regularly scheduled probes against the hosts that use or import volumes, and the storage subsystems upon which you want to report. You should run probes against storage subsystems after running them against any host.

- A disk must be partitioned before it is considered a logical volume by Disk Manager. If a disk is not partitioned, it will not be included in a file system or logical volumes report.

- Before you begin, you might want to refer to Generating and viewing reports. That section provides information on reports that applies all IBM Tivoli Storage Productivity Center reports. It describes the types of reports: system, user-defined, and batch. It also describes how to save reports and how to customize reports, such as filtering information and displaying the information in chart form.

You can view storage subsystem reports in the following ways:

Computer Views

Use these reports to view the relationships between the computer and the storage subsystems, volumes, and disks in your environment. The information is gathered from the computer's viewpoint. Note that, for reporting purposes, *computer* is defined as anything that is running Tivoli Storage Productivity Center.

Note: These reports are available only if the following steps have been performed:

- A data agent is installed on a machine where a volume is surfaced
- A probe has been run against the machine where the volume is located
- A probe has been run against the storage subsystem or IBM SAN Volume Controller cluster

By Computer

These reports view relationships between the computer and the storage subsystems, volumes, or disks in your environment. To generate a report by computer, select the **By Computer** node and choose an item from the **Relate Computers to:** field.

For example, to retrieve information on computers in your environment and their volumes, select the **By Computer** node and then choose **By volumes**. These settings will generate a report detailing all computers with agents on them, which use volumes from the monitored storage subsystems.

By Computer Group

These reports view relationships between the computer groups (both monitoring and reporting groups) and the storage subsystems, volumes, or disks in your environment. To generate a report by computer group, select the **By Computer group** node and choose an item from the **Relate Computers to:** field.

For example, to retrieve information on computers in a group that use storage from monitored storage subsystems, select the **By Computer Group** node and then choose **Storage Subsystems** from the **Relate Computers to:** field. These settings will generate a report detailing all computers in groups with agents on them, which use volumes from the monitored storage subsystems.

By Filesystem/Logical Volume

These reports will view the relationships between the filesystems or logical volumes and the storage subsystems, volumes, and disks in your environment. To generate a report for the relationship you want to view, select the **By Filesystem/Logical Volume** node and choose an item from the **Relate Filesystem/Logical Volume to:** field.

For example, to retrieve information on filesystems or logical volumes that use storage from monitored storage subsystems,

select the **By Filesystem/Logical Volume** node and choose **By Storage Subsystems**. These settings will generate a report detailing all filesystems or logical volumes which use storage from monitored storage subsystems.

By Filesystem Group
Use this report to view the relationships between the filesystem groups (both monitoring and reporting) and the storage subsystems, volumes, and disks in your environment. To generate a report for the relationship you want to view, select the **By Filesystem Group** node and choose an item from the **Relate Filesystem/Logical Volume to:** field.

For example, to retrieve information on filesystems in a group that use storage from monitored storage subsystems, select the **By Filesystem Group** node and choose **Storage Subsystems** from the **Relate Filesystem/Logical Volume to:** field. These settings will generate a report detailing all filesystems in a group which use storage from monitored storage subsystems.

Storage Subsystem Views
Use these reports to view the relationships among the storage-subsystem components (storage subsystems, volumes, and disks) to the computers and file systems or logical volumes in your environment.

Note: These reports are available only if the following steps have been performed:

- A data agent is installed on a machine where a volume is surfaced
- A probe has been run against the machine where the volume is located
- A probe has been run against the storage subsystem or IBM SAN Volume Controller cluster
- For IBM XIV Storage System, no relations between Filesystems/Logical Volumes or Computers and Disks can be displayed.

By Storage Subsystem
These reports view the relationships between the storage subsystems and the computers and file systems or logical volumes in your environment. To generate a report for the relationship you want to view, select the **By Storage Subsystem** node and choose an item from the **Relate Storage Subsystem to:** field.

For example, to view all storage subsystems which have storage being used by computers in the environment, select **By Storage Subsystem** and choose **Computers**. These settings will generate a report detailing all storage subsystems which have storage being used by computers in the environment.

By Storage Subsystem Group
These reports view the relationships between storage subsystem groups (both monitoring and reporting) and the computers and file systems or logical volumes in your environment. To generate a report for the relationship you want to view, select the **By Storage Subsystem Group** node and choose an item from the **Relate Storage Subsystem to:** field.

For example, to view all storage subsystem groups that have storage being used by computers in the environment, select **By Storage Subsystem Group** and choose **Computers** from the **Relate**

Storage Subsystem to: field. These settings will generate a report detailing all storage subsystem groups that have storage being used by computers in the environment.

By Volume

Use this report to view the relationships between the volumes and the computers and file systems or logical volumes in your environment. To generate a report for the relationship you want to view, select the **By Volume** node and then choose an item from the **Relate Volumes to:** field.

For example, to get a listing of all volumes in the environment on which file systems or logical volumes have been created, select **By Volume** and choose **Filesystems/Logical Volumes** from the **Relate Volumes to:** field. These settings will generate a report detailing all volumes in the environment on which file systems or logical volumes have been created.

To get a listing of all volumes in the environment which are being used by computers in the environment, select **By Volume** and choose **Computers**. These settings will generate a report detailing all volumes in the environment which are being used by computers in the environment.

Disks Use this report to view the relationships between disks and the computers and file systems or logical volumes in your environment. To generate a report for the relationship you want to view, select the **Disks** node and then choose an item from the **Relate Disks to:** field.

For example, to get a listing of all storage subsystem disks in the environment on which file systems or logical volumes have been created, select **Disks** and choose **Filesystems/Logical Volumes**. These settings will generate a report detailing all storage subsystem disks in the environment on which file systems or logical volumes have been created.

Volume to HBA Assignment

Use these reports to view all the logical unit numbers (volumes) of each probed storage subsystem that has been assigned to at least one host bus adaptor (HBA) port world wide name (WWN). These reports only retrieve data from machines that have HBAs.

The HBA is a fibre channel and volumes are physically attached to the storage subsystem.

By Storage Subsystem

This report shows all the volumes of each probed storage subsystem that has been assigned to at least one HBA Port WWN. This report is sorted by storage subsystem name.

By Volume Space

This report views all the volumes of each probed storage subsystem that has been assigned to at least one HBA Port WWN. This report is sorted by volume space.

Not on Monitored Computers

This report shows only volumes that have been assigned to HBA Port WWNs on hosts with no Data agent installed on them.

Not Visible to Monitored Server
Use this report to find volumes that have been assigned to hosts with Data agents, but not configured. This report can be used to generate a list of all volumes that need to be configured before they are used.

Volume to Backend Volume Assignment By Volume
Use this report to view the relationships between the virtual disks (including mirror copies) and managed disks that are associated with a virtualizer storage subsystem, such as the SAN Volume Controller, and the back-end storage volumes that are associated with those managed disks.

Note: A virtual disk with one mirror copy appears as one row in the table on this page. A virtual disk with multiple mirror copies appears in multiple rows in the table. Each row represents one of the copies.

Important: Your reports will contain inaccurate information if one of the following configurations are used:

- You used a SAN Volume Controller to virtualize an Tivoli Storage Enterprise Storage Server or FAStT or DS4xxxx subsystem
- You have configured Disk Manager to access the Common Information Model object managers (CIMOMs) of both an SAN Volume Controller and the Disk Array subsystems (Tivoli Storage Enterprise Storage Server and FAStT or DS4xxxx storage subsystems).

In this situation, disks and volumes will appear to belong to both the SAN Volume Controller and the Tivoli Storage Enterprise Storage Server subsystem (or the FAStT or DS4xxxx subsystem), and will be counted twice in system-wide reports. To avoid this situation, CIMOM access should be set up for either the SAN Volume Controller or for the Disk Array subsystems (Tivoli Storage Enterprise Storage Server or FAStT or DS4xxxx), which it virtualizes.

Generating storage-subsystem reports

Learn how to generate a storage subsystem report.

These reports are available only if the following steps have been performed:

1. An agent is installed on a machine where a volume is surfaced.
2. A probe has been run against the machine where the volume is located.
3. A probe has been run against the storage subsystem or SAN Volume Controller cluster.

To generate a storage-subsystems report, do the following:

1. In the left pane, expand **Disk Manager** → **Reporting** → **Storage Subsystems**.
2. Do one of the following:
   a. In the left pane, expand **Computer Views** and click one of the following:
      - By Computer
      - By Computer Group
      - By Filesystem/Logical Volume
      - By Filesystem Group
   b. In the left pane, expand **Storage Subsystems Views** and click one of the following:
      - By Storage Subsystem
      - By Storage Subsystem Group

- By Volume
- Disks

c. In the left pane, expand **Volume to HBA Assignment** and click one of the following:
   - By Storage System
   - By Volume Space
   - Not On Monitored Computers
   - Not Visible to Monitored Computers

d. In the left pane, expand **Volume to Backend Volume Assignment** and click the following:
   - By Volume

3. The **Storage Subsystem Performance** pane for your selected storage subsystem opens the Select Resources pane. All valid available resources are listed in Included Columns. In Available Columns and Included Columns, use  and  to select the columns you want to include in your report.
4. Filters can be applied to the selected columns. In the **Report Filter Specification** area, the **Filter** button opens the **Edit Filter** window.
   a. Click **Add** to add a column to the filter list. The columns will be added in the order they appear on your report.
   b. Click **Delete** to delete all columns. Click the X to the left of an individual column to delete it.
   c. In the **Records Must Meet** area, select **All Conditions** or **At Least One Condition**.
   d. Type the condition value in the **Value 1** text field.
      - If the text field is greyed-out, click **Edit**.
      - If you wish to add additional values for the same column using the same operator, click **Edit**.
   e. Select the Operator.
      - If you select the BETWEEN or NOT BETWEEN operators, two value text fields will open for a lower and upper limit. Click **Edit** to enter the upper and lower limit. You must input values in both **Value 1** and **Value 2** fields, if shown.
      - For numeric columns, such as Free Space, a pull-down menu at the end of the each field assists you with numerical units.
   f. Once your selections have been made, click **OK**.
5. Click **Generate Report**. The report appears under a new tab window. If there is no data, as in if a probe has not been run, the window will say "No Data To Display."
6. Click **File → Save**. The report is saved under **IBM Tivoli Storage Productivity Center → My Reports**. Reports can be regenerated from that node.

To change the parameters for your report, or create a new report of the same type, do the following:

1. Click the **Selection** tab before closing the Storage Subsystem Performance window. Make the necessary changes.
2. Click **Generate Report**.
3. Click **File → Save** and enter the name of the new report.
4. Click **OK**.

Creating a storage subsystem performance report

Learn how to create a storage-subsystem performance report.

Before you can view a performance report on a storage subsystem, that subsystem must have been discovered and had performance data successfully collected. For more information on storage-subsystem discovery, see "Discover storage subsystem, tape library, and fabric information" on page 197.

For a complete list of metrics available for performance reports, see "Performance metrics" on page 815.

1. In the left pane, expand **Disk Manager → Reporting → Storage Subsystem Performance**.
2. Click the type of performance report that you want to view. You can choose from the following performance reports:

   By Storage Subsystem
   This report contains performance information for storage subsystems.

   By Controller
   This report contains performance information for controllers.

   By I/O Group
   This report contains performance information for I/O groups.

   By Node
   This report contains historical performance data for SAN Volume Controller (SVC) nodes.

   By Array
   This report contains performance information for arrays.

   By Managed Disk Group
   This report contains performance information for managed-disk groups.

   By Volume
   This report contains performance information for volumes.

   By Managed Disk
   This report contains performance information for managed disks.

   By Port
   This report contains performance information for ports.

   In the right pane, the Selection page opens.
3. Select from the following options to customize the report:

   Selection button
   Click this button to determine the resources upon which you want to report, such as the storage subsystems, the fibre-channel ports, and so forth. The Select Resources window displays. The type of resources displayed for selection depends on the type of report you selected. For example, if you selected a By Storage Subsystem report, all displayed resources will be storage subsystems. If you clear the box for a resource on the Select Resources window, a pushpin icon displays on the Selection button.

   Filter button
   Click this button to further defines the data that you want to include in the report. Filters enable you to apply general rules to the report. When you click this button, the Edit Filter window displays.

For example, to select only data associated with subsystems that have a name ending with ″123–IBM″, select the ″Subsystem″ column, the ″LIKE″ operator, and type ″*123–IBM″ (without the quotation marks) in the Value 1 field. To select only data which has a Read I/O Rate greater than 0, select the ″Read I/O Rate (overall)″ column, the ″>″ operator, and type ″0″ (without the quotation marks) in the Value 1 field.

If a filter is defined for this report, a pushpin icon displays on the **Filter** button.

Time range selectors

All performance reports display historical performance data. Use the time range selectors to specify the exact time range that you want the report to cover. Click one of the following radio buttons to select the type of historical report you want:

Display latest performance data

This report contains only the latest available performance data for the selected resources. Note that even the latest data available might still be somewhat old, depending on whether data was recently collected for the resources. Including the time column in the report will help you determine how old the latest available data actually is. Regardless of the age of the data, clicking this button ensures that only a single row (representing a single point in time) for each selected resource is included in the report. If no data has ever been collected for a particular resource, then that resource is not included in the report. If no data has ever been collected for any resource, then the resulting report will be empty.

Display historic performance data using absolute time

This report contains only data that falls within a specific time range. Select the time range using the **From** and **To** selectors. The specified from and to values are interpreted as inclusive. The resulting report contains all the performance data that is available within the specified time range, for each selected resource. This means that the data can consist of potentially many rows per resource, each with a unique timestamp. If no data has been collected for a particular resource during the specified time range, that resource is not included in the report. If no data has been collected for any resource during the specified time range, then the resulting report will be empty.

To determine which historical performance records fall within the specified time range, the **beginning timestamp** of the collection intervals is used for comparison. If an historical record contains performance data for an interval of time from 1:00 PM to 1:05 PM, that record would be included in the report for a time range of 1:00 PM to 8:00 PM. However, that same record would not be included if a time range of 1:02 PM to 8:02 PM was specified, because it is the beginning of the time period (for example, 1:00 PM in this example) that is used to determine inclusion in the specified time interval.

Display historic performance data using relative time

This report contains only data that falls within a specific time range extending backward from the current date and time. Specify the number of days in the desired time range, using the **days ago until now** field. The report will be exactly the same

as a report using absolute time, with the **To** selector set to the current date and time, and the **From** selector set to the current date and time minus the number of **days ago until now**.

Note that the same criteria applies as if determining which historical performance records fall within the specified time range. Only the beginning timestamp of the collection intervals is used for comparison. This means that data for the current hour or the current day is not included in the report when selecting hourly or daily summation data, because data for the current hour or current day is only saved in the database when the hour or day is complete. As an alternative, consider using the first radio button, **Display latest performance data**, to see a report showing the most recent daily summation data.

Summation Level

Use this selector to select which type of performance data records to include in the report. The Tivoli Storage Productivity Center Performance Manager collects and stores three different types of performance data:

By sample

Represents the most detailed data. There is usually one sample record per resource for every **n** minutes, where **n** is the performance monitor interval length that was specified when the monitor was defined. Each sample represents the average performance of the resource over those **n** minutes. A sample record with a timestamp of **t** will be saved at time **t+n**. In other words, the beginning of each **n** minute time interval is used for display and for time range matching.

Hourly

Represents an hourly average of the performance of the specified resources. There will be one hourly record written every hour, per resource. Each such record will be written at or shortly after each hour mark on the (device) clock. For example, a record with a timestamp of 12:00 PM and interval length of 1 hour will be written at 1:00 PM.

Daily Represents a daily average of the performance of the resources. There will be one daily record written every day, per resource. Each daily record will be written at or shortly after midnight, according to the device clock. For example, a record with a timestamp of April 21, 12:00 AM and interval length of 24 hours will be written on April 22, 12:00 AM.

Note that hourly and daily records are usually retained in the database for longer periods of time, compared to By Sample records. As a result, if a generated report indicates that sample data is no longer available for a particular time range in the past, it is possible that the sample data has already been purged, in which case you can try generating the report using hourly or daily summation levels.

Be aware that when selecting a time range using the **Display historic performance data using relative time** radio button, it is necessary to select the length of the range in days. This can lead to unexpected results when also selecting the daily summation level. Remember that a daily record is not written until the day is complete, which means that at any given time during the day, the current day's data does not yet

exist in the database. Therefore, selecting the daily summation level as well as a relative time range of n **days ago until now**, will usually result in n-1 records being returned per resource, assuming that at least n-1 days worth of data has been previously collected for those resources. Specifying a relative time range of 1 **days ago until now** will result in an empty report.

Note: The user interface displays the date and time of data in the Tivoli Storage Productivity Center server time zone, but the daily and hourly aggregations are calculated based on the device time zones. Hourly records are written at the end of each hour according to the devices' clocks and the daily records are written at midnight according the the devices' clocks, but both are displayed according to the Tivoli Storage Productivity Center server's clock. For example, if a subsystem clock is one hour ahead of the Tivoli Storage Productivity Center server clock, the hourly data for 5:00 PM according to the device's clock is displayed as 4:00 PM in the user interface and the daily data for 12:00 AM according to the device's clock is displayed as 11:00 PM of the previous day.

Available Columns, Included Columns

Use these items to specify which columns to include and which columns to exclude from the generated report.

Available Columns

Displays the columns that you can include in the generated report that are not already included. If a column appears in the **Included Columns** list box, it will not appear in the **Available Columns** list box.

Included Columns

Displays the columns that will appear in the generated report. By default, all columns for a report are listed in this list box and will be displayed in the generated report. You can highlight the columns and use the up and down arrows to rearrange the order of the columns in the report. The order of the items in this list box determines the order in which the columns will appear in the generated report. The grayed-out column names shown at the top of the list box are the frozen columns, which cannot be selected, removed, or reordered for this type of report.

If a column appears in the **Available Columns** list box, it will not appear in the **Included Columns** list box.

To exclude a column from a report, click the name of the column in the **Included Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the left arrow button (<<). The selected columns are removed from the **Included Columns** list box, and will appear in the **Available Columns** list box.

To include a column in a report, click the name of the column in the Available Columns list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the right arrow button (>>). The selected columns will be removed from the **Available Columns** list box, and will appear in the **Included Columns** list box.

To rearrange the order of columns that will appear in the generated report, click the name of a column in the **Available Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the up arrow button (^) or the down arrow button (v) to move the selected columns up or down in the list, relative to the other columns.

4. Click **Generate Report**. A new tab opens and displays the report.
5. Optional: To view the report in chart format, click  , specify charting options, and then click **OK**. Right-click on the chart and select **Customize this Chart** to further customize how the chart is displayed.
6. Optional: To save the report settings to generate later, click  and enter a name for the report. The report is saved under theIBM Tivoli Storage Productivity Center node. See "Filtering the resources that appear on a report" on page 491 for instructions on modifying saved report settings.

If the report is greater than 2500 rows, a dialog box will appear which will allow you to retrieve additional rows by clicking **Display More Rows**. An additional 2500 rows will be retrieved from the server and displayed on the panel for each click of the button. Do not try to display too much data or there is the possibility of running out of memory.

Generating constraint violation reports

Learn how to generate a storage subsystem performance report.

Creating constraint violation reports:

Learn how to how to create constraint violation reports.

1. In the left pane, expand **Disk Manager** → **Reporting** → **Storage Subsystem Performance**.
2. Click **Constraint Violations**. The constraint violation report selection page appears in the right pane.
3. If you want to limit the report to a particular time period, select **Limit to this time range** and specify the time period.
4. Click **Selection** to choose which storage subsystems to include in the report.
5. Click **Filter** to modify the filter used in generating the report.
6. Customize the performance report. By default, all supported metrics are included. Click to move the selected metric up in the list and to move the selected metric down in the list.
   a. In the **Included Columns** list, select a metric and click  to remove it from the report.
   b. To change the order in which the columns appear in the report, select a metric and click either  or  .
7. Click **Generate Report**. The **Constraint Violations** tab opens and displays the report.

The report might show information about DS6000, DS8000, Tivoli Storage Enterprise Storage Server, SAN Volume Controller, and other subsystems all on the same report. By default, the report displays one subsystem per row. Results are aggregated for the current day.

Viewing detailed constraint violation information:

Learn how to view detailed information about constraint violations.

Before you can view detailed information about constraint violations, you must generate a constraint violation report (see "Creating constraint violation reports" on page 652).

1. Click on the **Constraint Violations** tab to view the report. The report might show information about DS6000, DS8000, Tivoli Storage Enterprise Storage Server,SAN Volume Controller, and other subsystems all on the same report. By default, the report displays one subsystem per row. Results are aggregated for the current day.
2. Click  to view a detailed report of that subsystem.
3. Click  , specify charting options, and then click **OK** to view the report in chart form. Right-click on the chart and select **Customize this Chart** to further customize how the chart is displayed.

Viewing affected volumes and affected hosts:

Learn how to view information on volumes and related hosts affected by constraint violations.

Before you can view affected volumes and hosts information, you must generate a constraint violation report (see "Creating constraint violation reports" on page 652) and view detailed information on a subsystem (see "Viewing detailed constraint violation information").

1. Click  beside a constraint violation to display the Affected Volumes Selection page.
2. Under Volumes, select whether you want the report to show all volumes or only the most active volumes associated with the subsystem and component and type in the maximum number of volumes to display in the report.
3. Under Performance Data, select whether you do not want the report to display performance data for the volumes or you do want the report to display historic performance data for volumes.
4. Customize the performance report. By default, all supported metrics are included. Click to move the selected metric up in the list and to move the selected metric down in the list.
   a. In the **Included Columns** list, select a metric and click  to remove it from the report.
   b. To change the order in which the columns appear in the report, select a metric and click either  or  .
5. Click **Generate Report**.

Fabric Manager reports

Fabric reports provide you with both overview and detailed information about your fabrics, switches, and ports. The reports are based on information collected by

probes and performance monitors. You can view reports as tables of data or as graphical charts. You can use pre-defined Fabric system reports or define your own switch performance reports for information by port or by constraint violations.

Before you can generate a report you need to gather information about your fabrics. You can do this through monitoring jobs, and discoveries.

Creating switch performance reports

After performance monitors are in operation, you can generate switch performance reports. The procedure is similar for each report.

"Creating a switch performance monitor" on page 241 describes how to create performance monitors for fabric switches.

For both reports, use the **Selection** tab to specify switches to report, report columns to include, and filters to apply. When you click **Generate Report**, the **Ports** page or **Constraint Violations** page is created and displays the report.

1. Expand **Fabric Manager → Reporting → Switch Performance**.
2. Click **By Switch**, **By Port** or **By Constraint Violations**. In the right pane, the Selection page opens.
3. Select from the following options to customize the report:

   Selection button
   Click this button to determine the resources upon which you want to report. The Select Resources window displays. The type of resources displayed for selection depends on the type of report you selected. If you clear the box for a resource on the Select Resources window, a pushpin icon displays on the Selection button.

   Filter button
   Click this button to further defines the data that you want to include in the report. Filters enable you to apply general rules to the report. When you click this button, the Edit Filter window displays.

   For example, to select only data associated with switches that have a name ending with ″123–IBM″, select the ″Switch″ column, the ″LIKE″ operator, and type ″*123–IBM″ (without the quotation marks) in the **Value 1** field.

   If a filter is defined for this report, a pushpin icon displays on the **Filter** button.

   Time range selectors
   All performance reports display historical performance data. Use the time range selectors to specify the exact time range that you want the report to cover. Click one of the following radio buttons to select the type of historical report you want:

   Display latest performance data
   This report contains only the latest available performance data for the selected resources. Note that even the latest data available might still be somewhat old, depending on whether data was recently collected for the resources. Including the time column in the report will help you determine how old the latest available data actually is. Regardless of the age of the data, clicking this button ensures that only a single row (representing a single point in time) for each selected resource is included in the report. If no data has ever been collected for

a particular resource, then that resource is not included in the report. If no data has ever been collected for any resource, then the resulting report will be empty.

Display historic performance data using absolute time
This report contains only data that falls within a specific time range. Select the time range using the **From** and **To** selectors. The specified from and to values are interpreted as inclusive. The resulting report contains all the performance data that is available within the specified time range, for each selected resource. This means that the data can consist of potentially many rows per resource, each with a unique timestamp. If no data has been collected for a particular resource during the specified time range, that resource is not included in the report. If no data has been collected for any resource during the specified time range, then the resulting report will be empty.

To determine which historical performance records fall within the specified time range, the **beginning timestamp** of the collection intervals is used for comparison. If an historical record contains performance data for an interval of time from 1:00 PM to 1:05 PM, that record would be included in the report for a time range of 1:00 PM to 8:00 PM. However, that same record would not be included if a time range of 1:02 PM to 8:02 PM was specified, because it is the beginning of the time period (for example, 1:00 PM in this example) that is used to determine inclusion in the specified time interval.

Display historic performance data using relative time
This report contains only data that falls within a specific time range extending backward from the current date and time. Specify the number of days in the desired time range, using the **days ago until now** field. The report will be exactly the same as a report using absolute time, with the **To** selector set to the current date and time, and the **From** selector set to the current date and time minus the number of **days ago until now**.

Note that the same criteria applies as if determining which historical performance records fall within the specified time range. Only the beginning timestamp of the collection intervals is used for comparison. This means that data for the current hour or the current day is not included in the report when selecting hourly or daily summation data, because data for the current hour or current day is only saved in the database when the hour or day is complete. As an alternative, consider using the first radio button, **Display latest performance data**, to see a report showing the most recent daily summation data.

Summation Level
Use this selector to select which type of performance data records to include in the report. The Tivoli Storage Productivity Center Performance Manager collects and stores three different types of performance data:

By sample
Represents the most detailed data. There is usually one sample record per resource for every **n** minutes, where **n** is the performance monitor interval length that was specified when the monitor was defined. Each sample represents the average

performance of the resource over those **n** minutes. A sample record with a timestamp of **t** will be saved at time **t+n**. In other words, the beginning of each **n** minute time interval is used for display and for time range matching.

Hourly
Represents an hourly average of the performance of the specified resources. There will be one hourly record written every hour, per resource. Each such record will be written at or shortly after each hour mark on the (device) clock. For example, a record with a timestamp of 12:00 PM and interval length of 1 hour will be written at 1:00 PM.

Daily Represents a daily average of the performance of the resources. There will be one daily record written every day, per resource. Each daily record will be written at or shortly after midnight, according to the device clock. For example, a record with a timestamp of April 21, 12:00 AM and interval length of 24 hours will be written on April 22, 12:00 AM.

Note that hourly and daily records are usually retained in the database for longer periods of time, compared to By Sample records. As a result, if a generated report indicates that sample data is no longer available for a particular time range in the past, it is possible that the sample data has already been purged, in which case you can try generating the report using hourly or daily summation levels.

Be aware that when selecting a time range using the **Display historic performance data using relative time** radio button, it is necessary to select the length of the range in days. This can lead to unexpected results when also selecting the daily summation level. Remember that a daily record is not written until the day is complete, which means that at any given time during the day, the current day's data does not yet exist in the database. Therefore, selecting the daily summation level as well as a relative time range of n **days ago until now**, will usually result in n-1 records being returned per resource, assuming that at least n-1 days worth of data has been previously collected for those resources. Specifying a relative time range of 1 **days ago until now** will result in an empty report.

Note: The user interface displays the date and time of data in the Tivoli Storage Productivity Center server time zone, but the daily and hourly aggregations are calculated based on the device time zones. Hourly records are written at the end of each hour according to the devices' clocks and the daily records are written at midnight according the the devices' clocks, but both are displayed according to the Tivoli Storage Productivity Center server's clock. For example, if a subsystem clock is one hour ahead of the Tivoli Storage Productivity Center server clock, the hourly data for 5:00 PM according to the device's clock is displayed as 4:00 PM in the user interface and the daily data for 12:00 AM according to the device's clock is displayed as 11:00 PM of the previous day.

Available Columns, Included Columns
Use these items to specify which columns to include and which columns to exclude from the generated report.

Available Columns
Displays the columns that you can include in the generated

report that are not already included. If a column appears in the **Included Columns** list box, it will not appear in the **Available Columns** list box.

Included Columns
Displays the columns that will appear in the generated report. By default, all columns for a report are listed in this list box and will be displayed in the generated report. You can highlight the columns and use the up and down arrows to rearrange the order of the columns in the report. The order of the items in this list box determines the order in which the columns will appear in the generated report. The grayed-out column names shown at the top of the list box are the frozen columns, which cannot be selected, removed, or reordered for this type of report.

If a column appears in the **Available Columns** list box, it will not appear in the **Included Columns** list box.

To exclude a column from a report, click the name of the column in the **Included Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the left arrow button (<<). The selected columns are removed from the **Included Columns** list box, and will appear in the **Available Columns** list box.

To include a column in a report, click the name of the column in the Available Columns list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the right arrow button (>>). The selected columns will be removed from the **Available Columns** list box, and will appear in the **Included Columns** list box.

To rearrange the order of columns that will appear in the generated report, click the name of a column in the **Available Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the up arrow button (^) or the down arrow button (v) to move the selected columns up or down in the list, relative to the other columns.

4. Click **Generate** to run the report.
   The report is displayed in the **Ports** tab. You can display the information in chart form by clicking  .
5. Click **File → Save as** to save the report in the **<user–id> Reports** node.

Switch Performance:

Use the Switch Performance report to generate and view reports that provide information about the free and occupied ports on the switches in your SAN. These reports provide information about the port connections and status, as well as the port and switch IDs, data and packet transfer rates, and other details. This report can only be generated for switches that have had performance monitors run on them.

This following table describes the columns on the Switch Performance report. All the columns in the report are for display only.

Switch	The ID of the switch.

Time	The date and time of the data collection.
Interval	The size of the sample interval, in seconds.
Port Send Packet Rate	The rate at which packets are sent through this port.
Port Receive Packet Rate	The rate at which packets are received through this port.
Total Port Packet Rate	The total rate at which packet operations are performed through this port.
Port Send Data Rate	The rate at which data is sent through this port.
Port Receive Data Rate	The rate at which data is received through this port.
Total Port Data Rate	The total rate at which data is processed through this port.
Port Peak Send Data Rate	The peak rate at which data is sent through this switch.
Port Peak Receive Data Rate	The peak rate at which data is received through this switch.
Port Send Packet Size	The size of the data block that is sent out through this switch.
Port Receive Packet Size	The size of the data block that is received in through this switch.
Overall Port Packet Size	The total size of the data block that is transferred by the switch.
Error Frame Rate	The rate at which error frames are generated by the switch.
Dumped Frame Rate	The rate at which frames are dumped by the switch.
Link Failure Rate	The rate at which failed links are generated by the switch.

Fabric Performance Report - By Port:

Click  to display detailed information per computer in a **Network** tab. The **Network** tab displays a file summary by computer.

The following fields on the Fabric Performance Report - By Port are for display only.

Port	The ID of the port.
WWPN	The World Wide Port Name (WWPN) of the selected port.
Slot	The number for the blade that the port is on.
Index	The unique number for the port across the whole switch.
Time	The date and time of the data collection.
Interval	The size of the sample interval, in seconds.
Port Send Packet Rate	The average number of packets sent per second, by a particular port over a particular time interval. A send operation is a read operation processed, or a write operation initiated by the particular port.
Port Receive Packet Rate	The average number of packets received per second, by a particular port over a particular time interval. A receive operation is a write operation processed, or a read operation initiated by the particular port
Total Port Packet Rate	The average number of packets sent and received per second, by a particular port over a particular time interval.
Port Send Data Rate	The average number of megabytes sent per second, by a particular port over a particular time interval.
Port Receive Data Rate	The average number of megabytes received per second, by a particular port over a particular time interval.
Total Port Data Rate	The average number of megabytes transferred per second, by a particular port over a particular time interval.
Port Peak Send Data Rate	The peak number of megabytes sent per second, by a particular port over a particular time interval.
Port Peak Receive Data Rate	The peak number of megabytes received per second, by a particular port over a particular time interval.
Port Send Packet Size	The size of the data block that is sent out through this switch.
Port Receive Packet Size	The size of the data block that is received in through this switch.
Overall Port Packet Size	The total size of the data block that is transferred by the switch

Error Frame Rate	The rate at which error frames are generated by the switch.
Dumped Frame Rate	The rate at which frames are dumped by the switch.
Link Failure Rate	The rate at which failed links are generated by the switch.
Loss of Sync Rate	The average number of times per second that synchronization was lost after the last reset of the device, for a particular component over a particular time interval.
Loss of Signal Rate	The average number of times per second that the signal was lost after the last reset of the device, for a particular component over a particular time interval.
CRC Error Rate	The average number of frames received per second in which the CRC in the frame did not match the CRC computed by the receiver, for a particular component over a particular time interval.
Short Frame Rate	The average number of frames received per second that were shorter than 28 octets (24 header + 4 CRC), not including any start of frame/end of frame bytes, for a particular component over a particular time interval.
Long Frame Rate	The average number of frames received per second that were longer than 2140 octets (24 header + 4 CRC + 2112 data), not including any start of frame/end of frame bytes, for a particular component over a particular time interval.
Encoding Disparity Error Rate	The average number of disparity errors received per second, for a particular component over a particular time interval.
Discarded Class3 Frame Rate	The average number of class-3 frames per second that were discarded by a particular component over a particular time interval.
F-BSY Frame Rate	The average number of F-BSY frames per second that were generated by a particular component over a particular time interval.
F-RJT Frame Rate	The average number of F-RJT frames per second that were generated by a particular component over a particular time interval.

Constraint Violations:

If you have defined switch performance alerts you can also generate the Constraint Violations report.

Click  to display detailed information per computer in a **Network** tab. The **Network** tab displays a file summary by computer.

The following fields on the Top Switch Ports Packet Rate Performance report are for display only (not editable).

Switch	The ID of the switch.
Total Port Data Rate Threshold	The total rate at which data is processed through this port.
Link Failure Rate Threshold	The rate at which failed links are generated by the switch.
Error Frame Rate Threshold	The rate at which error frames are generated by the switch.
Total Port Packet Rate Threshold	The total rate at which the packet operations are performed through this port.

Tape Manager reports

You can perform the following tape library reporting tasks:

- Create and view asset reports using data collected from your tape libraries; these reports can include such information as the library identifier and the status of the library
- Create and view capacity reports using data collected from your tape libraries; these reports can include such information as the number of drives in a library, the number of tape slots, the number of occupied tape slots

System reports

Use the information in this section to learn how to generate and view system reports. These reports are available in the **IBM Tivoli Storage Productivity Center → Reporting → System Reports** and **Data Manager for Databases → My Reports → System Reports** nodes in the navigation tree. Data for system reports are gathered every time that monitoring jobs are run against your storage resources.

Generating and viewing system reports

Learn how to generate system reports, drill down through and filter their data, and generate charts.

To display a report, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center → My Reports → System Reports → <manager>**.
2. Click the name of the report you want to view. The report will appear in the content pane to the right. To regenerate a report, select the report name, right-click, and select **Regenerate Report**.

Many of the rows that are presented in the generated system reports enable you to:

- drill down into a row to get more detailed information
- view report information in a chart format

To drill down for more detailed report information:

- Click  on the desired row. The drill down report is displayed and a new tab representing that report is added to the tabs at the top of the report.

To view report information in a chart format:

- Click  on the desired row. The chart is displayed and a new tab representing that chart is added to the tabs at the top of the report.

  You can customize the charts that you generate from system reports by using the Chart Customization window. You can access this window by right-clicking on a system report chart and selecting **Customize this chart** from the pop-up menu. See “Customizing charts” on page 479 for information on how to customize a system report chart.

To filter the data displayed on a report:

Click the **Filter** button, which displays an Edit Filter window. In this window, specify criteria for records (for example, Capacity in Data Reports), operators (for

example, =) and values for each ″column″ (whatever you want to type). For example, in the Disk Space Summary report, you can specify only disks with a capacity greater than 70 GB.

If you change any element of a report definition (for example, the maximum number of rows in the report), you can save the new report definition (and regenerate it later) by clicking **File → Save** and then typing a name for the report. The report is saved with the name you specify under the ***user_id's Reports*** node.

System reports - Data

Use these predefined reports to view statistics about your storage resources.

Access Time Summary

This report provides a summary of the number of files in your environment and when they were last accessed (for example, created or modified) during the last day, the last week, the last month, the last two months, the last three months, the last six months, the last nine months, the last year, and over a year.

The following table describes the columns on the Access Time Summary report. All the fields on the report are for display only.

Last Accessed <= 1 day	The number and total size of the files accessed within the last day.
Last Accessed 1 day — 1 week	The number and total size of the files accessed between 1 day and 1 week ago.
Last Accessed 1 week — 1 month	The number and total size of the files accessed between 1 week and 1 month ago.
Last Accessed 1 month — 2 months	The number and total size the files accessed between 1 month and 2 months ago.
Last Accessed 2 months — 3 months	The number and total size of the files accessed between 2 – 3 months ago.
Last Accessed 3 months — 6 months	The number and total size of the files accessed between 3 – 6 months ago.
Last Accessed 6 months — 9 months	The number and total size of the files accessed between 6 – 9 months ago.
Last Accessed 9 months — 1 year	The number and total size of the files accessed between 9 months and 1 year ago.
Last Accessed > 1 year	The number and total size of the files accessed over a year ago.
Overall File Count	Total number of files across a network.
Overall File Size	Total size of the space consumed by the files on a network.
Average Age	Average age of files on a network measured by days, hours, minutes, seconds.

- To view any of the following charts click  at the top of the report and select the corresponding option from the pop-up menu:
  - pie chart of storage space distribution for a selected row
  - pie chart of file counts for a selected row
  - pie chart of storage space distribution for all the rows on the report
  - pie chart of file counts for all the rows on a report
  - history chart of file counts
  - history chart of file count by percent
  - history chart of file sizes
  - history chart of files size by percent

- To view a pie chart of file access times, click  to left of a report row.
- To view detailed file access times for the computers on the network, click  to the left of a report row.

Disk Space Summary

This report shows storage capacity for the whole network. You can drill down through this report to view storage capacity for each of the following levels: per disk, per computer, per cluster, per computer group, and per domain.

The following table describes the columns on the Disk Space Summary report. All the fields on the report are for display only.

Disk Space	Total storage capacity of the disks on the computers within a network
Consumed File System Space	Amount of used storage space on the file systems within a network
Available File System Space	amount of unused storage space on the file systems within a network
Raw Volume Space	Space on host-side logical volumes that is not occupied by file systems
Unavailable Disk Space	Amount of storage subsystem volume or virtual storage volume space that is dedicated to redundancy. This value is dependent on the storage subsystem SMI-S provider returning a valid RAID value that is used to determine the overhead. If the value is zero for a storage subsystem, the overhead cannot be calculated. This amount does not include storage space information from storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
Available Disk Space	Space assigned to a (monitored) host that is not part of any logical volume.
Unknown Storage Subsystem Volume Space	Volume space of unknown usage.

- To view any of the following charts click  at the top of the report and select the corresponding option from the pop-up menu:
  - pie chart of network space distribution
  - pie chart of network row
  - history chart of the free space for a network
  - history chart of the percentage of free space for a network
  - history chart of the storage capacity on a network
- To view a pie chart of disk capacity, click  to the left of a report row.
- To view detailed disk capacity for the computers on a network, click  to the left of a report row.

Access File Summary

This report provides overview information for files in your storage environment.

The Access File Summary report provides overview information for files by directory, directory group, file system, file system group, cluster, computer, computer group, domain, and for the entire network. Through this report you can

view historically the number of files for each resource in the report. The historical chart can be generated to show daily, weekly, or monthly history.

The following table describes the columns on the Access File Summary report. All the fields on the report are for display only.

File Size	Total size of the storage space consumed by the files on a network
File Count	Total number of files on a network
Directory Count	Total number of directories on a network
Average File Size	Average storage space consumed by each of the files on a network
File System Size	Total storage capacity of the files on a network

- To view any of the following charts click  at the top of the report and select the corresponding option from the pop-up menu:
  - History chart of storage space usage for a selected row
  - History chart of file counts for a selected row
- To view a history chart of the storage space consumed by files on a network, click  to the left of a report row.
- To view detailed file storage usage for the computers on the network, click  to the left of a report row.

Storage Access Times

This report shows when files where last accessed and how long ago they were accessed.

The following table describes the columns on the Storage Access Times report. All the fields on the report are for display only.

Computer	Name of a computer against which the report was run
Last Accessed <= 1 day	Number of files that were accessed within the last 24 hours and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Accessed 1 day - 1 week	Number of files that were accessed between 1 day to 1 week previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Accessed 1 week - 1 month	Number of files that were accessed between 1 week to 1 month previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Accessed 1 month - 2 months	Number of files that were accessed between 1 month to 2 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Accessed 2 months - 3 months	Number of files that were accessed between 2 months to 3 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Accessed 3 months - 6 months	Number of files that were accessed between 3 months to 6 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)

Last Accessed 6 months - 9 months	Number of files that were accessed between 6 months to 9 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Accessed 9 months - 1 year	Number of files that were accessed between 9 months to 1 year previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Accessed > 1 year	Number of files that were accessed over one year previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Overall File Count	Total of all the counts
Overall File Size	Total of all the sizes
Average Age	Average time since the file was last accessed

To view graphical information about specific computers' access times:

- Click  . A new tab displaying a chart for each selected computer is displayed.

Disk Defects

This report shows any disk defects on the computers being monitored.

The following table describes the columns on the Disk Defects report. All the fields on the report are for display only.

Device	Name of the computer containing a disk defect
Path	Full path to the location of a disk defect
Primary Defects	Number of defects on the disk when it was new
Manufacturer	Name of the manufacturer for the disk drive containing the disk defect
Model	Model number of the disk drive containing a disk defect
In Storage Subsystem	Indicates Yes when a disk is an actual physical disk within the storage subsystem.
RPM	Speed that the disk turns (revolutions per minute)
Read Cache	Hardware caching turned on/off for reads
Write Cache	Hardware caching turned on/off for writes
Failure Predicted	Disk predicts a failure
Probe Last Run	Date and time when the last probe was run
Discovered Time	Date and time when a disk defect was discovered
Disk Space	Storage capacity of the disk drive containing a defect
Available Disk Space	Amount of unused storage space on the disk drive containing a defect
Grown Defects	Number of defects detected by the disk since new
Recovered Errors	Number of errors (read, write, verify) from which the disk was able to recover

Unrecovered Errors	Number of errors (read, write, verify) from which the disk was not able to recover
Bytes Written	Number of bytes written to the disk
Bytes Read	Number of bytes read from the disk
Automatic Write Reallocation	Disk attempts to move the data sector on the disk in the event of a write error
Desired Automatic Write Reallocation	Change automatic write allocation on the next probe
Automatic Read Reallocation	Disk attempts to move the data sector on the disk in the event of a read error
Desired Automatic Read Reallocation	Change automatic read allocation on the next probe
Read Continuous (no error correction)	No error correction (setting is ON=no error correction)
Desired Read Continuous (no error correction)	Change automatic read continuous on the next probe
Detectable	Indicates Yes when a storage volume is identified as being available to the storage subsystem following a successful probe.
Is Virtual	Indicates Yes when a disk is identified as being available to the server or the storage subsystem.
Mapped to Storage Subsystem	Indicates Yes when a computer disk or a backend storage volume can be mapped to a storage subsystem.

To view more information for a specific computer:

- Click  . Detailed information regarding the computer will be available in three tabs:
  - **General tab.** This tab displays general information about the computer, such as computer information, discovered date and time, disk allocation information, and so on.
  - **Latest Probe tab.** This tab displays information about the latest probe, and provides information such as sector size, probe last run, primary and grown defects, and so on.
  - **Probe History tab.** This tab displays probe history for the selected machine, and provides information such as probe times, record type, bytes written, and so on.

Storage Modification Times

This report shows information about files within the network that were modified with certain time frames.

The Storage Modification Times report provides information about files within the network that were modified:

- within the last 24 hours
- between 24 hours and one week previous
- between one week to one month previous
- between one month to two months previous
- between two months to three months previous
- between three months to six months previous

- between six months to nine months previous
- between nine months to one year previous
- more than one year previous

The following table describes the columns on the Storage Modification Times report. All the fields on the report are for display only.

Last Modified <= 1 day	Number of files that were modified in the last 24 hours and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Modified 1 day - 1 week	Number of files that were modified between 1 day to 1 week previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Modified 1 week - 1 month	Number of files that were modified between 1 week to 1 month previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Modified 1 month - 2 months	Number of files that were modified between 1 month to 2 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Modified 2 months - 3 months	Number of files that were modified between 2 months to 3 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Modified 3 months - 6 months	Number of files that were modified between 3 months to 6 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Modified 6 months - 9 months	Number of files that were modified between 6 months to 9 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Modified 9 months - 1 year	Number of files that were modified between 9 months to 1 year previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Last Modified > 1 year	Number of files that were modified over a year previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)
Overall File Count	Sum of the files
Overall File Size	Sum of the size of the files
Average Age	Average age since files were modified

To view graphical information about storage modification times:

- Click  . A new tab displaying a chart for the network is displayed.

Most At Risk Files

This report provides information on files that have been modified the longest time ago and have not yet been backed-up or archived since they were modified.

The following table describes the columns on the Most At Risk Files report. All the fields on the report are for display only.

Modification Time	Date and time when an at risk file was last modified
Computer	Name of the computer where an at risk file is stored
Filesystem	Filesystem where an at risk file is stored

Path	Full path to the location of an at risk file
Filename	The name of the at risk file
Owner	ID of the user who owns an at risk file
OS User Group	OS User group of the user who owns an at risk file
Logical Size	Logical size of an at risk file
Physical Size	Physical size of an at risk file (measured in kilobytes, megabytes, or gigabytes)
Access Time	Date and time when an at risk file was last accessed
Create Time	Date and time when an at risk file was created

To view more information about a specific file:

- Click  . Detailed information regarding the file will be displayed in a **File Detail** tab. The **File Detail** tab displays more information for the specific file and provides information such as file ownership, modification and creation time, and any file attributes that have been set.

Most Obsolete Files

This report provides information on files that have not been accessed or modified for the longest period of time.

The following table describes the columns on the Most Obsolete Files report. All the fields on the report are for display only.

Access Time	Date and time when an obsolete file was last accessed
Computer	Name of the computer where an obsolete file is stored
Filesystem	Filesystem where an obsolete file is stored
Path	Full path to the location of an obsolete file
Filename	File name of an obsolete file
Owner	ID of the user who owns an obsolete file
OS User Group	OS User Group to which the owner of the obsolete file belongs
Logical Size	Logical size of an obsolete file
Physical Size	Physical size of an obsolete file (measured in kilobytes, megabytes, or gigabytes)
Modification Time	Date and time when an obsolete file was last modified
Create Time	Date and time when an obsolete file was created

To view more information for a specific file:

- Click  . Detailed information regarding the file will be displayed in a **File Detail** tab. The **File Detail** tab displays more information for the specific file and provides information such as file ownership, modification and creation time, and any file attributes that have been set.

Oldest Orphaned Files

This report provides information on files that have the oldest creation date and no longer have the owners registered as users on the computer or network.

The following table describes the columns on the Oldest Orphaned Files report. All the fields on the report are for display only.

Access Time	Date and time when an orphaned file was last accessed
Computer	Name of the computer where an orphaned file is stored
Filesystem	Filesystem where an orphaned file is stored
Path	Full path to the location of an orphaned file
Filename	File name of an orphaned file
Owner	Operating system internal ID of the user who owned the orphaned file. This is the internal ID the operating system uses to identify the user, and *not* the user ID.
OS User Group	OS User group to which the owner of an orphaned file belongs
Logical Size	Logical size of an orphaned file
Physical Size	Physical size of an orphaned file (measured in kilobytes, megabytes, or gigabytes)
Modification Time	Date and time when an orphaned file was last modified
Create Time	Data and time when an orphaned file was created

To view more information for a specific file:

- Click  . Detailed information regarding the file will be displayed in a **File Detail** tab. The **File Detail** tab displays information about a specific file such as file ownership, modification and creation time, and any file attributes that have been set.

User Quota Violations

This report shows which users have violated a Data Manager storage usage quota.

The following table describes the columns on the User Quota Violations report. All the fields on the report are for display only.

Quota Creator	Creator of a quota
Quota Name	Name of a quota
Computer	Computer that violated a quota
Filesystem	File system that violated a quota
User Name	ID of the user that violated a quota
Space Used	Total amount of space used
Threshold	Threshold defined by a quota
% of Threshold	Percentage of quota; always > 100
Number of Violations	Number of violations by a user
Quota Type	Type of quota: Network Wide, Filesystem, or Computer

Computer Storage Availability

This report provides information about the availability of computers within your environment.

If the computer is not available (determined by a ping), the storage is likely not available for use. The following table describes the fields on the Storage Availability report. All the fields on the report are for display only.

Computer	Name of the computer
% Available	Percentage of successful pings
Transmitted	Number of pings transmitted
Received	Number of responses from the ping
Damaged	Number of responses with errors
Avg Transit Time	Average time it took to send and receive response from a ping
Max Transit Time	Maximum time it took to send and receive response from a ping
Min Transit Time	Minimum time it took to send and receive response from a ping
Report Start Time	Start time of information in this row
Report End Time	End time of information

To view the ping history for a specific computer:

- Click  . A new tab displaying ping history information for the selected computer is displayed.

Computer Disk Space

This report provides storage capacity information about each computer within your environment.

The following table describes the columns on the Computer Disk Space report. All the fields on the report are for display only.

Computer	Name of a computer against which the report was run
Disk Space	Total storage capacity for a computer
Available Disk Space	Amount of unused storage space on a computer (not in file systems seen by this operating system)

Owned Disk Space	Total amount of unique disk storage space for computers in a cluster that is owned by those computers. If Data Manager determines that a disk is configured for a different host, that disk's capacity is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column: • Computer disks on computers that have not been probed. • Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer. **Note:** • If the computer is an IBM SAN File System client, this value does not include the capacity of any volumes that are visible to it that are owned by the SAN File System. • If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads from the host ID. If this does not agree with the local host ID, the disk is classified as belonging to a different (Solaris) host.
Owned Available Disk Space	Total amount of unique disk storage space for computers in a cluster that is not allocated to any logical volume within the computers and is owned by the computers. If Data Manager determines that a disk is configured for a different computer, that disk's unallocated space is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column: • Computer disks on computers that have not been probed. • Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer. If the computer is an IBM SAN File System client, then this value excludes the unallocated space on any storage subsystem volumes visible to it that are owned by the SAN File System. If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product will examine the Veritas disk-label and read off the host-ID. If this does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.
Non-Fibre Channel Attached Disk Space	Total amount of unique disk space on the computers in a cluster that is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column: • Computer disks that reside behind a fibre channel port. • Computer disks on a computer that has not been probed. • Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.
Non-Fibre Channel Attached Available Disk Space	Total amount of disk storage space on the computers in a cluster that is not allocated to any logical volume within the computers and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column: • Computer disks that reside behind a fibre channel port. • Computer disks on a computer that has not been probed. • Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Fibre Channel Attached Disk Space	Total amount of unique disk space on the computers in a cluster that is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column: • Computer disks that do not reside behind a fibre channel port. • Computer disks that reside behind a fibre channel port that is not identified by a Data agent. • Computer disks on a computer that has not been probed. • Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.
Fibre Channel Attached Available Disk Space	Total amount of disk storage space on the computers in a cluster that is not allocated to any logical volume within the computers and is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column: • Computer disks that do not reside behind a fibre channel port. • Computer disks that reside behind a fibre channel port that is not identified by a Data agent. • Computer disks on a computer that has not been probed. • Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.
OS Type	Operating system running on a computer
Network Address	Network address of a computer
IP Address	IP address of a computer
Time Zone	Time zone in which a computer is running
CPU Architecture	Instruction set architecture on the machine where an agent is running.

Available File System Space

This report shows the total amount of unused storage across a network.

The following table describes the columns on the Total Freespace report. All the fields on the report are for display only.

Available File System Space	Total amount of available storage space available on a network
Percent Available File System Space	Percentage of total space that is unused on a network
Consumed File System Space	Amount of used storage space on a network
File System Space	Total amount (capacity) of storage space on a network
File Count	Total number of files on a network
Directory Count	Total number of directories on a network
Percent Free Inodes	Percent of free Inodes on a network
Used Inodes	Number of used Inodes on a network
Free Inodes	Number of free Inodes on a network

User Space Usage

This report provides information about storage statistics related to users within your environment.

The following table describes the columns on the User Space Usage report. All the fields on the report are for display only.

User Name	ID of a user
File Size	Total amount of space used by a user
File Count	Number of files owned or created by a user
Directory Count	Number of directories owned or created by a user
Largest File	Largest file owned by a user
2nd Largest File	Second-largest file owned by a user
Files < 1KB	Number and total space usage of files under 1KB in size
Files 1KB - 10KB	Number and total space usage of files between 1KB and 10KB in size
Files 10KB - 100KB	Number and total space usage of files between 10KB and 100KB in size
Files 100KB - 1MB	Number and total space usage of files between 100KB and 1MB in size
Files 1MB - 10MB	Number and total space usage of files between 1MB and 10MB in size
Files 10MB - 100MB	Number and total space usage of files between 10MB and 100MB in size
Files 100MB - 500MB	Number and total space usage of files between 100MB and 500MB in size
Files > 500MB	Number and total space usage of files over 500MB in size

Wasted Space

This report provides information about storage statistics on non-OS files not accessed in the first year and orphan files.

The following table describes the columns on the Wasted Space-File Summary by Computer report. All the fields on the report are for display only.

File Size	Total amount of space used by the obsolete and orphan files
File Count	Total number of obsolete and orphan files
Directory Count	Total number of orphan directories
Average File Size	Average size of obsolete and orphan files
File System Space	Amount of storage space available on a storage subsystem.

To view more information for a specific file:

- Click  . Detailed information per computer will be displayed in a **Network** tab. The **Network** tab displays file summary by computer.

System reports - Data Manager for Databases

Use these predefined reports to view statistics about your RDBMS storage resources. These reports are available in the **Data Manager for Databases** → **My Reports** → **System Reports** node of the navigation tree.

All DBMS - Database Storage By Computer

This report shows storage information about the databases in your environment sorted by the computers on which they are stored.

Computer	Name of the computer where the databases are located
Total Size	Amount of space consumed by the databases on the computers
File Capacity	Storage capacity of the data files within the databases
File Free space	Amount of free space available on the data files within the databases
DB-TS Count	Number tablespaces associated with the databases
File Count	Number of data files associated with the tablespaces in the databases
Log File Count	Number of log files associated with the databases

All DBMS - User Database Space Usage

This report shows information about the RDBMS storage space consumed by users across the network.

User Name	Name of a user upon whose storage usage information was collected
Total Size	Amount of space consumed by the objects owned by the user
Empty Used Space	Amount of empty used space residing on the objects owned by the user
Table Count	Number of tables owned by the user
Index Count	Number of indexes owned by the user
Other Count	Number of other objects owned by the user
Largest Object	Size of the largest object owned by the user
2nd Largest Object	Size of the second largest object owned by the user
Segment Size < 50KB	Number and total space usage of objects owned by the user that are under 50KB in size
Segment Size 50KB - 500KB	Number and total space usage of objects owned by the user that are between 50KB and 500KB in size
Segment Size 500KB - 1Mb	Number and total space usage of objects owned by the user that are between 500KB and 1MB in size
Segment Size 1Mb - 10Mb	Number and total space usage of objects owned by the user that are between 1MB and 10MB in size
Segment Size 10Mb - 100Mb	Number and total space usage of objects owned by the user that are between 10MB and 100MB in size
Segment Size 100Mb - 500Mb	Number and total space usage of objects owned by the user that are between 100MB and 500MB in size
Segment Size 500Mb - 1GB	Number and total space usage of objects owned by the user that are between 500MB and 1GB in size
Segment Size > 1Gb	Number and total space usage of objects owned by the user that are over 1GB in size

All DBMS - User Database Quota Violations

This report provides information on the users who violated any storage usage thresholds defined in quotas.

Quota Creator	ID of the user who created the quota
Quota Name	Name of the quota
Computer	Name of the computer where the quota violation occurred

RDBMS type	RDBMS to which an instance belongs, for example, Oracle, SQL Server, Sybase, or UDB
Instance	Name/SID of the Oracle Instance, SQL Server server, Sybase server, or UDB Instance where the quota violation occurred
Database-Tablespace	Name of the database where the tablespace resides and name of the tablespace whose quota was violated
User Name	ID of the user who violated the quota
Space Used	Amount of space consumed by the user
Threshold	Threshold of used space defined within the quota
% of Threshold	Percentage of space by which the user violated the quota
Number of Violations	Number of violations for the user
Quota Type	The type of quota: Database-Tablespace, Instance, or Network-wide

Oracle - Database Storage By Computer

This report shows storage information about the databases in your environment sorted by the computers on which they are stored.

Computer	Name of the computer where the databases are located
Total Size	Amount of space consumed by the databases on the computers
Data File Capacity	Storage capacity of the data files within the databases
Data File Free space	Amount of free space available on the data files within the databases
Tablespace Count	Number tablespaces associated with the databases
Data File Count	Number of data files associated with the tablespaces in the databases
Log File Count	Number of log files associated with the databases

Oracle - Total Database Freespace

This report provides information on the total free space for data files at a network-wide level, as well as the percentage of free space, the total used space, number of free extents, and number of data files on a network. This report is available for Oracle and Microsoft SQL Server.

Free Space	Amount of free space on all data files in the network.
Percent Free	Percent of free space available on the data files in the network.
Used Space	Amount of used space on the data files in the network.
Total Size	Total size of data files in the network.
Free Extents	Number of free extents on the data files in the network.
Coalesced Extents	Number of coalesced extents on the databases in the network.
Minimum Free Extent Size	Minimum size of a free extent in a database.
Maximum Free Extent Size	Maximum size of a free extent in a database.
Number Data Files	Number of data files on the databases in the network

Oracle - User Database Space Usage

This report shows information about the RDBMS storage space consumed by users across the network.

User Name	Name of a user upon whose storage usage information was collected
Total Size	Amount of space consumed by the objects owned by the user
Empty Used Space	Amount of empty used space residing on the objects owned by the user
Table Count	Number of tables owned by the user
Index Count	Number of indexes owned by the user
Other Count	Number of other objects owned by the user
Largest Segment	Size of the largest segment owned by the user
2nd Largest Segment	Size of the second largest segment owned by the user

Segment Size < 50KB	Number and total space usage of objects owned by the user that are under 50KB in size
Segment Size 50KB - 500KB	Number and total space usage of objects owned by the user that are between 50KB and 500KB in size
Segment Size 500KB - 1Mb	Number and total space usage of objects owned by the user that are between 500KB and 1MB in size
Segment Size 1Mb - 10Mb	Number and total space usage of objects owned by the user that are between 1MB and 10MB in size
Segment Size 10Mb - 100Mb	Number and total space usage of objects owned by the user that are between 10MB and 100MB in size
Segment Size 100Mb - 500Mb	Number and total space usage of objects owned by the user that are between 100MB and 500MB in size
Segment Size 500Mb - 1GB	Number and total space usage of objects owned by the user that are between 500MB and 1GB in size
Segment Size > 1Gb	Number and total space usage of objects owned by the user that are over 1GB in size

Oracle - Segments With Wasted Space

This report provides information on the segments containing allocated space that is currently empty/not being used. Use this information to gather information that can help you discover space that can be reclaimed and allocated to other objects.

Empty Used Space	Amount of empty used space within a segment (table, index, etc.)
Segment Creator	Owner of a segment
Segment Name	Name of a segment
Computer	Name of the computer on which the segment's Instance resides
Instance	SID of an Oracle Instance
Database	Name of the database to which the segment belongs
Tablespace	Name of the tablespace to which the segment belongs
Partition	Partition on which the segment is stored
Segment Type	Type of the segment, including: • TABLE • TABLE PARTITION • TABLE SUBPARTITION • NESTED TABLE • CLUSTER • INDEX • INDEX PARTITION • INDEX SUBPARTITION • LOBINDEX • LOBSEGMENT • LOB PARTITION • LOB SUBPARTITION
Parent Type	Subset of the segment type
Parent Creator	Owner of a segment
Parent Name	Name of a segment
Total Size	Amount of space allocated to a segment
Number of Extents	Number of extents allocated to a segment
Freelist Blocks	Number of blocks on the freelist chain
Initial Extent	Size of the first extent allocated to a segment
Next Extent	Amount of space Oracle will retrieve when allocating another extent
Maximum Extents	Maximum number of extents that Oracle would allocate to an object
Percent Increase	Percent increase is size Oracle will allocate for the next extent

SqlServer - Database Storage By Computer

This report shows storage information about the databases in your environment sorted by the computers on which they are stored.

Computer	Name of the computer on which SQL Server is running
Total Size	Total amount of storage space allocated to SQL Server on a computer
Data File Capacity	Amount of allocated storage space for data files that is unused on a computer where SQL Server is running
Data File Free space	Amount of free space available on the data files associated with SQL Server on a computer
Log File Capacity	Amount of allocated storage space for log files that is unused on a computer where SQL Server is running
Log File Free space	Amount of free space available on the log files associated with SQL Server on a computer
Database Count	Number of databases associated with SQL Server on a computer
Data File Count	Number of data files associated with SQL Server on a computer
Log File Count	Number of log files associated with SQL Server on a computer

SqlServer - Total Database Freespace

This report provides information on the total free space for data files at a network-wide level, as well as the percentage of free space, the total used space, number of free extents, and number of data files on a network. This report is available for Oracle and Microsoft SQL Server.

Data Free Space	Amount of free space on all data files in the network.
Percent Free	Percent of free space available on the data files in the network.
Data Used Space	Amount of used space on the data files in the network.
Data Size	Total size of data files in the network.
Log Free Space	Amount of free space.
Log Percent Used	Percent of used space.
Log Used Space	Amount of space used by the log.
Log Size	Size of the log.
Number Data Files	Number of data files on the databases in the network

SqlServer - User Database Space Usage

This report shows information about the RDBMS storage space consumed by users across the network.

User Name	Name of a user upon whose storage usage information was collected
Total Size	Amount of space consumed by the objects owned by the user
Empty Used Space	Amount of empty used space residing on the objects owned by the user
Table Count	Number of tables owned by the user
Index Count	Number of indexes owned by the user
Other Count	Number of other objects owned by the user
Largest Object	Size of the largest object owned by the user
2nd Largest Object	Size of the second largest object owned by the user
Object Size < 50KB	Number and total space usage of objects owned by the user that are under 50KB in size
Object Size 50KB - 500KB	Number and total space usage of objects owned by the user that are between 50KB and 500KB in size
Object Size 500KB - 1Mb	Number and total space usage of objects owned by the user that are between 500KB and 1MB in size
Object Size 1Mb - 10Mb	Number and total space usage of objects owned by the user that are between 1MB and 10MB in size

Object Size 10Mb - 100Mb	Number and total space usage of objects owned by the user that are between 10MB and 100MB in size
Object Size 100Mb - 500Mb	Number and total space usage of objects owned by the user that are between 100MB and 500MB in size
Object Size 500Mb - 1GB	Number and total space usage of objects owned by the user that are between 500MB and 1GB in size
Object Size > 1Gb	Number and total space usage of objects owned by the user that are over 1GB in size

Sybase - Database Storage By Computer

This report shows storage information about the databases in your environment sorted by the computers on which they are stored.

Computer	Name of the computer on which Sybase is running
Total Device Size	Total amount of storage space allocated to Sybase on a computer
Mirrored Device Size	Total storage capacity of the mirrored device for an instance
Device Free space	Amount of free space available on the data files associated with Sybase on a computer
Data Fragment Capacity	Amount of allocated storage space for log files that is unused on a computer where Sybase is running
Data Fragment Free space	Amount of free space available on the data fragments associated with Sybase on a computer
Log Fragment Capacity	Amount of allocated storage space for log fragments that is unused on a computer where Sybase is running
Log Fragment Free space	Amount of free space available on the log files associated with Sybase on a computer
Device Count	Number of devices associated with Sybase on a computer
Database Count	Number of databases associated with Sybase on a computer
Fragment Count	Number of fragments associated with Sybase on a computer
Log File Count	Number of log files associated with Sybase on a computer

Sybase - Total Device Freespace

This report provides information on the total free space for Sybase devices at a network-wide level, as well as the percentage of free space, the total used space, number of fragments, and number of devices.

Free Space	Amount of free space on all the devices in the network
Percent Free	Percent of free space available on the devices in the network
Used Space	Amount of used space on the devices in the network
Total Size	Total size of devices in the network
Number Fragments	Number of fragments associated with the devices in the network
Number Devices	Number of devices in the network

Sybase - Total Database Freespace

This report provides information on the total free space for data files at a network-wide level, as well as the percentage of free space, the total used space, number of free extents, and number of data files on a network. This report is available for Oracle and Microsoft SQL Server.

Data Free Space	Amount of free space on all data files in the network.
Percent Free	Percent of free space available on the data files in the network.
Data Used Space	Amount of used space on the data files in the network.
Data Size	Total size of data files in the network.
Log Free Space	Amount of free space available to a database log

Log Percent Used	Percentage of allocated space that a database log is consuming
Log Used Space	Amount of storage space consumed by a database log
Log Size	Total amount of space allocated to a database log
Number Fragments	nNumber of fragments associated with a database

Sybase - User Database Space Usage

This report shows information about the RDBMS storage space consumed by users across the network.

User Name	Name of a user upon whose storage usage information was collected
Total Size	Amount of space consumed by the objects owned by the user
Empty Used Space	Amount of empty used space residing on the objects owned by the user
Table Count	Number of tables owned by the user
Index Count	Number of indexes owned by the user
Other Count	Number of other objects owned by the user
Largest Object	Size of the largest object owned by the user
2nd Largest Object	Size of the second largest object owned by the user
Object Size < 50KB	Number and total space usage of objects owned by the user that are under 50KB in size
Object Size 50KB - 500KB	Number and total space usage of objects owned by the user that are between 50KB and 500KB in size
Object Size 500KB - 1Mb	Number and total space usage of objects owned by the user that are between 500KB and 1MB in size
Object Size 1Mb - 10Mb	Number and total space usage of objects owned by the user that are between 1MB and 10MB in size
Object Size 10Mb - 100Mb	Number and total space usage of objects owned by the user that are between 10MB and 100MB in size
Object Size 100Mb - 500Mb	Number and total space usage of objects owned by the user that are between 100MB and 500MB in size
Object Size 500Mb - 1GB	Number and total space usage of objects owned by the user that are between 500MB and 1GB in size
Object Size > 1Gb	Number and total space usage of objects owned by the user that are over 1GB in size

UDB - Instance Storage By Computer

This report shows information about the UDB instances in your environment sorted by the computers on which they are stored.

Computer	Name of the computer on which a UDB Instance(s) resides
Total Size	Total amount of storage space allocated UDB on a computer
Container Capacity	Amount of storage space allocated to the containers that reside on a computer
Container Free space	Amount of free space available on the containers that reside on a computer
Log File Capacity	Amount of storage space allocated to the log files that reside on a computer
Tablespace Count	Number of tablespaces that reside on a computer
Container Count	Number of containers that reside on a computer
Log File Count	Number of log files that reside on a computer

UDB - Total DMS Container Freespace

This report shows the total free space for the containers associated with DMS tablespaces on UDB instances within your environment.

Free Space	Amount of free space available on the DMS containers within a network
Percent Free	Percentage of free space available on the DMS containers within a network

Used Space	Amount of storage space consumed on the DMS containers within a network
Total Size	Total amount of space on the DMS containers within a network
Number Containers	Number of DMS containers within a network

UDB - User Database Space Usage

This report shows information about the RDBMS storage space consumed by users across the network.

User Name	Name of a user upon whose storage usage information was collected
Total Size	Amount of space consumed by the objects owned by the user
Empty Used Space	Amount of empty used space residing on the objects owned by the user
Table Count	Number of tables owned by the user
Index Count	Number of indexes owned by the user
Other Count	Number of other objects owned by the user
Largest Object	Size of the largest segment owned by the user
2nd Largest Object	Size of the second largest segment owned by the user
Object Size < 50KB	Number and total space usage of objects owned by the user that are under 50KB in size
Object Size 50KB - 500KB	Number and total space usage of objects owned by the user that are between 50KB and 500KB in size
Object Size 500KB - 1Mb	Number and total space usage of objects owned by the user that are between 500KB and 1MB in size
Object Size 1Mb - 10Mb	Number and total space usage of objects owned by the user that are between 1MB and 10MB in size
Object Size 10Mb - 100Mb	Number and total space usage of objects owned by the user that are between 10MB and 100MB in size
Object Size 100Mb - 500Mb	Number and total space usage of objects owned by the user that are between 100MB and 500MB in size
Object Size 500Mb - 1GB	Number and total space usage of objects owned by the user that are between 500MB and 1GB in size
Object Size > 1Gb	Number and total space usage of objects owned by the user that are over 1GB in size

System Reports - Disk Manager

Use these predefined reports to view performance information about storage subsystems.

The following system reports are available for Disk Manager .

Array Performance

This report shows performance data for arrays. This report supports DS6000, DS8000, and Tivoli Storage Enterprise Storage Server models only.

Column	Description
Subsystem	ID of the storage subsystem
Array	The ID of the specified array.
Time	Date and time that the data was collected.

Column	Description
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.
Disk Utilization Percentage	The percentage of utilization for the disk the on which the array is located.
Sequential I/O Percentage	Percentage of all I/O operations that were sequential operations.
Backend Read I/O Rate	Average number of I/O operations per second for read operations.
Backend Write I/O Rate	Average number of I/O operations per second for write operations.
Total Backend I/O Rate	Average number of I/O operations per second for read and write operations.
Backend Read Data Rate	Average number of Mbps that were transferred for read operations.
Backend Write Data Rate	Average number of Mbps that were transferred for write operations.
Total Backend Data Rate	Average number of Mbps that were transferred for read and write operations.
Backend Read Response Time	Average number of milliseconds that it took to respond to each read operation. For SAN Volume Controller models, this is the external response time of the managed disks (MDisks).
Backend Write Response Time	Average number of milliseconds that it took to respond to each write operation. For SAN Volume Controller models, this is the external response time of the managed disks (MDisks).
Overall Backend Response Time	Average number of milliseconds that it took to respond to each I/O operation (read and write). For SAN Volume Controller models, this is the external response time of the managed disks (MDisks).
Write-cache Delay Percentage	Percentage of I/O operations that were delayed due to write-cache space constraints or other conditions. Write-cache Delay Percentage is the ratio of delayed I/O operations to total I/O operations.
Write-cache Delay I/O Rate	Average number of I/O operations per second that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval.

Controller Cache Performance

This report shows performance data for storage-subsystem controller caches. This report supports DS6000, DS8000, and Tivoli Storage Enterprise Storage Server models only.

Column	Description
Subsystem	ID of the storage subsystem
Controller	The ID of the controller.
Time	Date and time that the data was collected.
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.
Read I/O Rate (normal)	Average number of I/O operations per second for nonsequential read operations.
Read I/O Rate (sequential)	Average number of I/O operations per second for sequential read operations.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read operations.
Write I/O Rate (normal)	Average number of I/O operations per second for nonsequential write operations.
Write I/O Rate (sequential)	Average number of I/O operations per second for sequential write operations.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential write operations.
Total I/O Rate (normal)	Average number of I/O operations per second for nonsequential read and write operations.
Total I/O Rate (sequential)	Average number of I/O operations per second for sequential read and write operations.
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Cache Holding Time	The duration allocated for cache holding (in seconds).
Read Cache Hits Percentage (overall)	Overall read cache percentage for both sequential and nonsequential read operations.
Write Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and nonsequential write operations.
Total Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and nonsequential read and write operations.

Column	Description
Read Cache Hits Percentage (normal)	The normal read cache percentage allocated for the controller.
Write Cache Hits Percentage (normal)	The normal write cache percentage allocated for the controller.
Total Cache Hits Percentage (normal)	The total normal cache percentage (read and write) allocated for the controller.
Read Cache Hits Performance (sequential)	The sequential read cache allocated for the controller.
Write Cache Hits Performance (sequential)	The sequential write cache allocated for the controller.
Total Cache Hits Performance (sequential)	The total sequential cache percentage (read and write) allocated for the controller.
Record Mode Read I/O Rate	The record mode read I/O rate for the controller.
Record Mode Read Cache Hits Percentage	The record mode read I/O cache for the controller.
Disk to Cache Transfer Rate	The disk to cache I/O rate for the controller.
Cache to Disk TransferRate	The cache to disk I/O rate for the controller.
Write-cache Delay Percentage	Percentage of I/O operations that were delayed due to write-cache space constraints or other conditions. Write-cache Delay Percentage is the ratio of delayed I/O operations to total I/O operations.
Write-cache Delay I/O Rate	The write-cache delayed I/O rate for the controller.

Controller Performance

This report shows performance data for storage-subsystem controllers. This report supports DS6000, DS8000, and Tivoli Storage Enterprise Storage Server models only.

Column	Description
Subsystem	ID of the storage subsystem
Controller	The ID of the controller.
Time	Date and time that the data was collected.
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.
Read I/O Rate (overall)	Average number of Mbps that were transferred for read operations.
Write I/O Rate (overall)	Average number of Mbps that were transferred for write operations.

Column	Description
Total I/O Rate (overall)	Average number of megabytes per second that were transferred for read and write operations.
Read Data Rate	Average number of I/O operations per second for both sequential and nonsequential read operations.
Write Data Rate	Average number of I/O operations per second for both sequential and nonsequential write operations.
Total Data Rate	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Response Time	Average number of milliseconds that it took to respond to each read operation.
Write Response Time	Average number of milliseconds that it took to respond to each write operation.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write).

I/O Group Performance

This report shows performance data for I/O groups. This report supports SAN Volume Controller models only.

Column	Description
Subsystem	ID of the storage subsystem
I/O Group	The ID of the specified I/O group.
Time	Date and time that the data was collected.
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.
CPU Utilization Percentage	Average utilization percentage of the CPUs.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read operations.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential write operations.
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Data Rate	Average number of Mbps that were transferred for read operations.

Column	Description
Write Data Rate	Average number of Mbps that were transferred for write operations.
Total Data Rate	Average number of megabytes per second that were transferred for read and write operations.
Read Response Time	Average number of milliseconds that it took to respond to each read operation.
Write Response Time	Average number of milliseconds that it took to respond to each write operation.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write).
Port to Host Send I/O Rate	Average number of exchanges (I/Os) per second sent to host computers.
Port to Host Receive I/O Rate	Average number of exchanges (I/Os) per second received from host computers.
Total Port to Host I/O Rate	Average number of exchanges (I/Os) per second transmitted between host computers.
Port to Disk Send I/O Rate	Average number of exchanges (I/Os) per second sent to storage subsystems
Port to Disk Receive I/O Rate	Average number of exchanges (I/Os) per second received from storage subsystems.
Total Port to Disk I/O Rate	Average number of exchanges (I/Os) per second transmitted between storage subsystems.
Port to Host Send Data Rate	Average number of megabytes per second sent to host computers.
Port to Host Receive Data Rate	Average number of megabytes per second received from host computers.
Total Port to Host Data Rate	Average number of megabytes per second transmitted between host computers.
Port to Disk Send Data Rate	Average number of megabytes per second sent to storage subsystems.
Port to Disk Receive Data Rate	Average number of megabytes per second received from storage subsystems.
Total Port to Disk Data Rate	Average number of megabytes per second transmitted between storage subsystems.

Managed Disk Group Performance

This report shows performance data for managed-disk (MDisk) groups. This report supports SAN Volume Controller models only.

Column	Description
Subsystem	ID of the storage subsystem
Managed Disk Group	The ID of the specified managed disk.
Time	Date and time that the data was collected.

Column	Description
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.
Read Transfer Size	Average number of KB per I/O for read operations, for a particular component over an interval.
Write Transfer Size	Average number of KB per I/O for write operations, for a particular component over an interval.
Overall Transfer Size	Average number of KB per I/O for read and write operations, for a particular component over an interval.
Backend Read I/O Rate	Average number of I/O operations per second for read operations.
Backend Write I/O Rate	Average number of I/O operations per second for write operations.
Total Backend I/O Rate	Average number of I/O operations per second for read and write operations.
Backend Read Data Rate	Average number of Mbps that were transferred for read operations.
Backend Write Data Rate	Average number of Mbps that were transferred for write operations.
Total Backend Data Rate	Average number of Mbps that were transferred for read and write operations.

Node Cache Performance report

This report lists the most recently gathered performance data for SAN Volume Controller nodes, which is aggregated to the node level from the per-node VDisk sample data.

Column	Description
Subsystem	ID of the storage subsystem
Node	The ID of the specified node.
Time	Date and time that the data was collected.
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.

Column	Description
Read Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and nonsequential read operations, for a particular component over an interval.
Write Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and nonsequential write operations, for a particular component over an interval.
Total Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and nonsequential read and write operations, for a particular component over an interval.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read operations.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential write operations.
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Disk to Cache Transfer Rate	Average number of I/O operations (track transfers) per second for disk to cache transfers, for a particular component over an interval.
Cache to Disk Transfer Rate	Average number of I/O operations (track transfers) per second for cache to disk transfers, for a particular component over an interval.
Write-cache Delay Percentage	Percentage of I/O operations that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval. (The ratio of delayed operations to total I/Os.)
Write-cache Overflow Percentage	Percentage of write operations that were delayed because of lack of write-cache space, for a particular component over an interval.
Write-cache Flush-through Percentage	Percentage of write operations that were processed in flush-through write mode, for a particular component over an interval.
Write Cache Write-through Percentage	Percentage of write operations that were processed in write-through write mode, for a particular component over an interval.
Write-cache Delay I/O Rate	Average number of I/O operations per second that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval.
Write-cache Overflow I/O Rate	Average number of tracks per second that were delayed because of lack of write-cache space, for a particular component over an interval.

Column	Description
Write Cache Flush-through I/O Rate	Average number of tracks per second that were processed in flush-through write mode, for a particular component over an interval.
Write Cache Write-through I/O Rate	Average number of tracks per second that were processed in write-through write mode, for a particular component over an interval.
Readahead Percentage of Cache Hits	Percentage of all read cache hits which occurred on prestaged data.
Dirty Write Percentage of Cache Hits	Percentage of all write cache hits which occurred on already dirty data in the cache.

Port Performance

This report shows performance data for ports. This report supports DS6000, DS8000, Tivoli Storage Enterprise Storage Server, and other storage subsystems .

Column	Description
Subsystem	ID of the storage subsystem
Port	The ID of the specified port.
WWPN	The World Wide Port Name (WWPN) of the selected port.
Time	Date and time that the data was collected.
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.
Port Send I/O Rate	The rate (in bytes) of send ops for the port.
Port Receive I/O Rate	The rate (in bytes) of receive ops for the port.
Total Port I/O Rate	The total rate (in bytes) of send and receive ops for the port.
Port Send Data Rate	The send data rate (in bytes) for the port.
Port Receive Data Rate	The receive data rate (in bytes) for the port.
Total Port Transfer Rate	The total rate (in bytes) of send and receive data rates for the port.
Port to Host Send I/O Rate	The time (in seconds) of the send response time for the port.
Port to Host Receive I/O Rate	The time (in seconds) of the receive response time for the port.
Total Port to Host I/O Rate	The average number of exchanges (I/Os) per second transmitted between host computers.
Port to Disk Send I/O Rate	The average number of exchanges (I/Os) per second sent to storage subsystems.

Column	Description
Total Port to Remote Node I/O Rate	The average number of exchanges (I/Os) per second transmitted between storage subsystems.
Port to Host Send Data Rate	The send data rate (in bytes) for the port.
Port to Host Receive Data Rate	The receive data rate (in bytes) for the port.
Total Port to Host Data Rate	The total rate (in bytes) of send and receive data rates for the port.
Port to Disk Send Data Rate	The time (in seconds) of the send response time for the port.
Port to Disk Receive Data Rate	The time (in seconds) of the receive response time for the port.
Total Port Data Rate	The total (in seconds) of send and receive response times for the port.
Port Send Response Time	The size (in bytes) of data being transferred from the port.
Port Receive Response Time	The size (in bytes) of data being transferred to the port.
Total Port Response Time	The total (in seconds) of send and receive response times for the port.
Port Send Transfer Size	The size (in bytes) of data being transferred from the port.
Port Receive Transfer Size	The size (in bytes) of data being transferred to the port.
Total Transfer Size	The size (in bytes) of data being transferred to and from the port.

Subsystem Performance

This report shows performance data for storage subsystems. This report supports all subsystems.

Column	Description
Subsystem	ID of the storage subsystem
Time	Date and time that the data was collected.
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read operations.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential write operations.

Column	Description
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Data Rate	Average number of Mbps that were transferred for read operations.
Write Data Rate	Average number of Mbps that were transferred for write operations.
Total Data Rate	Average number of megabytes per second that were transferred for read and write operations.
Read Response Time	Average number of milliseconds that it took to respond to each read operation.
Write Response Time	Average number of milliseconds that it took to respond to each write operation.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write).

Top Active Volumes Cache Hit Performance

This report summarizes the primary metrics for active volumes, prioritized by cache hit ratio. Only volumes with an I/O rate that is greater than 0 will be included in this report. If all volumes are inactive, this report will be empty. This report supports all storage subsystems.

Column	Description
Subsystem	ID of the storage subsystem
Volume	The ID of the volume.
Time	Date and time that the data was collected.
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.
Read Cache Hits Percentage (overall)	Overall read cache percentage for both sequential and nonsequential read operations.
Write Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and nonsequential write operations.
Total Cache Hits Percentage (overall)	Percentage of cache hits for both sequential and nonsequential read and write operations.
Readahead Percentage of Cache Hits	Percentage of all read cache hits which occurred on prestaged data.
Dirty Write Percentage of Cache Hits	Percentage of all write cache hits which occurred on already dirty data in the cache.

Column	Description
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read operations.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential write operations.
Total I/O Rate (overall	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Data Rate	Average number of Mbps that were transferred for read operations.
Write Data Rate	Average number of Mbps that were transferred for write operations.
Total Data Rate	Average number of megabytes per second that were transferred for read and write operations.
Read Response Time	Average number of milliseconds that it took to respond to each read operation.
Write Response Time	Average number of milliseconds that it took to respond to each write operation.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write).

Top Volumes Data Rate Performance

This report summarizes the primary performance metrics for volumes, prioritized by data rate. This report supports all storage subsystems.

Column	Description
Subsystem	ID of the storage subsystem
Volume	The ID of the specified volume.
Time	Date and time that the data was collected.
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.
Read Data Rate	Average number of Mbps that were transferred for read operations.
Write Data Rate	Average number of Mbps that were transferred for write operations.
Total Data Rate	Average number of megabytes per second that were transferred for read and write operations.

Column	Description
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read operations.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential write operations.
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Response Time	Average number of milliseconds that it took to respond to each read operation.
Write Response Time	Average number of milliseconds that it took to respond to each write operation.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write).

Top Volumes Disk Performance

This report summarizes the primary performance metrics for volumes, prioritized by disk rate. This report supports all storage subsystems.

Column	Description
Subsystem	ID of the storage subsystem
Volume	The ID of the volume.
Time	Date and time that the data was collected.
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.
Disk to Cache Transfer Rate	Average number of I/O operations (track transfers) per second for disk to cache transfers, for a particular component over an interval.
Cache to Disk Transfer Rate	Average number of I/O operations (track transfers) per second for cache to disk transfers, for a particular component over an interval.
Write-cache Delay Percentage	Percentage of I/O operations that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval. (The ratio of delayed operations to total I/Os.)
Write-cache Overflow Percentage	Percentage of write operations that were delayed because of lack of write-cache space, for a particular component over an interval.

Column	Description
Write-cache Flush-through Percentage	Percentage of write operations that were processed in flush-through write mode, for a particular component over an interval.
Write Cache Write-through Percentage	Percentage of write operations that were processed in write-through write mode, for a particular component over an interval.
Write-cache Delay I/O Rate	Average number of I/O operations per second that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval.
Write-cache Overflow I/O Rate	Average number of tracks per second that were delayed because of lack of write-cache space, for a particular component over an interval.
Write Cache Flush-through I/O Rate	Average number of tracks per second that were processed in flush-through write mode, for a particular component over an interval.
Write Cache Write-through I/O Rate	Average number of tracks per second that were processed in write-through write mode, for a particular component over an interval.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read operations.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential write operations.
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Data Rate	Average number of Mbps that were transferred for read operations.
Write Data Rate	Average number of Mbps that were transferred for write operations.
Total Data Rate	Average number of megabytes per second that were transferred for read and write operations.
Read Response Time	Average number of milliseconds that it took to respond to each read operation.
Write Response Time	Average number of milliseconds that it took to respond to each write operation.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write).

Top Volumes I/O Rate Performance

This report summarizes the primary performance metrics for volumes, prioritized by I/O rate. This report supports all storage subsystems.

Column	Description
Subsystem	ID of the storage subsystem
Volume	The ID of the specified volume.
Time	Date and time that the data was collected.
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read operations.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential write operations.
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Data Rate	Average number of Mbps that were transferred for read operations.
Write Data Rate	Average number of Mbps that were transferred for write operations.
Total Data Rate	Average number of megabytes per second that were transferred for read and write operations.
Read Response Time	Average number of milliseconds that it took to respond to each read operation.
Write Response Time	Average number of milliseconds that it took to respond to each write operation.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write).

Top Volumes Response Performance

This report summarizes the primary performance metrics for volumes, prioritized by response rate. This report supports all storage subsystems.

Column	Description
Subsystem	ID of the storage subsystem
Volume	The ID of the volume.
Time	Date and time that the data was collected.

Column	Description
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.
Read Response Time	Average number of milliseconds that it took to respond to each read operation.
Write Response Time	Average number of milliseconds that it took to respond to each write operation.
Overall Response Time	Average number of milliseconds that it took to service each I/O operation (read and write).
Peak Read Response Time	The peak (worst) response time for disk readsfor a component over a time interval.
Peak Write Response Time	The peak (worst) response time for disk writes for a component over a time interval.
Read I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read operations.
Write I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential write operations.
Total I/O Rate (overall)	Average number of I/O operations per second for both sequential and nonsequential read and write operations.
Read Data Rate	Average number of Mbps that were transferred for read operations.
Write Data Rate	Average number of Mbps that were transferred for write operations.
Total Data Rate	Average number of megabytes per second that were transferred for read and write operations.

System reports - Fabric Manager

Use these predefined reports to view performance information about fabrics.

Port Connections

Use the Port Connections report to view information about the ports on the switches in your storage area network (SAN). This report includes information about the port connections and status, as well as the port and switch IDs.

Fabric ID	The ID of the fabric that the switch is in.
Switch	The label value of the switch.
Switch Port	The ID of the port (*domainID,port_number*).
Slot	The number for the blade that the port is on.
Index	The unique number for the port across the whole switch.

Switch Port State	The state of the switch port: Normal, Warning, or Critical.
Connection State	The state of the connection between the two ports: Normal, Warning, or Critical.
Connected Port	The port that the switch port is connected to. This value is the WWN of the device connected to that port.
Connected Port State	The state of the connected port: Normal, Warning, or Critical.
Device	The device that the connected port is on.
Device State	The state of the device that the connected port is on: Normal, Warning, or Critical.

SAN Assets (ALL)

Use the SAN Assets (All) report to generate and view reports that provide information about all the assets discovered by IBM Tivoli Storage Productivity Center.

The following table describes the columns on the SAN Assets (ALL) report. All the columns in the report are for display only.

Fabric ID	The ID of the fabric that the asset is in.
Type	The type of asset discovered.
Label	The label of the asset.
Identifier	The unique identifier of the storage asset. For example, a switch WWN.
Status	The state of the device, either Normal, Warning, or Critical.

SAN Assets (Connected Devices)

Use the SAN Assets (Connected Devices) report to generate and view reports that provide information about all the connected devices discovered by IBM Tivoli Storage Productivity Center.

The following table describes the columns on the SAN Assets (Connected Devices) report. All the columns in the report are for display only.

Fabric ID	The ID of the fabric that the asset is in.
Type	The type of asset that was discovered.
Label	The label of the asset.
Identifier	The unique identifier of the asset. For example, the device worldwide name.
Status	The status of the asset: Normal, Warning, or Critical.

SAN Assets (Switches)

Use the SAN Assets (Switches) report to generate and view reports that provide information about switches in your SAN. These reports provide information about the switch ID, type, manufacturer, WWN, location, and other details.

The following table describes the columns on the SAN Assets (Switches) report. All the columns in the report are for display only.

Fabric ID	The ID of the fabric that the switch is in.
Type	The type is always ″Switch″.
Switch Type	The type of switch discovered. For example, a physical switch.

Label	The label of the switch.
Status	The status of the switch. The possible states are **Normal** and **Missing**.
IP Address	The IP address of the switch.
Vendor ID	The manufacturer of the switch.
Model	The model name and number of the switch.
Version	The version information of the switch.
Serial Number	The serial number of the switch.
Object ID	The object ID of the switch.
Management ID	The management ID of the switch.
Management Address	The URL of the switch element manager.
Domain	The ID of the domain that the switch is in.
Contact	The contact information for the switch.
WWN	The world-wide name (WWN) of the switch.
Location	The physical location of the switch.
Physical Switch WWN	The world-wide name (WWN) of the physical switch.

Switch Performance

Use the Switch Performance report to generate and view reports that provide information about the free and occupied ports on the switches in your SAN. These reports provide information about the port connections and status, as well as the port and switch IDs, data and packet transfer rates, and other details. This report can only be generated for switches that have had performance monitors run on them.

This following table describes the columns on the Switch Performance report. All the columns in the report are for display only.

Switch	The ID of the switch.
Time	The date and time of the data collection.
Interval	The size of the sample interval, in seconds.
Port Send Packet Rate	The rate at which packets are sent through this port.
Port Receive Packet Rate	The rate at which packets are received through this port.
Total Port Packet Rate	The total rate at which packet operations are performed through this port.
Port Send Data Rate	The rate at which data is sent through this port.
Port Receive Data Rate	The rate at which data is received through this port.
Total Port Data Rate	The total rate at which data is processed through this port.
Port Peak Send Data Rate	The peak rate at which data is sent through this switch.
Port Peak Receive Data Rate	The peak rate at which data is received through this switch.
Port Send Packet Size	The size of the data block that is sent out through this switch.
Port Receive Packet Size	The size of the data block that is received in through this switch.
Overall Port Packet Size	The total size of the data block that is transferred by the switch.
Error Frame Rate	The rate at which error frames are generated by the switch.
Dumped Frame Rate	The rate at which frames are dumped by the switch.
Link Failure Rate	The rate at which failed links are generated by the switch.

Switch Port Errors

Use the Switch Port Errors report to generate and view reports that provide information about the errors being generated by ports on the switches in your SAN. These reports provide information about the error frames, dumped frames, link failures, port and switch IDs, and other details. This report can only be generated for switches that have had performance monitors run on them.

To view more information for a specific file:

- Click  . Detailed information per computer will be displayed in a **Network** tab. The **Network** tab displays file summary by computer.

This table describes the fields on the Switch Port Errors report. All the fields on the report are for display only.

Switch	The ID of the switch.
Port	The ID of the port.
WWPN	The World Wide Port Name (WWPN) of the selected port.
Slot	The number for the blade that the port is on.
Index	The unique number for the port across the whole switch.
Time	The date and time of the data collection.
Interval	The size of the sample interval, in seconds.
Port Send Packet Size	The size of the data block that is sent out through this switch.
Port Receive Packet Size	The size of the data block that is received in through this switch.
Overall Port Packet Size	The total size of the data block that is transferred by the switch
Error Frame Rate	The rate at which error frames are generated by the switch.
Dumped Frame Rate	The rate at which frames are dumped by the switch.
Link Failure Rate	The rate at which failed links are generated by the switch.
Loss of Sync Rate	The average number of times per second that synchronization was lost after the last reset of the device, for a particular component over a particular time interval.
Loss of Signal Rate	The average number of times per second that the signal was lost after the last reset of the device, for a particular component over a particular time interval.
CRC Error Rate	The average number of frames received per second in which the CRC in the frame did not match the CRC computed by the receiver, for a particular component over a particular time interval.
Short Frame Rate	The average number of frames received per second that were shorter than 28 octets (24 header + 4 CRC), not including any start of frame/end of frame bytes, for a particular component over a particular time interval.
Long Frame Rate	The average number of frames received per second that were longer than 2140 octets (24 header + 4 CRC + 2112 data), not including any start of frame/end of frame bytes, for a particular component over a particular time interval.
Encoding Disparity Error Rate	The average number of disparity errors received per second, for a particular component over a particular time interval.

This table describes columns that are not part of the Switch Port Errors system report. However, you can include them, and save the resulting report in the ***<user-id> Reports*** node.

Port Send Packet Rate	The rate at which packets are sent through this port.
Port Receive Packet Rate	The rate at which packets are received through this port.
Total Port Packet Rate	The total rate at which packet operations are performed through this port.
Port Send Data Rate	The rate at which data is sent through this port.
Port Receive Data Rate	The rate at which data is received through this port.

Total Port Data Rate	The total rate at which data is processed through this port.
Port Peak Send Data Rate	The peak rate at which data is sent through this switch.
Port Peak Receive Data Rate	The peak rate at which data is received through this switch.

Top Switch Ports Data Rate Performance

Use the Top Switch Port Data Rate Performance report to generate and view reports that provide information about the switches in your SAN that have the top data rates. This report can only be generated for switches that have had performance monitors run on them.

This report only displays rows for the 25 highest data rates from the switch ports monitored. It displays those rows in descending order by Total Port Data Rate. You can adjust the number of rows shown on the **Selection** tab.

To view more information for a specific file:

- Click  . Detailed information per computer will be displayed in a **Network** tab. The **Network** tab displays the file summary by computer.

This section describes the fields on the Top Switch Ports Data Rate Performance report. All the fields on the report are for display only.

Switch	The ID of the switch.
Port	The ID of the port.
Time	The date and time of the data collection.
Interval	The size of the sample interval, in seconds.
Total Port Data Rate	The total rate at which data is processed through this port.
Total Port Packet Rate	The total rate at which the packet operations are performed through this port.

This table describes columns that are not part of the Top Switch Ports Data Rate Performance report. However, you can include them, and save the resulting report in the ***<user> Reports*** node.

Port Send Packet Rate	The rate at which packets are sent through this port.
Port Receive Packet Rate	The rate at which packets are received through this port.
Port Send Data Rate	The rate at which data is sent through this port.
Port Receive Data Rate	The rate at which data is received through this port.
Port Peak Send Data Rate	The peak rate at which data is sent through this switch.
Port Peak Receive Data Rate	The peak rate at which data is received through this switch.
Port Send Packet Size	The size of the data block that is sent out through this switch.
Port Receive Packet Size	The size of the data block that is received in through this switch.
Overall Port Packet Size	The total size of the data block that is transferred by the switch
Error Frame Rate	The rate at which error frames are generated by the switch.
Dumped Frame Rate	The rate at which frames are dumped by the switch.
Link Failure Rate	The rate at which failed links are generated by the switch.
WWPN	The World Wide Name identifier for the FC port.
Slot	The number for the blade that the port is on.
Index	The unique number for the port across the whole switch.

Top Switch Ports Packet Rate Performance

Use the Top Switch Ports Packet Rate Performance report to generate and view reports that provide information about the switches in your SAN that have the top packet rates. This report can only be generated for switches that have had performance monitors run on them.

This report only displays rows for the 25 highest packet rates from the switch ports monitored. It displays those rows in descending order by Total Port Packet Rate. You can adjust the number of rows shown on the Selection tab.

To view more information for a specific file:

- Click  . Detailed information per computer will be displayed in a **Network** tab. The **Network** tab displays file summary by computer.

This section describes the fields on the Top Switch Ports Packet Rate Performance report. All the fields on the report are for display only.

Switch	The ID of the switch.
Port	The ID of the port.
Time	The date and time of the data collection.
Interval	The size of the sample interval, in seconds.
Total Port Packet Rate	The total rate at which the packet operations are performed through this port.
Total Port Data Rate	The total rate at which data is processed through this port.

This table describes columns that are not part of the Top Switch Ports Packet Rate Performance report. However, you can include them, and save the resulting report in the ***<user> Reports*** node.

Port Send Packet Rate	The rate at which packets are sent through this port.
Port Receive Packet Rate	The rate at which packets are received through this port.
Port Send Data Rate	The rate at which data is sent through this port.
Port Receive Data Rate	The rate at which data is received through this port.
Port Peak Send Data Rate	The peak rate at which data is sent through this switch.
Port Peak Receive Data Rate	The peak rate at which data is received through this switch.
Port Send Packet Size	The size of the data block that is sent out through this switch.
Port Receive Packet Size	The size of the data block that is received in through this switch.
Overall Port Packet Size	The total size of the data block that is transferred by the switch
Error Frame Rate	The rate at which error frames are generated by the switch.
Dumped Frame Rate	The rate at which frames are dumped by the switch.
Link Failure Rate	The rate at which failed links are generated by the switch.
WWPN	The World Wide Name identifier for the FC port.
Slot	The number for the blade that the port is on.
Index	The unique number for the port across the whole switch.

Batch reports

Use the information in this section to learn how to generate and view batch reports. These reports are available in the IBM Tivoli Storage Productivity Center > **Reporting > My Reports** > **Batch Reports** node in the navigation tree.

A batch report represents any Tivoli Storage Productivity Center report that you want to run on a regular basis and save its data to a file. Schedule batch reports to

run immediately, run once at a specified date and time, or run repeatedly at a specified date, time, cycle, and day. You can also select the time zone to use when running the report.

You can view the information in a generated batch report file directly or import it into an application of your choice. Batch reports provide a convenient and powerful way for you to save report definitions and schedule when to run those reports.

Creating batch reports

Create a batch report.

To create a batch report follow this procedure:

1. Expand **IBM Tivoli Storage Productivity Center → Reporting → My Reports**.
2. Right-click the **Batch Reports** node and select **Create Batch Report**.

   Note: If the batch report you want to create is similar to an existing one, click the name of the existing report, edit it and save it using a different name.
3. Click the **Report** tab and select the type of report you want to create and move it from **Available** to **Current Selections** by clicking  . See "Selecting batch report options" on page 701 for details of the tab contents.
4. Click the **Selection** tab and refine the criteria for selecting and displaying report data. Some options might not be available for all reports. Use the **Selection** page to:
   - Select a profile for controlling which data is collected. This option is not available for all reports. See "Using profiles" on page 218 for information about profiles.
   - Click **Selection** to specify the resources (for example, computers) that will be displayed in the report. This option is not available for all reports. For details about selecting resources, see "Selecting resources to appear on a report" on page 488.
   - Click **Filters** to apply filters the data that you want to display (for example, show only the machines that contain over 80% used space). For details about applying filters, see "Filtering the resources that appear on a report" on page 491.
   - Select the columns you want to include in the report and move them from **Available Columns** to **Included Columns** by clicking  .

   Note:
   - The values you set in **Available Columns** and **Included Columns** are not used for the following types of batch reports: History CSV File, PDF Chart, and HTML Chart.
   - The value you set in the **Summation Level** field is not used for the following types of batch reports: History CSV File, PDF Chart, and HTML Chart.
5. Click the **Options** tab of the Batch Report job window, specify the following:
   - The machine on which the report file will be saved
   - The report format (for example, HTML)
   - Whether to run a script when the report process is complete
   - The format for the name of the batch report

   See "Selecting batch report options" on page 701 for details of the tab contents.

6. In the **When to Run** tab of the Batch Report job window, specify when and how often to run the batch report. See "Scheduling a batch report" on page 703 for details of the tab contents.
7. In the **Alert** tab of the Batch Report job window, define the conditions that will trigger an alert and the method of notification (for example, e-mail). See "Defining alerts for batch reports" on page 704 for details of the tab contents.
8. Click **File → Save** and type a name for the batch report. The batch report will be saved using your user ID as a prefix. For example, if you are logged onto Tivoli Storage Productivity Center as "Smith," the name of the batch report will be Smith.*report_name*.

Selecting batch report options

Define on what agent computer the batch report will run and the format of the file to which it will be saved.

Use the **Options** page to define where the batch report will run and information about the output file.

To define an Agent computer and output file specification:

1. In the **Agent Computer Specification** section, choose the agent machine where the batch report will run and where the output file will be created. The report file will be saved on the agent machine you choose.

   Note: If the agent you specify in this field is removed from the computer before the batch report is run, an error occurs and the batch report is not generated. If this occurs, edit the definition of the batch report to specify an existing agent and then run the report again.
2. In the **Report Type Specification** section, select the type of batch report you want to generate. You will be limited to one type of output per batch report job and can select from the following output formats:
   - **Formatted File.** Select this option to create a formatted text file of the report.
   - **HTML File**. Select this option to create an HTML version of the report. After the batch report is run, you can view the resulting HTML file in the reports subdirectory within the IBM Tivoli Storage Productivity Center agent install directory.

     If you upgraded from Tivoli Storage Productivity Center v2.x to v3.1, the default install directory is:
     - For Windows: `\Program files\Tivoli\ep\subagents\TPC\Data\log\computername\reports`
     - For UNIX and Linux: `/opt/tivoli/ep/subagents/TPC/Data/log/computername/reports`

     If Tivoli Storage Productivity Center v3.1 is a new installation, the default install directory is:
     - For Windows `\Program Files\IBM\TPC\ca\subagents\TPC\Data\log\computername\reports`
     - For UNIX and Linux: `/opt/IBM/TPC/ca/subagents/TPC/Data/log/computername/reports`
   - **History CSV File**. For reports with historical trending, you can export the historical data to a CSV file. The CSV file will have the following format: <*object name, date, amount, date, amount*,> and so on, where <*date, amount*> corresponds to each data point on the history chart.

- **Include Headers. Click this option to include headers in the generated file.** The column headings that appear in the header reflect the type of report from which you are exporting history data.
- **Historical data list field.** Depending on the type of report you are generating, select the type of history file for the report, such as File Count, Free Space, Free Space %, Used, Space, Used Space%, etc. from the list field. See "Trending with historical data" on page 483 for more information on the types of history files.

Note: If the CSV file type is associated with a spreadsheet program (such as Excel) in your environment, that program might display data in the wrong columns for a CSV file that: 1) was exported from Tivoli Storage Productivity Center, 2) contains commas, and 3) was opened/formatted in the spreadsheet program by double clicking it. To successfully view data in a CSV file that was exported from Tivoli Storage Productivity Center, we recommend you: 1) start the spreadsheet program, 2) open the CSV file using the spreadsheet program's open command.

- **PDF Chart.** Select this option if you want to create a chart for the batch report in PDF format. Both history and pie/bar charts are supported. This option will be available for reports that support charting only and for agents that run on Windows machines--for other reports and non-Windows agents this option will be grayed out.

  When you select a graphing button, the **Agent Computer Specification** box will show agents at the IBM Tivoli Storage Productivity Center for Data v.2.1 or above version level only. If an unacceptable agent has been selected, an error message will be displayed.

  After the batch report is run, you can view the resulting PDF chart file in the reports subdirectory within the Tivoli Storage Productivity Center agent install directory.

  If you upgraded from Tivoli Storage Productivity Center v2.x to v3.1, the default install directory is:

  - For Windows: `\Program files\Tivoli\ep\subagents\TPC\Data\log\computername\reports`
  - For UNIX and Linux: `/opt/tivoli/ep/subagents/TPC/Data/log/computername/reports`

  If Tivoli Storage Productivity Center v3.1 is a new installation, the default install directory is:

  - For Windows `\Program Files\IBM\TPC\ca\subagents\TPC\Data\log\computername\reports`
  - For UNIX and Linux: `/opt/IBM/TPC/ca/subagents/TPC/Data/log/computername/reports`
- **HTML Chart.** Select this option if you want to create a chart for the batch report in HTML format. Both history and pie/bar charts are supported. This option will be available for reports that support charting only and for agents that run on Windows machines--for other reports and non-Windows agents this option will be grayed out.

  When you select a graphing button, the **Agent Computer Specification** box will show agents at the IBM Tivoli Storage Productivity Center for Data v.2.1 or above version level only. If an unacceptable agent has been selected, an error message will be displayed.

  After the batch report is run, you can view the resulting PDF chart file in the reports subdirectory within the Tivoli Storage Productivity Center agent install directory.

If you upgraded from Tivoli Storage Productivity Center v2.x to v3.1, the default install directory is:

- For Windows: `\Program files\Tivoli\ep\subagents\TPC\Data\log\computername\reports`
- For UNIX and Linux: `/opt/tivoli/ep/subagents/TPC/Data/log/computername/reports`

If Tivoli Storage Productivity Center v3.1 is a new installation, the default install directory is:

- For Windows `\Program Files\IBM\TPC\ca\subagents\TPC\Data\log\computername\reports`
- For UNIX and Linux: `/opt/IBM/TPC/ca/subagents/TPC/Data/log/computername/reports`

- **Chart type list field.** For reports with charts, the graphic representation of a chart is generated by selecting one of a report's charts.
- **Customize this chart.** Click this button to use the Chart Customization window to customize the chart that is generated for the batch report, which includes the ability to select a chart type (such as pie, bar, stacked bar, stacked area, etc.) and chart orientation (such as vertical or horizontal). If there is already a chart customization for the corresponding online report, it will be used for this batch report. However, you can override that customization by using this function.

  The button is active only when PDF Chart or HTML Chart is selected. See "Customizing charts" on page 479 for more information on how to customize a batch report.

3. In the **Classic Column Names Specification** section, check **Use Classic Column Names** to use the original names of report columns in the output for batch reports. If you do not select this option, the new names of columns are used. For your reference, the online help displays the original names of those columns within parentheses ( ) next to their new names.

   Note: The titles of some reports have been updated. These new titles are always used for batch reports
4. In the **Script** section, check **Run Script** to run a script when the batch report is completed. In the **Script Name** field, enter the name of the script you want to run. For example, when a batch report completes, you may want to run a script to process the generated file. Here, there is only one parameter that is passed to the script: the `<path>` to the generated batch report.
5. In the **Output File Specification** section, define the output file naming format for the batch report log. You can select the following replaceable parameters:
   - report creator
   - name
   - title
   - run number
   - run date
   - run time
   - Unique number

Scheduling a batch report

Determine when and how often to run a batch report.

Use the **When to Run** page to specify when and how often to run the batch reporting job.

You have the following options for determining when to generate a report:

- Generate the report immediately
- Generate the report once at a specified time and date
- Generate the report repeatedly according to a schedule you specify

Select the option you want and then edit the time and date information accordingly. You can also use the fields in the How to handle time zones section to indicate the time zone that IBM Tivoli Storage Productivity Center should use when scheduling and running the batch report. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run the batch report at:

- 9:00 AM local time of the agent running the batch report
- 9:00 AM of the time zone where the server resides
- 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the batch report will run at 7:00 AM on an agent located in the PST time zone.

Local time in each timezone. Select this option to use the time zone of the location where the agent running the batch report is located.

Same Global time across all timezones. Use the options in this section to apply the same global time across all the time zones where the probe is being run.

- **Use the timezone that the server runs in.** Select this option to use the time zone of the location where the Tivoli Storage Productivity Center server resides.
- **Use this timezone.** Select this option to indicate a specific time zone for when to run the job. When you select this option, the list box to the right of the field will activate and enable you to select from a wide range of global time zones.

Defining alerts for batch reports

Define the alert that will be triggered if the batch report fails to run.

Use the **Alert** tab to define an alert that will be triggered if the report does not run on schedule.

To define alert values for a batch report, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center → Reporting -> My Reports** and right-click on **Batch Reports**.
2. Select **Create Batch Reports**. The Create Batch Reports window opens.
3. For information on creating the batch report, see "Batch reports" on page 699.
4. Click the **Alerts** tab.
5. Define the conditions that will trigger the alert. The default setting is **Report Failed**. No other alerting conditions are currently available.
6. Select alert mechanisms in addition to writing an error to the error log:
   - If you want to use SNMP traps, IBM Tivoli Enterprise Console events, or e-mail as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services → Configuration → Alert Dispositions**
   - If you want a script to be run after an event has been triggered, you must create the script before you create the alert.
7. Save the file:
   a. Click **File → Save**.

b. The **Save As** window opens. Enter a name for the batch report and click **OK**.

Editing alert values for a batch report

1. Expand **IBM Tivoli Storage Productivity Center → My Reports → Batch Reports**. The list of existing batch reports is displayed.
2. Click the batch report. The **Edit Batch Report** window opens.
3. In the **Alerts** tab, make any necessary changes.
4. Click **File → Save**.

Related topics:

"Batch reports" on page 699

Creating batch reports for storage subsystem or switch performance

This topic describes how to create a batch report for storage subsystem performance or switch performance.

To create a batch report for storage subsystem performance or switch performance, follow these steps:

1. Expand **IBM Tivoli Storage Productivity Center > Reporting > My Reports**.
2. Right-click the **Batch Reports** node and select **Create Batch Report**.
3. Click the **Report** tab and select the type of report you want to create and move it from **Available** to **Current Selections** by clicking the right arrow button (>>). Select an item from **Storage Subsystem Performance** or **Switch Performance**.
4. Click the **Selection** tab. The Report Filter Specifications window displays.

   Selection button

   This determines the resources upon which you want to report, such as the storage subsystems, the fibre-channel ports, and so forth. The Select Resources window displays. The type of resources displayed for selection depends on the type of report you previously selected on the **Report** tab. For example, if you selected a By Storage Subsystem report, all displayed resources will be storage subsystems. If you clear the box for a resource on the Select Resources window, a pushpin icon displays on the Selection button.

   Filter button

   This further defines the data that you want to include in the report. Filters enable you to apply general rules to the report. When you click this button, the Edit Filter window displays.

   For example, to select only data associated with subsystems that have a name ending with ″123–IBM″, select the ″Subsystem″ column, the ″LIKE″ operator, and type ″*123–IBM″ (without the quotation marks) in the Value 1 field. To select only data which has a Read I/O Rate greater than 0, select the ″Read I/O Rate (overall)″ column, the ″>″ operator, and type ″0″ (without the quotation marks) in the Value 1 field.

   If a filter is defined for this report, a pushpin icon displays on the **Filter** button.

   Time range selectors

   All performance reports display historical performance data. Use the time range selectors to specify the exact time range that you want the report to cover. Click one of the following radio buttons to select the type of historical report you want:

Display latest performance data

This report contains only the latest available performance data for the selected resources. Note that even the latest data available might still be somewhat old, depending on whether data was recently collected for the resources. Including the time column in the report will help you determine how old the latest available data actually is. Regardless of the age of the data, clicking this button ensures that only a single row (representing a single point in time) for each selected resource is included in the report. If no data has ever been collected for a particular resource, then that resource is not included in the report. If no data has ever been collected for any resource, then the resulting report will be empty.

Display historic performance data using absolute time

This report contains only data that falls within a specific time range. Select the time range using the **From** and **To** selectors. The specified from and to values are interpreted as inclusive. The resulting report contains all the performance data that is available within the specified time range, for each selected resource. This means that the data can consist of potentially many rows per resource, each with a unique timestamp. If no data has been collected for a particular resource during the specified time range, that resource is not included in the report. If no data has been collected for any resource during the specified time range, then the resulting report will be empty.

To determine which historical performance records fall within the specified time range, the **beginning timestamp** of the collection intervals is used for comparison. If an historical record contains performance data for an interval of time from 1:00 PM to 1:05 PM, that record would be included in the report for a time range of 1:00 PM to 8:00 PM. However, that same record would not be included if a time range of 1:02 PM to 8:02 PM was specified, because it is the beginning of the time period (for example, 1:00 PM in this example) that is used to determine inclusion in the specified time interval.

Specifying an absolute time range is not suitable for batch reports that are run on a repeating schedule, because the specified time range is fixed and will not automatically be updated based on the most current date and time. The result will be that the same report is generated every time that the batch report is run. Instead, use a relative time range for this type of report.

Display historic performance data using relative time

This report contains only data that falls within a specific time range extending backward from the current date and time. Specify the number of days in the desired time range, using the **days ago until now** field. The report will be exactly the same as a report using absolute time, with the **To** selector set to the current date and time, and the **From** selector set to the current date and time minus the number of **days ago until now**. However, when regenerating this report periodically, as the current date and time advances, the data included in the report

will advance as well. As a result, this is the suitable option to select when defining batch reports that are run on a repeated schedule.

Note that the same criteria applies as if determining which historical performance records fall within the specified time range. Only the beginning timestamp of the collection intervals is used for comparison. This means that data for the current hour or the current day is not included in the report when selecting hourly or daily summation data, because data for the current hour or current day is only saved in the database when the hour or day is complete. As an alternative, consider using the first radio button, **Display latest performance data**, to see a report showing the most recent daily summation data.

Summation Level

Use this selector to select which type of performance data records to include in the report. The Tivoli Storage Productivity Center Performance Manager collects and stores three different types of performance data:

By sample

Represents the most detailed data. There is usually one sample record per resource for every **n** minutes, where **n** is the performance monitor interval length that was specified when the monitor was defined. Each sample represents the average performance of the resource over those **n** minutes. A sample record with a timestamp of **t** will be saved at time **t+n**. In other words, the beginning of each **n** minute time interval is used for display and for time range matching.

Hourly

Represents an hourly average of the performance of the specified resources. There will be one hourly record written every hour, per resource. Each such record will be written at or shortly after each hour mark on the (device) clock. For example, a record with a timestamp of 12:00 PM and interval length of 1 hour will be written at 1:00 PM.

Daily Represents a daily average of the performance of the resources. There will be one daily record written every day, per resource. Each daily record will be written at or shortly after midnight, according to the device clock. For example, a record with a timestamp of April 21, 12:00 AM and interval length of 24 hours will be written on April 22, 12:00 AM.

Note that hourly and daily records are usually retained in the database for longer periods of time, compared to By Sample records. As a result, if a generated report indicates that sample data is no longer available for a particular time range in the past, it is possible that the sample data has already been purged, in which case you can try generating the report using hourly or daily summation levels.

Be aware that when selecting a time range using the **Display historic performance data using relative time** radio button, it is necessary to select the length of the range in days. This can lead to unexpected results when also selecting the daily summation level. Remember that a daily record is not written until the day is complete, which means that at any given time during the day, the current day's data does not yet

exist in the database. Therefore, selecting the daily summation level as well as a relative time range of n **days ago until now**, will usually result in n-1 records being returned per resource, assuming that at least n-1 days worth of data has been previously collected for those resources. Specifying a relative time range of 1 **days ago until now** will result in an empty report.

Available Columns, Included Columns
Use these items to specify which columns to include and which columns to exclude from the generated report.

Available Columns
Displays the columns that you can include in the generated report that are not already included. If a column appears in the **Included Columns** list box, it will not appear in the **Available Columns** list box.

Included Columns
Displays the columns that will appear in the generated report. By default, all columns for a report are listed in this list box and will be displayed in the generated report. You can highlight the columns and use the up and down arrows to rearrange the order of the columns in the report. The order of the items in this list box determines the order in which the columns will appear in the generated report. The grayed-out column names shown at the top of the list box are the frozen columns, which cannot be selected, removed, or reordered for this type of report.

If a column appears in the **Available Columns** list box, it will not appear in the **Included Columns** list box.

To exclude a column from a report, click the name of the column in the **Included Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the left arrow button (<<). The selected columns are removed from the **Included Columns** list box, and will appear in the **Available Columns** list box.

To include a column in a report, click the name of the column in the Available Columns list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the right arrow button (>>). The selected columns will be removed from the **Available Columns** list box, and will appear in the **Included Columns** list box.

To rearrange the order of columns that will appear in the generated report, click the name of a column in the **Available Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the up arrow button (^) or the down arrow button (v) to move the selected columns up or down in the list, relative to the other columns.

5. Click the **Options** tab of the Batch Report job window, specify the following:
   - The machine on which the report file will be saved
   - The report format (for example, HTML)
   - Whether to run a script when the report process is complete
   - The format for the name of the batch report
6. In the **When to Run** tab of the Batch Report job window, specify when and how often to run the batch report.

7. In the **Alert** tab of the Batch Report job window, define the conditions that will trigger an alert and the method of notification.
8. Click **File > Save** and type a name for the batch report. The batch report will be saved using your user ID as a prefix. For example, if you are logged onto Tivoli Storage Productivity Center as ″Smith,″ the name of the batch report will be Smith.*report_name.*

Editing batch reports

Edit existing batch report definitions.

To edit a batch-report definition, follow this procedure:

1. Expand **IBM Tivoli Storage Productivity Center → Reporting → My Reports → Batch Reports**.
2. Click the report name, make your changes and then click **File → Save**.

Displaying a list of batch job reports

List the currently defined batch report jobs in the navigation tree.

To display a list of batch job reports, follow this procedure:

1. Expand **IBM Tivoli Storage Productivity Center → Reporting → My Reports → Batch Reports**.
2. Click the name of a batch report. A list of jobs is displayed
3. Click a job.

Displaying batch report log files

Display the log file for a batch report job.

To display the log file for a batch-report job, follow this procedure:

1. Expand **IBM Tivoli Storage Productivity Center → Reporting → My Reports → Batch Reports**.
2. Expand a batch report node and click a job.
3. Click the magnify button.

Displaying batch reports from the saved directory

Access the output file generated by a batch report job. This output file will contain the data for the report you selected to run as a batch report.

After you run a batch report, output files are saved in the agent's install directory on the machine you specified in the **Options** tab of the Create Batch Report window. The default install directories for agents are the following:

- Windows:\Program Files\IBM\TPC\ca\subagents\TPC\Data\log\computername\reports

  If you upgraded to v3.1 or higher from a previous version of IBM Tivoli Storage Productivity Center, the default location for batch reports is:\Program Files\IBM\TPC\ep\subagents\TPC\Data\log\computername\reports If you defined the output file to be formatted as a .txt file, you can view the report with Note Pad. If you defined the output file as .xls, you can import the data to MS Excel, by selecting **Data ► Get External Data ► Import Data**
- UNIX, Linux: /opt/IBM/TPC/ca/subagents/TPC/Data/log/computername/reports

If you upgraded to v3.1 or higher from a previous version of Tivoli Storage Productivity Center, the default location for batch reports is:/opt/tivoli/ep/subagents/TPC/Data/log/computername/reports

Viewing CSV batch reports with spreadsheet applications

View CSV-formatted batch reports in third-party spreadsheet applications.

If you are using a spreadsheet program (such as Excel) to view a CSV-type batch report file, that program might display data in the wrong columns either because the data was exported from Data Manager, the data contains commas, or the file was opened or formatted in the spreadsheet program by double-clicking it.

If data is displaying in the wrong columns, open the CSV file using the spreadsheet program's open command.

Exporting and importing settings that you define

This topic describes how to export batch report definitions, constraints, groups, and profiles.

You can export settings that you define from one instance of IBM Tivoli Storage Productivity Center and import the settings into another instance. To export and import your definitions, complete the following steps:

1. On the system running the server component, change to the following directory:

On Linux and UNIX	`/opt/IBM/TPC/Data/server/tools/`
On Windows	`C:\Program files\IBM\TPC\Data\server\tools\`

2. To export the definitions from a command prompt, type the following command and press **Enter**:

   ```
   export -e <file>
   ```

   where <file> is the fully qualified name of the file where the definition information is written.
3. On the system to which you want to migrate the definitions, change to the **tools** directory.
4. To import the definitions, type the following command and press Enter:

   ```
   export -i <file>
   ```

   where <file> is the fully qualified name of the file that you created when exporting.

User ID (saved) reports

Use the information in this section to learn how to generate, save, and view user ID reports. You can create a user ID reports based on a system report or any of the online reports available in Data Manager, Data Manager for Databases, Disk Manager, or Fabric Manager. Once saved, user ID reports are available in the **IBM Tivoli Storage Productivity Center** > **Reporting >My Reports** > **user ID name** node in the navigation tree.

Creating a user-defined report

Create a custom report based on a system report or any of the online reports available in Data Manager, Data Manager for Databases, Disk Manager, or Fabric Manager.

To create a custom report, follow this procedure:

1. Specify new report criteria by doing one of the following:
   - Access a system report and change its format according to your organization's needs.
   - Create a new report using the reporting component in any of the IBM Tivoli Storage Productivity Center managers. For example, **Data Manager → Reporting → Usage → Access Load → Access Time → By Filesystem**.
2. Click **File → Save** and a type a name for the report. The report is saved in the **Tivoli Storage Productivity Center > Reporting > My Reports > *user name's*** node, where ***user name's*** is your login name. For example, if you are logged in as "Smith", your custom reports appear under the **Smith's Reports** node.

Displaying user-defined reports

Display a custom report saved under your user ID.

To display a list of custom reports, expand the **IBM Tivoli Storage Productivity Center > Reporting > My Reports > <*your user ID*>** node. Click the name of the report you want to view.

Editing user-defined reports

Change the settings for an existing user ID report and save those changes.

To edit a custom report definition, follow this procedure:

1. Expand **My Reports → *user name's* Reports**.
2. Click the report name, make your changes and then click **File → Save**.

Rollup reports

Use rollup reports to view the storage information collected by the master IBM Tivoli Storage Productivity Center server from its associated subordinate servers. Rollup reports enable you to have a network-wide perspective of storage usage in an environment where multiple Tivoli Storage Productivity Center servers are deployed and collecting storage metrics.

Before you can generate rollup reports, you must perform the following:

- Associate subordinate servers with the master server on the **Administrative Services → Data Sources → Tivoli Storage Productivity Center Servers window**.
- Run discoveries, probes, and scans from the subordinate servers to gather information about their locally-monitored storage entities.
- Run discoveries, probes, and scans from the master server to gather information about its locally-monitored storage entities. Note that information about the storage entities monitored by a master server are automatically shown in rollup reports.
- Run TPC Server probes from the master against the subordinate servers to collect their storage information for the rollup reports. Use the Tivoli Storage Productivity Center **Monitoring → TPC Server Probes** node to define and schedule probe jobs.

Note:

- If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.
- After upgrading to Tivoli Storage Productivity Center V4.1 or later from a previous version of the application, 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**.

Tivoli Storage Productivity Center includes different rollup reporting categories that provide you with the flexibility to view data about your storage resources according to the needs of your environment. The following rollup reporting categories are available:

Generating rollup reports

Learn how to generate and view rollup reports.

1. Expand the navigation tree to display the type of report you want to generate. For example, if you want to generate a database capacity report, expand **IBM Tivoli Storage Productivity Center → Reporting → Rollup Reports → Database Capacity**.
2. Highlight a By Clause to determine how you want to view a report. For example, to view rollup capacity information according to computer group, click **Tivoli Storage Productivity Center → Reporting → Rollup Reports → All DBMSs → By Computer Group**. The **Selection** page displays.
3. Use the **Selection** page to determine what columns appear within a report, what storage assets to report upon (Capacity and Database Capacity reports only), and how to filter that values that appear report columns.
4. (Capacity and Database Capacity reports only) Click **Selection...** to select the objects that you want to report upon from the Select Resources window. See "Selecting resources to appear on a report" on page 488 for information.
5. Click **Filter...** to further filter the objects that appear in a report. Filters enable you to apply general rules to the report based on the rows in that report. See "Filtering the resources that appear on a report" on page 491 for more information on filtering.
6. Use the **Available Columns** and **Included Columns** list boxes to determine what columns are displayed in a generated report.
7. Click **Generate Report**. A new tab will be added to the tab dialog representing the report that you generated.
8. Click the new tab to view the report. Reports are tabular in format and composed of rows and columns. You can scroll the report up and down and to the right and left to view the entire report. Use the **View** menu to hide/show the navigation tree to increase the viewable area of the report or drag the divider bar in the middle of the screen back and forth to reallocate the amount of space that is given each pane.

Rollup reports - fragmentation

Fragmentation occurs when related storage entities are being monitored by different subordinate IBM Tivoli Storage Productivity Center servers.

The following example shows when fragmentation can occur using subordinate servers and rollup reports:

- **Example:** Host A is assigned a volume from storage subsystem 1. Host A is being monitored only by one Tivoli Storage Productivity Center subordinate

server (Server 1) and storage subsystem 1 is being monitored by another subordinate server (Server 2). Server 1 and Server 2 are subordinate servers to the master server (Server 3).

Reason for fragmentation: In this example, the host and subsystem are fragmented because they are not being monitored by the same Tivoli Storage Productivity Center server. This will cause the **Correlated Volume Space** value in the rollup reports for storage subsystems and volumes to be incorrect.

To prevent fragmentation when using master and subordinate server relationships to generate rollup reports, we recommend the following:

- Do not monitor a storage subsystem and the hosts to which its volumes are assigned with different subordinate servers that are associated with the same master server.
- Do not monitor a fabric and its related switches with different subordinate servers that are associated with the same master server.
- If an agent machine is using storage from a subsystem, then both the agent and subsystem should be monitored by the same subordinate server.
- Configure all the subordinate servers to discover all the storage subsystems within your environment. Note that if a subordinate server is monitoring a subsystem with an agent other than the agent that is using storage on that subsystem, the value for the **Correlated Volume Space** column in rollup reports will be incorrect. For example:
  - Host A has storage from Subsystem 1
  - Subordinate server B is Monitoring Host A and Subsystem 1
  - Subordinate server C is Monitoring Subsystem 1 and probed the subsystem more recently than server B.
  - The master server is monitoring subordinate servers B and server C.

  The master server will report the storage subsystem data from server C. Because the master server does not know about Host A, the value for the **Correlated Volume Space** column will be incorrect.

Rollup Reports - Asset

Use these reports to view detailed statistics about agents (IBM Tivoli Storage Productivity Center), computers, storage subsystems, disk and volume groups, disks, filesystems, logical volumes, volumes, and fabrics that are monitored by subordinate servers.

The information shown in these reports is gathered by the master server when it runs Tivoli Storage Productivity Center server probes against associated subordinate servers. If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

Rollup Reports - Asset - Agents

Use this report to view information about Data agents and Device agents that are associated with subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Agent: Shows agents sorted according to the name of the machine on which they are installed.

- By OS Type: Shows agents sorted according to the operating system under which they run. Click the magnifying glass icon next to an operating system type to view more detailed information about each of the agents running under that operating system.
- By TPC Server: Shows agents sorted according to the subordinate server that manages them. Click the magnifying glass icon next to subordinate server to view more detailed information about each of the agents associated with that server.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Agent

Agent Name
Name of the machine on which an agent is installed.

Data/SR Agent Version
Tivoli Storage Productivity Center version number of a Data agent (for example, 3.3.0).

Fabric Agent Version
Tivoli Storage Productivity Center version number of a Fabric agent (for example, 3.3.0.86).

OS Type
Operating system on the machine where an agent is installed.

Data/SR Agent Status
Status of a Data agent. Possible statuses include Up and Down. A down agent represents an agent that is still running but with which the server is not communicating.

Fabric Agent Status
Status of a Fabric agent. Possible statuses include Up and Down. A down agent represents an agent that is still running but with which the server is not communicating.

Data/SR Agent Port
Port number on which a Data agent listens for requests. Port 9549 is registered for use by IBM.

Fabric Agent Port
Port number on which a Fabric agent listens for requests. Port 9550 is registered for use by IBM.

Host Address
Network name or IP address of an agent as seen by subordinate server.

Time Zone
Time zone where the agent is installed (for example, CDT).

CPU Architecture
Instruction set architecture on the machine where an agent is running.

By OS Type

OS Type
Operating system on the machine where an agent is installed.

Number of Data Agents/SR Agents
Total number of Data agents running on the operating system.

Number of Data Agents/SR Agents (Status: Up)
Number of Data agents running on the operating system that have a status of Up.

Number of Data Agents/SR Agents (Status: Down)
Number of Data agents running on the operating system that have a status of Down. A down agent represents an agent that is still running but with which the server is not communicating.

Number of Data Agents/SR Agents (Status: Defected)
Number of Data agents running on the operating system that have been reassigned to another server.

Number of Data Agents/SR Agents (Status: Upgrading)
Number of Data agents running on the operating system that are currently in the process of upgrading to the same version as the subordinate server.

Number of Data Agents/SR Agents (Status: Unreachable)
Number of Data agents running on the operating system that are unreachable. This status indicates that an agent that is unreachable and not responding to the server. This might occur when the agent is not running but did not inform the server that it was shutting down, or when a network problem has occurred.

Number of Data Agents/SR Agents (Status: Upgrade Required)
Number of Data agents running on the operating system that require an upgrade to the same version as the subordinate server.

Number of Data Agents/SR Agents (Status: Deleted)
Number of Data agents running on the operating system that have a status of Deleted.

Number of Fabric Agents
Total number of Fabric agents running on the operating system.

Number of Fabric Agents (Status: Up)
Number of Fabric agents running on the operating system that have a status of Up.

Number of Fabric Agents (Status: Down)
Number of Fabric agents running on the operating system that have a status of Down. A down agent represents an agent that is still running but with which the server is not communicating.

By TPC Server

TPC Server
Name of the TPC server on which an agent is associated.

Number of Data Agents
Total number of Data agents associated with the TPC server.

Number of Data Agents/SR Agents (Status: Up)
Number of Data agents associated with the TPC server that have a status of Up.

Number of Data Agents/SR Agents (Status: Down)
Number of Data agents associated with the TPC server that have a status of Down. A down agent represents an agent that is still running but with which the server is not communicating.

Number of Data Agents/SR Agents (Status: Defected)
Number of Data agents associated with the TPC server that have been reassigned to another server.

Number of Data Agents/SR Agents (Status: Upgrading)
Number of Data agents associated with the TPC server that are currently in the process of upgrading to the same version as the subordinate server.

Number of Data Agents/SR Agents (Status: Unreachable)
Number of Data agents associated with the TPC server that are unreachable. This status indicates that an agent that is unreachable and not responding to the server. This might occur when the agent is not running but did not inform the server that it was shutting down, or when a network problem has occurred.

Number of Data Agents/SR Agents (Status: Upgrade Required)
Number of Data agents associated with the TPC server that require an upgrade to the same version as the subordinate server.

Number of Data Agents/SR Agents (Status: Deleted)
Number of Data agents associated with the TPC server that have a status of Deleted.

Number of Fabric Agents
Total number of Fabric agents associated with the TPC server.

Number of Fabric Agents (Status: Up)
Number of Fabric agents associated with the TPC server that have a status of Up.

Number of Fabric Agents (Status: Down)
Number of Fabric agents associated with the TPC server that have a status of Down.

Rollup Reports - Asset - Computers

Use this report to view information about the hosts and computers (including NetApp and NetWare) that are monitored by subordinate servers in your environment. This includes VMWare-related entities such as hypervisors and virtual machines.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

You can generate and sort the data within this report using the following by clauses:

- By Computer: Shows computers monitored by subordinate servers. Sorted by computer name.
- By TPC Server: Shows subordinate servers that are monitoring computers. Sorted by subordinate server name. Click the magnifying glass icon next to subordinate server to view more detailed information about each of the computers associated with that server.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Computer

Computer
Name of a computer monitored by a subordinate server.

OS Type
Operating system running on a computer.

OS Version
Version of the operating system running on a computer.

RAM Amount of RAM installed on a computer.

Disk Space (Capacity)
Total storage capacity for a computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Note: If the computer is an IBM Tivoli Storage SAN File Systems client, then this value will also include the capacity of any storage subsystem volumes visible to this computers that are owned by the SAN File System.

Number of Disks
Total number of disks for a computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value in this column:

- Computer disks on a computer that has not been probed by a Data agent.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Processor Count
Number of CPUs running on a computer.

IP Address
IP address of a computer.

Serial Number
Serial number of a computer.

Model Model number or name of a computer.

Manufacturer
Manufacturer of a computer (for example, IBM).

Time Zone
Time zone where the computer is located.

Available File System Space (Filesystem Free Space)
Total amount unique file system storage space that is not used or available to the computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Network Address
Network name of a computer.

Processor Type
Type of processor running on a computer (for example, GenuineIntel: x86 Family 15).

Processor Speed
Speed at which a computer's processor is running (in MHz).

CPU Arch
Architecture of the processor (for example, IA32).

Swap Space
Amount of swap space (virtual memory) defined on a computer.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring computers in your environment.

# of Computers
Number of computers monitored by the subordinate server.

Disk Space (Capacity)
Total storage capacity of the computers monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Note: If a computer is an IBM Tivoli Storage SAN File Systems client, then this value will also include the capacity of any storage subsystem volumes visible to this computers that are owned by the SAN File System.

Available File System Space (Filesystem Free Space)
Total amount unique file system storage space that is not used or available to the computers that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

Number of Disks
Number of disks associated with the computers monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value in this column:

- Computer disks on a computer that has not been probed by the Data agent.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Rollup Reports - Asset - Storage Subsystems

Use this report to view information about the storage subsystems that are monitored by subordinate servers in your environment.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

You can generate and sort the data within this report using the following by clauses:

- By Storage Subsystem: Sorted according to storage subsystems name. The names of storage subsystems monitored by subordinate servers appear in the first column and remain displayed when scrolling through the other columns on the report.
- By Disk Space: Sorted according to the storage capacity of storage subsystems. The names of storage subsystems monitored by subordinate servers and their disk capacity appear in the first two columns and remain displayed when scrolling through the other columns on the report.
- By TPC Server: Shows storage subsystems aggregated at subordinate server level. The names of subordinate servers that are monitoring storage subsystems appear in the first column and remain displayed when scrolling through the other columns on the report. Click the magnifying glass icon next to subordinate server to view more detailed information about each of the storage subsystems monitored by that server.

Note:

- Storage subsystems that have been discovered but not probed will not appear in these reports.
- This report includes information about storage subsystem Volumes. For SAN Volume Controller subsystems, Tivoli Storage Productivity Center uses the term ″Volume″ to represent ″virtual disks″ for those subsystems.

By Storage Subsystem, and By Disk Space

Name Name or alias of a storage subsystem that is monitored by a subordinate server or the master server.

Disk Space (Capacity)
Total raw (unformatted) disk capacity of a storage subsystem. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Capacity of spare disks identified on IBM Tivoli Storage Enterprise Storage Server, IBM DS8000, and IBM DS6000 storage subsystems.
- Capacity of storage subsystem disks that become missing after a storage subsystem probe.

Note: For Hitachi Data Systems storage subsystems, this is the capacity of all the Volumes in the storage subsystem.

Available Disk Space (Unallocated Space)
Total unformatted disk free space of the storage subsystem. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- RAID overhead on storage subsystems that are identified by Tivoli Storage Productivity Center.
- External back-end storage for TagmaStore subsystems.

- Storage space that is consumed by volumes that reside on external disks when Hitachi Data Systems (HDS) storage subsystems are configured as storage virtualizers.

Note: For Hitachi Data Systems storage subsystems, internal disk information is not available. This field will be 0.

Number of Disks
Total number of disks in the storage subsystem. Tivoli Storage Productivity Center does not include disks that become missing after a successful storage subsystem probe in its calculation of the value for this column. For Hitachi Data Systems storage subsystems, internal disk information is not available. This field will be set to 1.

Number of Volumes (LUNs)
Number of storage volumes and virtual storage volumes discovered by a probe. Tivoli Storage Productivity Center does not include storage volumes or virtual storage volumes that become missing after a successful probe in its calculation of the value for this column.

Volume Space (LUN Capacity)
Total amount of unique storage subsystem volume space or virtual storage volume space discovered on monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include storage volumes or virtual storage volumes that become missing after a successful probe in its calculation of the value for this column.

Note: For space-efficient volumes, the Available Space is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage space reported for a storage subsystem using space-efficient volumes. This also applies to other space calculations, such as the calculations for the Storage Subsystem's Consumable Volume Space and FlashCopy Target Volume Space.

Consumable Volume Space (Usable LUN Capacity)
Total amount of unique storage subsystem volume space or virtual storage volume space discovered on monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following storage entities in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a probe
- Storage subsystem volumes or virtual storage volumes that are used as flash copy targets.
- Storage subsystem volumes or virtual storage volumes that are identified as a Business Continuance Volume extended (BCVx) (EMC only).
- Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database (EMC only).

FlashCopy Target Volume Space (Capacity)
Total amount of unique storage subsystem volume space and virtual storage volume space discovered on monitored storage subsystem arrays and has been identified as flash copy target space. Tivoli Storage Productivity Center does not include the following storage entities in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a probe.

- Storage subsystem volumes or virtual storage volumes that are not used as flash copy targets.

Note: Tivoli Storage Enterprise Storage Server subsystems must have at least the following microcode levels: ess800mincodelevel = 2.4.3.56, essf20mincodelevel = 2.3.3.89. IBM Tivoli Storage Productivity Center will not report FlashCopy information if the Tivoli Storage Enterprise Storage Server subsystems do not meet this requirement.

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Total amount of storage subsystem volume space or virtual storage volume space on systems where the systems disks are identified by Tivoli Storage Productivity Center as residing within the monitored storage subsystems in a group. Tivoli Storage Productivity Center does not include the following in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Disks that are assigned to a monitored storage subsystem but are not identified by Tivoli Storage Productivity Center as a disk residing within the storage subsystem. Some reasons why these disks are not identified as residing within the storage subsystem include:
  - Tivoli Storage Productivity Center does not support the identification of disks for some storage subsystem types.
  - The connection of the disk to the storage subsystem is lost and is no longer detectable during a probe.
- Storage subsystem volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers. For example, IBM SAN Volume Controller.

Volume Space Not Correlated (LUN Capacity Not Visible to Data Manager)
Total amount of unique storage subsystem volume space and virtual storage volume space discovered on monitored storage subsystem arrays in a group that can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Disks that reside within a monitored storage subsystem.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers.

Unavailable Disk Space (Overhead)
Amount of storage subsystem volume or virtual storage volume space that is dedicated to redundancy. This value is dependent on the storage subsystem SMI-S provider returning a valid RAID value that can be used to determine the overhead. If the value is zero for a storage subsystem, the overhead cannot be calculated.Tivoli Storage Productivity Center does not include the storage subsystem volumes or virtual storage volumes that become missing after a probe in its calculation of the value for this column.

Formatted Space (Capacity)
Total amount of formatted storage space associated with the storage subsystem that is used or can be used for subsystem storage volumes or virtual storage volumes. This value is dependent on the storage subsystem SMI-S provider returning a formatted value.

For some storage subsystems, this is the total managed space which includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Formatted Space with No Volumes (Formatted Space with No Volumes)
Total amount of formatted storage space associated with the storage subsystem that can be used for subsystem storage volumes or virtual storage volumes. This value is dependent on the storage subsystem SMI-S provider returning a formatted value.

For some subsystems, this value represents the remaining managed space that includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Overall Unavailable Disk Space (Total Overhead)
(IBM Tivoli Storage Enterprise Storage Server, IBM DS6000, and IBM DS8000 storage subsystems only) Total amount of unformatted space within disk groups or array sites.

Available Disk Group or Array Site Space (Undefined Disk Group or Array Site Capacity)
This value does not apply to all storage subsystems. For IBM Tivoli Storage Enterprise Storage Server, IBM DS6000, IBM DS8000 storage subsystems, this is the total raw disk space of any unformatted disk groups or array sites. For other storage subsystems that contain a valid value, this is the total disk space that is not within a storage pool. This column displays **N/A** for the storage subsystems where this value is not applicable. For IBM XIV Storage System, the total available physical (hard) space that is left for pool creation.

Cache Amount of RAM or other cache for a storage subsystem.

Note: This column is blank for SAN Volume Controllers.

Manufacturer
Manufacturer of a storage subsystem.

Model Model number/name of a storage subsystem.

Serial Number
Serial number of a storage subsystem.

Firmware Revision
Microcode level of a storage subsystem.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage subsystem volume space within the storage subsystem that is mapped or assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a probe.

- Storage subsystem volumes that are not mapped or are not assigned to host systems.

Unassigned Volume Space (Subsystem Unassigned LUN Capacity)
Total storage subsystem volume space within a storage subsystem that is not mapped or not assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a probe.
- Storage subsystem volumes that are mapped or assigned to host systems.

Unassigned and Available Formatted Space (Subsystem Total Formatted Free Space) Total storage subsystem space that is not mapped or not assigned to host systems and the total amount of formatted storage space associated with the storage subsystem that can be used for subsystem storage volumes.

This total amount of formatted storage space is dependent on the storage subsystem SMI-S provider returning a formatted value. For some subsystems, this value represents the remaining managed space that might include space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes that become missing after a probe.
- Storage subsystem volumes that are mapped or assigned to host systems.

zOS Volume Space (LUN Capacity)
Total storage subsystem space that is reserved for use on mainframe storage such as used by zOS. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are reserved for use on open system server or non-mainframe storage.

Open System Volume Space (LUN Capacity)
Total storage subsystem space that is reserved for use on open system storage such as used by operating systems like LINUX, AIX, HP/UX, SUN Solaris, and Windows. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are reserved for use on mainframe storage.

Last Probed Timestamp
Date and time when a storage subsystem was last probed by the subordinate or master server.

Physical Disk Space (Capacity)
Total amount of physical disk space discovered on the monitored storage

subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, IBM SAN Volume Controller) that are attached from a storage subsystem.

Backend Volume Space (LUN Capacity)
Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage subsystem virtualizer (for example, IBM SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are back-end disks on a storage subsystem virtualizer that is not monitored.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring storage subsystems in your environment.

Number of Subsystems (Total)
Number of storage subsystem monitored by the subordinate server.

Disk Space (Total Disk Capacity)
Total raw (unformatted) disk capacity of the storage subsystems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Capacity of spare disks identified on IBM Tivoli Storage Enterprise Storage Server, IBM DS8000, and IBM DS6000 storage subsystems.
- Capacity of storage subsystem disks that become missing after a storage subsystem probe.

Note: For Hitachi Data Systems storage subsystems, this is the capacity of all the Volumes in the storage subsystem.

Available Disk Space (Total Disk Unallocated Space)
Total unformatted disk free space of the storage subsystems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- RAID overhead on storage subsystems that are identified by Tivoli Storage Productivity Center.
- External back-end storage for TagmaStore subsystems.
- Storage space that is consumed by volumes that reside on external disks when Hitachi Data Systems (HDS) storage subsystems are configured as storage virtualizers.

Note: For Hitachi Data Systems storage subsystems, internal disk information is not available. This field will be 0.

Number of Disks (Total)
Total number of disks in the storage subsystems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include disks that become missing after a successful storage subsystem probe in its calculation of the value for this column. For Hitachi Data Systems storage subsystems, internal disk information is not available. This field will be set to 1.

Number of Volumes (Total Number of LUNs)
Total number of storage volumes and virtual storage volumes discovered by the subordinate server. Tivoli Storage Productivity Center does not include storage volumes or virtual storage volumes that become missing after a successful probe in its calculation of the value for this column.

Volume Space (Total LUN Capacity)
Total amount of unique storage subsystem volume space or virtual storage volume space discovered on storage subsystem arrays that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include storage volumes or virtual storage volumes that become missing after a successful probe in its calculation of the value for this column.

Note: For space-efficient volumes, the Available Space is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage space reported for a storage subsystem using space-efficient volumes. This also applies to other space calculations, such as the calculations for the Storage Subsystem's Consumable Volume Space and FlashCopy Target Volume Space.

Consumable Volume Space (Total Usable LUN Capacity)
Total amount of unique storage subsystem volume space or virtual storage volume space discovered on storage subsystem arrays that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following storage entities in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a probe
- Storage subsystem volumes or virtual storage volumes that are used as flash copy targets.
- Storage subsystem volumes or virtual storage volumes that are identified as a Business Continuance Volume extended (BCVx) (EMC only).
- Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database (EMC only).

FlashCopy Target Volume Space (Total FlashCopy Target Capacity)
Total amount of unique storage subsystem volume space and virtual storage volume space discovered on storage subsystem arrays that are monitored by the subordinate server and have been identified as flash copy target space. Tivoli Storage Productivity Center does not include the following storage entities in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a probe.
- Storage subsystem volumes or virtual storage volumes that are not used as flash copy targets.

Note: Tivoli Storage Enterprise Storage Server subsystems must have at least the following microcode levels: ess800mincodelevel = 2.4.3.56, essf20mincodelevel = 2.3.3.89. Tivoli Storage Productivity Center will not report FlashCopy information if the Tivoli Storage Enterprise Storage Server subsystems do not meet this requirement.

Unavailable Disk Space (Overhead)
Amount of storage subsystem volume or virtual storage volume space that is dedicated to redundancy and is monitored by the subordinate server. This value is dependent on the storage subsystem SMI-S provider returning a valid RAID value that can be used to determine the overhead. If the value is zero for a storage subsystem, the overhead cannot be calculated.Tivoli Storage Productivity Center does not include the storage subsystem volumes or virtual storage volumes that become missing after a probe in its calculation of the value for this column.

Correlated Volume Space (Total LUN Capacity Visible to Data Manager)
Total amount of storage subsystem volume space or virtual storage volume space on systems where the systems disks are identified by Tivoli Storage Productivity Center as residing within the storage subsystems that are monitored by the subordinate server. The value in this column reflects the subsystems that are monitored by the subordinate server.Tivoli Storage Productivity Center does not include the following in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Disks that are assigned to a monitored storage subsystem but are not identified by Tivoli Storage Productivity Center as a disk residing within the storage subsystem. Some reasons why these disks are not identified as residing within the storage subsystem include:
  - Tivoli Storage Productivity Center does not support the identification of disks for some storage subsystem types.
  - The connection of the disk to the storage subsystem is lost and is no longer detectable during a probe.
- Storage subsystem volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers. For example, IBM SAN Volume Controller.

Volume Space Not Correlated (Total LUN Capacity Not Visible to Data Manager)
Total amount of unique storage subsystem volume space and virtual storage volume space discovered on monitored storage subsystem arrays that can be assigned or are assigned to servers within the network. The value in this column reflects the storage subsystems that are monitored by the subordinate server.Tivoli Storage Productivity Center does not include the following in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Disks that reside within a monitored storage subsystem.

- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Centeridentifies as back-end storage volumes on monitored storage subsystem virtualizers.

Overall Unavailable Disk Space (Total Formatted Capacity)
Total amount of formatted storage space that is used or can be used for subsystem storage volumes or virtual storage volumes. The value in this column reflects the storage subsystems that are monitored by the subordinate server. This value is dependent on the storage subsystem SMI-S provider returning a formatted value.

For some storage subsystems, this is the total managed space which includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Note: This value does not apply to IBM SAN Volume Controller.

Formatted Space (Total Formatted Space with No Volumes)
Total amount of formatted storage space that can be used for subsystem storage volumes or virtual storage volumes. The value in this column reflects the storage subsystems that are monitored by the subordinate server. This value is dependent on the storage subsystem SMI-S provider returning a formatted value.

For some subsystems, this value represents the remaining managed space that includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Formatted Space with No Volumes (Total Subsystem Total Formatted Free Space) Total storage subsystem space that is not mapped or not assigned to host systems and the total amount of formatted storage space associated with the storage subsystem that can be used for subsystem storage volumes. The value in this column reflects the storage subsystems that are monitored by the subordinate server.

This total amount of formatted storage space is dependent on the storage subsystem SMI-S provider returning a formatted value. For some subsystems, this value represents the remaining managed space that might include space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes that become missing after a probe.
- Storage subsystem volumes that are mapped or assigned to host systems.

Capacity Overhead (Total)
(IBM Tivoli Storage Enterprise Storage Server only) Sum of all the space in a disk group minus the formatted capacity of the storage subsystems monitored by a subordinate server.

Available Disk Group or Array Site Space (Total Undefined Disk Group or Array Site Capacity)
This value does not apply to all storage subsystems. For IBM Tivoli Storage Enterprise Storage Server, IBM DS6000, IBM DS8000 storage subsystems, this is the total raw disk space of any unformatted disk groups or array

sites that are monitored by the subordinate server. For XIV Storage System, the total available physical (hard) space that is left for pool creation. For other storage subsystems that contain a valid value, this is the total disk space that is monitored by the subordinate server but is not within a storage pool. This column displays **N/A** for the storage subsystems where this value is not applicable.

Cache (Total)
Amount of RAM or other cache for the storage subsystems that are monitored by the subordinate server.

Assigned Volume Space
Total storage subsystem volume space within the storage subsystem that is mapped or assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are not mapped or are not assigned to host systems.

Unassigned Volume Space (Total Subsystem Assigned LUN Capacity)
Total storage subsystem volume space within storage subsystems that is mapped or assigned to host systems. The value in this column reflects the storage subsystems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a probe.
- Storage subsystem volumes that are not mapped or are not assigned to host systems.

Unassigned and Available Formatted Disk Space (Total Subsystem Unassigned LUN Capacity)
Total storage subsystem volume space within storage subsystems that is not mapped or not assigned to host systems. The value in this column reflects the storage subsystems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a probe.
- Storage subsystem volumes that are mapped or assigned to host systems.

zOS Volume Space (Total zOS LUN Capacity)
Total storage subsystem space that is reserved for use on mainframe storage such as used by zOS. The value in this column reflects the storage subsystems that are monitored by the subordinate server.Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are reserved for use on open system server or non-mainframe storage.

Open System Volume Space (Total Open System LUN Capacity)
Total storage subsystem space that is reserved for use on open system

storage such as used by operating systems like LINUX, AIX, HP/UX, SUN Solaris, and Windows. The value in this column reflects the storage subsystems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes that are reserved for use on mainframe storage.

Physical Disk Space (Capacity)
Total amount of physical disk space discovered on the storage subsystem arrays that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, IBM SAN Volume Controller) that are attached from a storage subsystem.

Backend Volume Space (LUN Capacity)
Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage subsystem virtualizer (for example, IBM SAN Volume Controller). The value in this column reflects the volumes that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are back-end disks on a storage subsystem virtualizer that is not monitored.

Rollup Reports - Asset - Disk/Volume Groups

Use this report to view information about the disk and volume groups that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Computer: Shows disk and volume groups aggregated at the computer level and sorted by computer name.
- By Disk/Volume Group Space: Shows disk and volume groups aggregated at the computer level and sorted by storage capacity.
- By Available Disk/Volume Group Space: Shows disk and volume groups aggregated at the computer level and sorted by storage free space.
- By TPC Server: Shows disk and volume groups aggregated at the subordinate server level and sorted by subordinate server name. Click the magnifying glass icon next to the name of a subordinate server to view information about the computers it monitors.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Computer

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

Computer Name
Name of a computer with disk and volume groups that is monitored by a subordinate server.

Number of Disk/Volume Groups (Volume Groups)
Number of volume groups on a computer.

Disk/Volume Group Space (Total Capacity)
Total storage capacity of the disks and volume groups on a computer.

Available Disk/Volume Group Space (Total Free Space)
Amount of storage space available on the disk and volume groups that are monitored by a subordinate server.

OS Type
Operating system running on a computer.

Number of Disks
Number of disks on a computer.

Last Probe Time
Date and time when a computer was last probed.

By Disk/Volume Group Space (Freespace)

Computer Name
Name of a computer with disk and volume groups that are monitored by a subordinate server.

Disk/Volume Group Space (Total Free Space)
Amount of storage space available on the disks and volume groups associated with a computer.

Number of Disk/Volume Groups (Number of Volume Groups)
Number of volume groups on a computer.

Available Disk/Volume Group Space (Total Capacity)
Total storage capacity of the disks and volume groups on a computer.

OS Type
Operating system running on a computer.

Number of Disks
Number of disks on a computer.

Last Probe Time
Date and time when a computer was last probed.

By Available Disk/Volume Group Space (Capacity)

Computer Name
Name of a computer with disk and volume groups that is monitored by a subordinate server.

Available Disk/Volume Group Space (Total Capacity)
Total storage capacity of the disks and volume groups on a computer.

Number of Disk/Volume Groups (Volume Groups)
Number of volume groups on a computer.

Disk/Volume Group Space (Total Free Space)
Amount of storage space available on the disks and volume groups associated with a computer.

OS Type
Operating system running on a computer.

Number of Disks
Number of disks on a computer.

Last Probe Time
Date and time when a computer was last probed.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring computers with disk and volume groups.

Number of Disk/Volume Groups
Number of disk and volume groups that are monitored by a subordinate server.

Disk/Volume Group Space (Total Free Space)
Amount of storage space available on the disks and volume groups that are monitored by a subordinate server.

Available Disk/Volume Group Space (Total Capacity)
Total storage capacity of the disks and volume groups that are monitored by a subordinate server.

Rollup Reports - Asset - Disks

Use this report to view information about the disks that are monitored by subordinate servers in your environment.

You can generate and sort the data within Disk reports using the following by clauses:

- By Available Disk Space (Computers): Shows disks aggregated at the computer level and sorted by unallocated space. The names of computers and their system type appear in the first two columns and remain displayed when scrolling through the other columns on the report.
- By Available Disk Space (Subsystems): Shows disks aggregated at the storage subsystem level and sorted by unallocated space. The names of storage subsystems and their system type appear in the first two columns and remain displayed when scrolling through the other columns on the report.
- By TPC Server (computers): Shows computer disks aggregated at the subordinate server level. The names of subordinate servers appear in the first column and remain displayed when scrolling through the other columns on the report. Click the magnifying glass icon next to a subordinate server name to view more information about the computers monitored by that subordinate server.
- By TPC Server (subsystems): Shows storage subsystem disks aggregated at the subordinate server level. The names of subordinate servers appear in the first column and remain displayed when scrolling through the other columns on the

report. Click the magnifying glass icon next to a subordinate server name to view more information about the storage subsystems monitored by that subordinate server.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

By Available Disk Space (Computers)

TOTAL
Shows the system-wide totals for each of the columns in the report.

System Name
Name of a computer with disks that are monitored by a subordinate server.

OS Type
Operating system running on a computer.

Available Disk Space (Total Unallocated Space)
Amount of storage space on disks that is not allocated to any logical volume. Note: the value that appears in the TOTAL row for this column includes the unallocated capacity of all disks within the computers monitored by subordinate servers.

Disk Space (Total Capacity)
Amount of storage space available on the disks associated with a computer.

Number of Disks
Number of disks on a computer.

By Available Disk Space (Subsystems)

System Name
Name of a storage subsystem with disks that are monitored by a subordinate server.

Subsystem Type
Operating system running on a storage subsystem.

Available Disk Space (Total Unallocated Space)
Amount of space on the disk that is not allocated to any volumes for disks within storage subsystems). Note: the value that appears in the TOTAL row for this column includes the unallocated capacity of all disks within the storage subsystem, including the capacity of spare disks.

Disk Space (Total Capacity)
Amount of storage space available on the disks associated with a storage subsystem.

Note: For space-efficient volumes, the Available Space is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage space reported for a storage subsystem using space-efficient volumes. This also applies to other space calculations, such as the calculations for the Storage Subsystem's Consumable Volume Space and FlashCopy Target Volume Space.

Number of Disks
Total number of disks in a storage subsystem.

Physical Disk Space
Total amount of physical disk space discovered on the monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, IBM SAN Volume Controller) that are attached from a storage subsystem.

By TPC Server (Computers)

TOTAL
Shows the system-wide totals for each of the columns in the report.

TPC Server
Name of a subordinate server that is monitoring disks on computers.

Available Disk Space (Total Unallocated Space)
Amount of storage space on computer disks that is not allocated to any logical volume and are monitored by a subordinate server.

Disk Space (Total Capacity)
Amount of storage space available on all the computer disks that are monitored by a subordinate server.

Note: For space-efficient volumes, the Available Space is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage space reported for a storage subsystem using space-efficient volumes. This also applies to other space calculations, such as the calculations for the Storage Subsystem's Consumable Volume Space and FlashCopy Target Volume Space.

Number of Disks
Number of computer disks monitored by a subordinate server.

By TPC Server (Subsystems)

TPC Server
Name of a subordinate server that is monitoring disks on storage subsystems.

Available Disk Space (Total Unallocated Space)
Amount of storage space on storage subsystem disks that is not allocated to any volume and are monitored by a subordinate server.

Disk Space (Total Capacity)
Amount of storage space available on all the storage subsystem disks that are monitored by a subordinate server.

Number of Disks
Number of storage subsystem disks monitored by a subordinate server.

Physical Disk Space
Total amount of physical disk space discovered on the monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.

- Back-end storage volumes on monitored storage subsystem virtualizers (for example, IBM SAN Volume Controller) that are attached from a storage subsystem.

Rollup Reports - Asset - File Systems or Logical Volumes

Use this report to view information about the filesystems or logical volumes that are monitored by subordinate servers in your environment.

You can generate and sort the data within File Systems or Logical Volumes reports using the following by clauses:

- By Available File System or Logical Volume Space: Shows filesystems aggregated at the computer level and sorted by free space and computer name.
- By TPC Server: Shows filesystems aggregated at the subordinate server level and sorted by subordinate server name. Click the magnifying glass icon next to a subordinate server name to view more information about the filesystems or logical volumes monitored by that subordinate server.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

By Available File System or Logical Volume Space

Computer Name
Name of a computer with filesystems or logical volumes that are monitored by a subordinate server.

Available File System/Logical Volume Space (Free Space)
Amount of storage space available in the filesystems or logical volumes on a computer.

Consumed File System/Logical Volume Space (Used Space)
Amount of storage space consumed by the filesystems or logical volumes on a computer.

File System/Logical Volume Space (Capacity)
Total storage capacity of the filesystems or logical volumes on a computer.

OS Type
Operating system running on a computer.

Number of File Systems/Logical Volumes (Number of FSs/LVs)
Number of filesystems or logical volumes on a computer.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring the filesystems and logical volumes in your environment.

Available File System/Logical Volume Space (Total Free Space)
Amount of storage space available in all the filesystems or logical volumes on the computers monitored by the subordinate server.

Consumed File System/Logical Volume Space (Total Used Space)
Amount of storage space consumed by all the filesystems or logical volumes on the computers monitored by the subordinate server.

File System/Logical Volume Space (Total Capacity)
Total storage capacity of all the filesystems or logical volumes on the computers monitored by the subordinate server.

Number of File Systems/Logical Volumes (Number of FSs/LVs)
Number of filesystems or logical volumes on the computers monitored by the subordinate server.

Rollup Reports - Asset - Storage Subsystem Volumes

Use this report to view information about the Storage Subsystem Volumes that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Storage Subsystem: Shows Storage Subsystem Volume capacity and count aggregated at the storage subsystem level.
- By Storage Subsystem Volume Space: Shows Storage Subsystem Volume capacity and count aggregated at the subsystem level. Sorted according to the storage capacity of Storage Subsystem Volumes.
- By TPC Server: Shows Storage Subsystem Volume capacity and count aggregated at the subordinate server level. Sorted by subordinate server name. Click the magnifying glass icon next to a subordinate server to view information about the storage subsystems it monitors.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

By Storage Subsystem

Storage Subsytem Name
Name of a storage subsystem containing Storage Subsystem Volumes that is monitored by a subordinate server.

Storage Subsystem Type
Type of storage subsystem, such as the IBM SAN Volume Controller.

Volume Space (LUN Capacity)
Total storage capacity of the Storage Subsystem Volumes in a storage subsystem.

Consumable Volume Space (Usable LUN Capacity)
Storage capacity of Storage Subsystem Volumes in the created disk array (not including BCVs, VCM database, and Replication targets).

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Capacity of all Storage Subsystem Volumes visible to a Data agent associated with a storage subsystem.

Formatted Space (Formatted Capacity)
(IBM Tivoli Storage Enterprise Storage Server, IBM FAStT, HP StorageWorks, and Hitachi Data Systems only) Total amount of formatted storage space in the storage pools associated with a storage subsystem. For example, a RAID volume space has more capacity than a stripe-mirror volume space.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage capacity of all Storage Subsystem Volumes that the storage subsystem knows are assigned to a host computer. Note: the column will appear blank for SAN Volume Controller.

zOS Volume Space (zOS LUN Capacity)
Total storage capacity for all the Storage Subsystem Volumes allocated to zOS. Data for this column is calculated for Tivoli Storage Enterprise Storage Server only--the value for this column is set to N/A for other storage subsystems.

Number of Volumes (Number of LUNs)
Number of Storage Subsystem Volumes associated with a storage subsystem.

Backend Volume Space
Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage subsystem virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are back-end disks on a storage subsystem virtualizer that is not monitored.

By Storage Subsystem Volume Space

Volume Space (LUN Capacity)
Total storage capacity of the Storage Subsystem Volumes in a storage subsystem.

Consumable Volume Space (Usable LUN Capacity)
Storage capacity of Storage Subsystem Volumes in the created disk array (not including BCVs, VCM database, and Replication targets).

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Capacity of all Storage Subsystem Volumes visible to a Data agent associated with a storage subsystem.

Formatted Space (Formatted Capacity)
(IBM Tivoli Storage Enterprise Storage Server, IBM FAStT, HP StorageWorks, and Hitachi Data Systems only) Total amount of formatted storage space in the storage pools associated with a storage subsystem. For example, a RAID volume space has more capacity than a stripe-mirror volume space.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage capacity of all Storage Subsystem Volumes that the storage subsystem knows are assigned to a host computer.

zOS Volume Space (zOS LUN Capacity)
Total storage capacity for all the storage subsystem volumes allocated to zOS. Data for this column is calculated for Tivoli Storage Enterprise Storage Server and SAN Volume Controller only--the value for this column is set to N/A for other storage subsystems.

Number of Volumes (Number of LUNs)
Number of storage subsystem volumes associated with a storage subsystem.

Storage Subsystem Name
Name of a storage subsystem containing storage subsystem volumes that is monitored by a subordinate server.

Storage Subsystem Type
Type of storage subsystem, such as the IBM SAN Volume Controller.

Backend Volume Space
Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage subsystem virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are back-end disks on a storage subsystem virtualizer that is not monitored.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring storage subsystems.

Storage Subsystem Count
Number of storage subsystems that are monitored by a subordinate server.

Volume Space (LUN Capacity)
Total storage capacity of the Storage Subsystem Volumes that are monitored by a subordinate server.

Consumable Volume Space (Usable LUN Capacity)
Storage capacity of Storage Subsystem Volumes in the created disk array (not including BCVs, VCM database, and Replication targets) that are monitored by a subordinate server.

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Capacity of all Storage Subsystem Volumes visible to a Data agent associated with a subordinate server.

Formatted Space (Formatted Capacity)
(IBM Tivoli Storage Enterprise Storage Server, IBM FAStT, HP StorageWorks, and Hitachi Data Systems only) Total amount of formatted storage space in the storage pools associated with the storage subsystems that are monitored by a subordinate server. For example, a RAID volume space has more capacity than a stripe-mirror volume space.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage capacity of all Storage Subsystem Volumes that the monitored storage subsystems know are assigned to a host computer.

zOS Volume Space (zOS LUN Capacity)
Total storage capacity for all the Storage Subsystem Volumes allocated to zOS that are monitored by a subordinate server. Data for this column is calculated for Tivoli Storage Enterprise Storage Server only--the value for this column is set to N/A for other storage subsystems.

Number of Volumes (Number of LUNs)
Number of Storage Subsystem Volumes associated with the storage subsystems monitored by a subordinate server.

Backend Volume Space
Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage subsystem virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are back-end disks on a storage subsystem virtualizer that is not monitored.

Rollup Reports - Asset - Fabrics

Use this report to view information about the fabrics that are monitored by subordinate servers in your environment.

You view the following reports about fabrics:

- SAN Assets (ALL): Use this report to view information about all the SAN assets for each subordinate server.
- SAN Assets (Connected Dev): Use this report to view information about all the connected devices for each subordinate server.
- SAN Assets (Switches): Use this report to view information about switches in your SAN for each subordinate server.

SAN Assets (ALL)

TPC Server
Name of a subordinate server that is monitoring fabrics.

Fabric ID
The ID of the fabric that the asset is in.

Type The type of asset discovered. For example: switch, host, host bus adapter, storage subsystem.

Label The label of the asset. For example: the host name, switch name, storage subsystem name, node name of the port.

Identifier
The unique identifier of the switch. For example, the switch WWN.

Status The status of the asset: Normal, Warning, or Critical.

SAN Assets (Connected Dev)

TPC Server
Name of a subordinate server that is monitoring a connected device.

Fabric ID
The ID of the fabric that the asset is in.

Type The type of asset discovered. For example: switch, host, host bus adapter, storage subsystem.

Label The label of the asset. For example: the host name, switch name, storage subsystem name, node name of the port.

Identifier
The unique identifier of the switch. For example, the switch WWN.

Status The status of the asset: Normal, Warning, or Critical.

SAN Assets (Switches)

TPC Server
Name of a subordinate server that is monitoring a switch.

Fabric ID
The ID of the fabric that a switch is in.

Type The type of switch discovered.

Switch Type
The switch type, such as physical.

Label The label of a switch.

Status The status of a switch: Normal, Warning, or Critical.

IP Address
The IP address of a switch.

Vendor ID
The manufacturer of the switch.

Model The model name and number of a switch.

Version
The version information of the switch.

Serial Number
The serial number of the switch.

Object ID
The object ID of the switch.

Management ID
The management ID of the switch.

Management Address
The URL of the switch element manager.

Domain
The name of the domain that the switch is in.

Contact
The contact information for the switch.

WWN The world-wide name (WWN) of the switch.

Location
The physical location of the switch.

Physical Switch WWN
The world-wide name (WWN) of the physical switch.

Rollup Reports - Asset - Hypervisors

Use this report to view information about the hypervisors that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clause:

- By Hypervisor: Sorted according to hypervisor name. The names of hypervisors monitored by subordinate servers appear in the first column and remain displayed when scrolling through the other columns on the report.

By Hypervisors

Hypervisor Name
Name of a hypervisor monitored by a subordinate server.

OS Version
Version of the operating system running on a computer.

RAM The amount of RAM installed on a computer.

Disk Space
Total raw (unformatted) disk capacity of a storage subsystem. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Capacity of spare disks identified on IBM Tivoli Storage Enterprise Storage Server, IBM DS8000, and IBM DS6000 storage subsystems.
- Capacity of storage subsystem disks that become missing after a storage subsystem probe.

Note: For Hitachi Data Systems storage subsystems, this is the capacity of all the Volumes in the storage subsystem.

Number of Disks
Total number of disks for a computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value in this column:

- Computer disks on a computer that has not been probed by a Data agent.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

Processor Count
Number of CPUs running on a computer.

IP Address
IP address of a computer.

Serial Number
Serial number of a computer.

Model Model number or name of a computer.

Manufacturer
Manufacturer of a computer (for example, IBM).

Time Zone
Time zone where the computer is located.

Available File System Space
Total amount unique file system storage space that is not used or available to the computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

OS Type
Operating system running on a computer.

Network
Network name of a computer.

Processor Type
Type of processor running on a computer (for example, GenuineIntel: x86 Family 15).

Processor Speed
Speed at which a computer's processor is running (in MHz).

CPU Architecture
Architecture of the processor (for example, IA32).

Swap Space
Amount of swap space (virtual memory) defined on a computer.

Rollup Reports - Database Asset

Use these reports to view detailed statistics that have been collected by subordinate servers about the RDBMSs in your environment, including Oracle, SQL Server, Sybase, and UDB/DB2.

The information shown in these reports is gathered by the master server when it runs Tivoli Storage Productivity Center server probes against associated subordinate servers. If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

Rollup Reports - Database Asset - All DBMSs

Use this report to view information about the RDBMSs (including Oracle, Sybase, SQL Server, and UDB/DB2) that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Instance: Shows databases for each computer sorted by computer name.
- By Version: Shows databases for each computer sorted by RDBMS version.
- By DBMS Type: Shows databases for each computer sorted by RDBMS type.
- By TPC Server: Shows databases aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

By Instance

Computer Name
Name of the computer on which an Instance resides.

Instance
Name of an instance.

RDBMS Type
RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

RDBMS Version
Version of the RDBMS under which an instance is running.

Port Port on which an instance listens for requests.

Home Directory
Path of the home directory for an instance.

Logonid
ID used by Tivoli Storage Productivity Center to log in to an instance.

Probe Status
Status of the last probe run against an instance.

Last Start Time
Date and time when an instance was last started.

Last Probe Time
Date and time when an instance was last probed.

Tablespace count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

By Version

RDBMS Version
Version of the RDBMS under which an instance is running.

Computer Name
Name of the computer on which an instance resides.

Instance
Name of an instance.

RDBMS Type
RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

Port Port on which an instance listens for requests.

Home Directory
Path of the home directory for an instance.

Logonid
ID used by Tivoli Storage Productivity Center to log in to an instance.

Probe Status
Status of the last probe run against an instance.

Last Start Time
Date and time when an instance was last started.

Last Probe Time
Date and time when an instance was last probed.

Tablespace Count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

By DBMS Type

RDBMS Type
RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

Computer Name
Name of the computer on which an instance resides.

Instance
Name of an instance.

RDBMS Version
Version of the RDBMS under which an instance is running.

Port Port on which an instance listens for requests.

Home Directory
Path of the home directory for an instance.

Logonid
ID used by Tivoli Storage Productivity Center to log in to an instance.

Probe Status
Status of the last probe run against an instance.

Last Start Time
Date and time when an instance was last started.

Last Probe Time
Date and time when an instance was last probed.

Tablespace count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring DBMSs in your environment.

Total Number of DBMSs
Number of DBMSs monitored by a subordinate server.

Number of UDB Instances
Number of UDB instances monitored by a subordinate server.

Number of Oracle Instances
Number of Oracle instances monitored by a subordinate server.

Number of SQL/Server Instances
Number of SQL/Server instances monitored by a subordinate server.

Number of Sybase Instances
Number of Sybase instances monitored by a subordinate server.

Tablespace Count
Number of tablespaces associated with the instances monitored by a subordinate server.

File Count
Number of data files associated with the instances monitored by a subordinate server.

Log File Count
Number of log files associated with the instances monitored by a subordinate server.

File Free Space
Amount of storage free space available on the data files associated with the instances monitored by a subordinate server.

Total Size
Amount of storage space allocated to the instances monitored by a subordinate server.

Total File Capacity (File Capacity)
Amount of storage space allocated to the data files associated with the instances monitored by a subordinate server.

Rollup Reports - Database Asset - Oracle, SQL/Server, Sybase, UDB

Use this report to view information about Oracle, Sybase, SQL Server, and UDB/DB2 instances that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Instance: Shows databases for each computer sorted by computer name.
- By Version: Shows databases for each computer sorted by the version of the RDBMS under which a database is running.
- By TPC Server: Shows databases aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Instance

Computer Name
Name of the computer on which an Instance resides.

Instance
Name of an instance.

RDBMS Type
RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

RDBMS Version
Version of the RDBMS under which an instance is running.

Port Port on which an instance listens for requests.

Home Directory
Path of the home directory for an instance.

Logonid
ID used by Tivoli Storage Productivity Center to log in to an instance.

Probe Status
Status of the last probe run against an instance.

Last Start Time
Date and time when an instance was last started.

Last Probe Time
Date and time when an instance was last probed.

Tablespace count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

By Version

RDBMS Version
Version of the RDBMS under which an instance is running.

Computer Name
Name of the computer on which an instance resides.

Instance
Name of an instance.

RDBMS Type
RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

Port Port on which an instance listens for requests.

Home Directory
Path of the home directory for an instance.

Logonid
ID used by Tivoli Storage Productivity Center to log in to an instance.

Probe Status
Status of the last probe run against an instance.

Last Start Time
Date and time when an instance was last started.

Last Probe Time
Date and time when an instance was last probed.

Tablespace Count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring DBMSs in your environment.

Number of Oracle Instances
Number of Oracle instances monitored by a subordinate server.

Number of SQL/Server Instances
Number of SQL/Server instances monitored by a subordinate server.

Number of Sybase Instances
Number of Sybase instances monitored by a subordinate server.

Number of UDB Instances
Number of UDB instances monitored by a subordinate server.

Tablespace count
Number of tablespaces associated with the instances monitored by a subordinate server.

File Count
Number of data files associated with the instances monitored by a subordinate server.

Log File Count
Number of log files associated with the instances monitored by a subordinate server.

File Free Space
Amount of storage free space available on the data files associated with the instances monitored by a subordinate server.

Total Size
Amount of storage space allocated to the instances monitored by a subordinate server.

Total File Capacity
Amount of storage space allocated to the data files associated with the instances monitored by a subordinate server.

Rollup Reports - TPC Server Storage Space

Use these reports to view storage metrics related to the disk space, filesystem space, consumed filesystem space, and available filesystem space gathered by the subordinate servers in your environment.

The information shown in these reports is gathered by the master server when it runs "Collecting storage statistics from IBM Tivoli Storage Productivity Center servers (TPC server probes)" on page 207 against associated subordinate servers. If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

Rollup Reports - TPC Server Storage Space - Disk Space

Use this report to view disk capacity information about the computers and storage subsystems that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Storage Subsystem: Shows disk capacity aggregated at the storage subsystem level and sorted by storage subsystem name. Storage subsystems that have been discovered but not probed will not appear in this report.
- By Storage Subsystem Group: Shows disks capacity aggregated at the storage subsystem monitoring group and storage subsystem reporting group levels and sorted by group name. Storage subsystems that have been discovered but not probed will not be included in the values shown in this report.
- By Computer: Shows disks capacity aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows disks capacity aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.

- By Cluster: Shows disks capacity aggregated at the cluster level and sorted by cluster name.
- By TPC Server (Computers): Shows computer disk capacity aggregated at the subordinate server level and sorted by subordinate server name.
- By TPC Server (Storage Subsystems): Shows storage subsystem disk capacity aggregated at the subordinate server level and sorted by subordinate server name. Storage subsystems that have been discovered but not probed will not be included in the values shown in this report.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

Note: The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

By Storage Subsystem

Subsystem Name
Name of a storage subsystem that is monitored by a subordinate server.

Subsystem Type
Type of storage subsystem, such as such as the IBM SAN Volume Controller.

Disk Space (Disk Capacity)
Total raw (unformatted) disk capacity of a storage subsystem.

Available Disk Space (Available Capacity)
Amount of storage space available on the disks associated with a storage subsystem. For Hitachi Data Systems storage subsystems, internal disk information is not available. This field will be 0.

Number of Disks
Number of disks associated with a storage subsystem. For Hitachi Data Systems storage subsystems, internal disk information is not available. This field will be set to 1.

Number of Volumes (Number of LUNs)
Number of volumes associated with a storage subsystem.

Storage Subsystem Volume Space (LUN capacity)
Amount of storage space allocated to the volumes associated with the storage subsystem.

Consumable Storage Subsystem Volume Space (Usable LUN Capacity)
Storage capacity of volumes in the created disk array (not including BCVs, VCM database, and Replication targets).

Correlated Volume Space (Visible LUN Capacity)
Capacity of all volumes visible to a Data agent associated with a storage subsystem.

FlashCopy Target Storage Subsystem Volume Space (FlashCopy Target Capacity)
The total capacity of all FlashCopy target volumes in the subsystem. Note: Tivoli Storage Enterprise Storage Server subsystems must have at least the following microcode levels: ess800mincodelevel = 2.4.3.56, essf20mincodelevel = 2.3.3.89. Tivoli Storage Productivity Center will not report FlashCopy information if the Tivoli Storage Enterprise Storage Server subsystems do not meet this requirement.

Volume Space Not Correlated (Not Visible LUN Capacity)
Usable volume capacity minus assigned volume capacity.

Unavailable Disk Space (Overhead)
Amount of capacity that is dedicated to redundancy.

Formatted Disk Space (Formatted Capacity)
(IBM Tivoli Storage Enterprise Storage Server, IBM FAStT, HP StorageWorks, and Hitachi Data Systems only) Total amount of formatted storage space in the storage pools associated with the storage subsystem. For example, a RAID volume space has more capacity than a stripe-mirror volume space.

Formatted Disk Space with No Volumes (Total Formatted Space with No Volumes)
(IBM Tivoli Storage Enterprise Storage Server, IBM FAStT, HP StorageWorks, and Hitachi Data Systems only) Amount of storage space available to allocate to volumes in the storage pools associated with the storage subsystem.

Available Disk Group or Array Site Space (Undefined Disk Group or Array Site Capacity)
Raw disk space in ranks on which no volume space has been defined. For IBM Tivoli Storage Enterprise Storage Server subsystems, this is the total raw capacity of any unformatted disk groups. For IBM FAStT, HP StorageWorks, and Hitachi Data Systems storage subsystems this is the total capacity of the disks that are not in a storage pool.

Cache Amount of RAM or other cache for a storage subsystem.

Note: This column is blank for SAN Volume Controllers.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage capacity of all volumes that the storage subsystem knows are assigned to a host computer. Note: the column will appear blank for SAN Volume Controller.

Unassigned Volume Space (Subsystem Unassigned LUN Capacity)
Total storage capacity of all volumes that the storage subsystem knows are not assigned to a host computer. Note: the column will appear blank for SAN Volume Controller.

zOS Volume Space (zOS LUN Capacity)
Total storage capacity for all the volumes allocated to zOS. Data for this column is calculated for Tivoli Storage Enterprise Storage Server only--the value for this column is set to N/A for other storage subsystems.

Open System Volume Space (Open System LUN Capacity)
Total volume storage capacity (excluding the zOS volume space).

Physical Disk Space
Total amount of physical disk space discovered on the monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller) that are attached from a storage subsystem.

Backend Volume Space
Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage subsystem virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are back-end disks on a storage subsystem virtualizer that is not monitored.

By Storage Subsystem Group

Subsystem Group Name
Name of a storage subsystem monitoring group or storage subsystem reporting group.

TPC Server
Name of a subordinate server to which the group is associated.

Disk Space (Disk Capacity)
Total raw (unformatted) disk capacity of the storage subsystems in a group.

Available Disk Space (Available Capacity)
Amount of storage space available on the disks associated with the storage subsystems in a group.

Number of Disks
Number of disks associated with the storage subsystems in a group.

Number of Volumes (Number of LUNs)
Number of volumes associated with the storage subsystems in a group.

Storage Subsystem Volume Space (LUN capacity)
Amount of storage space allocated to the volumes associated with the storage subsystems in a group.

Consumable Storage Subsystem Volume Space (Usable LUN Capacity)
Storage capacity of volumes in the created disk array (not including BCVs, VCM database, and Replication targets).

FlashCopy Target Storage Subsystem Volume Space (FlashCopy Target Capacity)
The total capacity of all FlashCopy target volumes in the storage subsystems in a group.

Correlated Volume Space (LUN Capacity Visible to Data Manager)
Capacity of all volumes visible to a Data agent associated with the storage subsystems in a group.

Volume Space Not Correlated (LUN Capacity Not Visible to Data Manager)
Usable volume capacity minus assigned volume capacity on the storage subsystems in a group.

Unavailable Volume Space (Overhead)
Amount of capacity that is dedicated to redundancy in the storage subsystems in a group.

Formatted Space (Formatted Capacity)
(IBM Tivoli Storage Enterprise Storage Server, IBM FAStT, HP

StorageWorks, and Hitachi Data Systems only) Total amount of formatted storage space in the storage pools associated with the storage subsystems in a group.

Total Formatted Space with no Volumes (Total Formatted Space with No Volumes)
(IBM Tivoli Storage Enterprise Storage Server, IBM FAStT, HP StorageWorks, and Hitachi Data Systems only) Amount of storage space available to allocate to volumes in the storage pools associated with the storage subsystems in a group.

Available Disk Group or Array Site Space (Total Overhead)
(IBM Tivoli Storage Enterprise Storage Server only) Sum of all the space in a disk group minus the formatted capacity on the storage subsystems in a group.

Cache Amount of RAM or other cache for the storage subsystems in a group.

Note: This column does not include values for SAN Volume Controllers.

Assigned Volume Space (Subsystem Assigned LUN Capacity)
Total storage capacity of all volumes that the storage subsystems in a group know are assigned to a host computer.

Unassigned Volume Space (Subsystem Unassigned LUN Capacity)
Total storage capacity of all volumes that the storage subsystems in a group know are not assigned to a host computer.

zOS Volume Space (zOS LUN Capacity)
Total storage capacity for all the volumes on the storage subsystems in a group allocated to zOS.

Open System Volume Space (Open System LUN Capacity)
Total volume storage capacity (excluding the zOS volume capacity) of the storage subsystems in a group.

Physical Disk Space
Total amount of physical disk space discovered on the monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller) that are attached from a storage subsystem.

Backend Volume Space
Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage subsystem virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are back-end disks on a storage subsystem virtualizer that is not monitored.

By Computer

Computer Name
Name of a computer.

Disk Space (Total Capacity)
Total storage capacity for a computer. If the computer is an IBM Tivoli Storage SAN File Systems client, then this value will also include the capacity of any volumes visible to this computer that are owned by the SAN File System.

Available Disk Space (Total Unallocated Space)
Amount of unused storage space on a computer (not in filesystems seen by this operating system). If the computer is an IBM Tivoli Storage SAN File Systems client, then this value excludes the unallocated space on any volumes visible to it that are owned by the SAN File System.

Number of Disks
Number of disks on a computer.

Owned Disk Space (Owned Capacity)
Total storage capacity for a computer that is owned by that computer. If Data Manager can determine that a disk was configured for a different host, that disk's capacity will not be counted in the ″Owned Capacity″.

If the computer is an IBM Tivoli Storage SAN File Systems client, then this value does not include the capacity of any volumes visible to it that are owned by the SAN File System. If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product will examine the Veritas disk-label and read off the host-ID. If this does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

Other Available Disk Space (Other Unallocated Space)
Total unallocated space on disks owned by the computer. This column does not display the total unallocated disk space on disks not owned by the computer. If the computer is an IBM Tivoli Storage SAN File Systems client, then this value excludes the unallocated space on any volumes visible to it that are owned by the SAN File System.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)
Total disk capacity of all non-Fibre attached disks for a computer.

Fibre Channel Attached Disk Space (Total Other Attached Unallocated Space)
Total unallocated space of all non-Fibre attached disks for a computer.

Non-Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Capacity)
Total disk capacity of all fibre channel attached disks for a computer.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)
Total unallocated space of all fibre channel attached disks for a computer.

By Computer Group

Computer Group Name
Name of a computer monitoring group or computer reporting group.

Disk Space (Total Capacity)
Total storage capacity for the computers in a group. If the computer is an

IBM Tivoli Storage SAN File Systems client, then this value will also include the capacity of any volumes visible to this computer that are owned by the SAN File System.

Available Disk Space (Total Unallocated Space)
Amount of unused storage space on the computers in a group.

Number of Disks
Number of disks on the computers in a group.

Owned Disk Space (Owned Capacity)
Total storage capacity for the computers in a group that is owned by those computers. If Data Manager can determine that a disk was configured for a different host, that disk's capacity will not be counted in the "Owned Capacity".

If any computers in the group are IBM Tivoli Storage SAN File Systems clients, then this value does not include the capacity of any volumes visible to it that are owned by the SAN File System. If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product will examine the Veritas disk-label and read off the host-ID. If this does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

Other Available Disk Space (Other Unallocated Space)
Total unallocated space on disks owned by the computers in a group. This column does not display the total unallocated disk space on disks not owned by computers. If any of the computers in the group are IBM Tivoli Storage SAN File Systems clients, then this value excludes the unallocated space on any volumes visible to it that are owned by the SAN File System.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)
Total disk capacity of all non-Fibre attached disks on the computers in a group.

Fibre Channel Attached Disk Space (Total Other Attached Unallocated Space)
Total unallocated space of all non-Fibre attached disks on the computers in a group.

Non-Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Capacity)
Total disk capacity of all fibre channel attached disks on the computers in a group.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)
Total unallocated space of all fibre channel attached disks on the computers in a group.

By Cluster

Cluster Name
Name of a cluster.

Disk Space (Total Capacity)
Total storage capacity for a cluster.

Available Disk Space (Total Unallocated Space)
Amount of unused storage space on a cluster (not in filesystems seen by this operating system).

Owned Disk Space (Owned Capacity)
Total storage capacity for a cluster that is owned by that cluster. If Data

Manager can determine that a disk was configured for a different host, that disk's capacity will not be counted in the ″Owned Capacity″.

If the computer is an IBM Tivoli Storage SAN File Systems client, then this value does not include the capacity of any volumes visible to it that are owned by the SAN File System. If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product will examine the Veritas disk-label and read off the host-ID. If this does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

Other Available Disk Space (Other Unallocated Space)
Total unallocated space on disks owned by the cluster. This column does not display the total unallocated disk space on disks not owned by the cluster. If a computer in a cluster is an IBM Tivoli Storage SAN File Systems client, then this value excludes the unallocated space on any volumes visible to it that are owned by the SAN File System.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)
Total disk capacity of all non-Fibre attached disks in a cluster.

Fibre Channel Attached Disk Space (Total Other Attached Unallocated Space)
Total unallocated space of all non-Fibre attached disks in a cluster.

Non-Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Capacity)
Total disk capacity of all fibre channel attached disks in a cluster.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)
Total unallocated space of all fibre channel attached disks in a cluster.

By TPC Server (Computers)

TPC Server
Name of a subordinate server that is monitoring computer disks in your environment.

Disk Space (Total Capacity)
Total storage capacity of all the computer disks monitored by a subordinate server.

Available Disk Space (Total Unallocated Space)
Amount of unused storage space on all the computer disks monitored by a subordinate server.

Number of Disks
Number of computer disks monitored by a subordinate server.

Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)
Total disk capacity of all non-Fibre attached disks monitored by a subordinate server.

Fibre Channel Attached Disk Space (Total Other Attached Unallocated Space)
Total unallocated space of all non-Fibre attached disks monitored by a subordinate server.

Non-Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Capacity)
Total disk capacity of all fibre channel attached disks monitored by a subordinate server.

Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)
Total unallocated space of all fibre channel attached disks monitored by a subordinate server.

By TPC Server (Storage Subsystems)

TPC Server
Name of a subordinate server that is monitoring storage subsystem disks in your environment.

Disk Space (Total Capacity)
Total raw (unformatted) disk capacity of the storage subsystems monitored by a subordinate server.

Available Disk Space (Total Unallocated Space)
Amount of storage space available on the disks associated with the storage subsystems monitored by a subordinate server.

Number of Disks
Number of storage subsystem disks monitored by a subordinate server.

Physical Disk Space
Total amount of physical disk space discovered on the monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller) that are attached from a storage subsystem.

Rollup Reports - TPC Server Storage Space - File System Space

Use this report to view file system space information about the file systems that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Computer: Shows file system capacity aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows file system capacity aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By Cluster: Shows file system capacity aggregated at the cluster level and sorted by cluster name.
- By TPC Server: Shows file system capacity aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Computer

Computer Name
Name of the computer on which a file system resides.

Total Capacity
Total amount (capacity) of storage space in the file systems on a computer.

Note that the value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**, the difference representing the space wasted due to the system overhead.

Total Used Space
Total amount of used storage space in the file systems on a computer. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space in the file systems on a computer.

Number of FS's/LV's
Number of file systems or logical volumes on a computer.

By Computer Group

Computer Group
Name of a computer monitoring group or computer reporting group.

Total Capacity
Total amount (capacity) of storage space on the file systems in a group. Note that the value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**, the difference representing the space wasted due to the system overhead.

Total Used Space
Total amount of used storage space on the file systems in a group. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space on the file systems in a group.

Number of FS's/LV's
Number of file systems or logical volumes in a group.

By Cluster

Cluster
Name of a cluster.

Total Capacity
Total storage capacity of the file systems in a cluster. Note that the value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**, the difference representing the space wasted due to the system overhead.

Total Used Space
Total amount of used storage space on the file systems in a cluster. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space on the file systems in a cluster.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring file systems in your environment.

Total Capacity
Total storage capacity of all the file systems monitored by a subordinate server.

Total Used Space
Total amount of used storage space on the file systems monitored by a subordinate server. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space on the file systems monitored by a subordinate server.

Total Number of FS's/LV's
Total number of file systems or logical volumes monitored by a subordinate server.

Rollup Reports - TPC Server Storage Space - Consumed File System Space

Use this report to view information about the used storage space on filesystems that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Computer: Shows filesystem used space aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows filesystem used space aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By Cluster: Shows filesystem used space aggregated at the cluster level and sorted by cluster name.
- By TPC Server: Shows filesystem used space aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Computer

Computer Name
Name of the computer on which a filesystem resides.

Total Used Space
Total amount of used storage space in the filesystems on a computer. On some UNIX filesystems, this value does not include storage space reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space in the filesystems on a computer.

Total Capacity
Total amount (capacity) of storage space in the filesystems on a computer. Note that the value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**, the difference representing the space wasted due to the system overhead.

Number of FS's/LV's
Number of filesystems or logical volumes on a computer.

By Computer Group

Computer Group
Name of a computer monitoring group or computer reporting group.

Total Capacity
Total amount (capacity) of storage space on the filesystems in a group. Note that the value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**, the difference representing the space wasted due to the system overhead.

Total Used Space
Total amount of used storage space on the filesystems in a group. On some UNIX filesystems, this value does not include storage space reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space on the filesystems in a group.

Number of FS's/LV's
Number of filesystems or logical volumes in a computer group.

By Cluster

Cluster
Name of a cluster.

Total Used Space
Total amount of used storage space on the filesystems in a cluster. On some UNIX filesystems, this value does not include storage space reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space on the filesystems in a cluster.

Total Capacity
Total storage capacity of the filesystems in a cluster. Note that the value in

this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**, the difference representing the space wasted due to the system overhead.

Number of FS's/LV's
Number of filesystems or logical volumes in a cluster.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring filesystems in your environment.

Total Used Space
Total amount of used storage space on the filesystems monitored by a subordinate server. On some UNIX filesystems, this value does not include storage space reserved by the operating system (overhead).

Total Free Space
Total amount of unused storage space on the filesystems monitored by a subordinate server.

Total Capacity
Total storage capacity of all the filesystems monitored by a subordinate server.

Total Number of FS's/LV's
Number of filesystems or logical volumes monitored by a subordinate server.

Rollup Reports - TPC Server Storage Space - Available File System Space

Use this report to view information about the available storage space on filesystems that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Computer: Shows filesystem free space aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows filesystem free space aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By Cluster: Shows filesystem free space aggregated at the cluster level and sorted by cluster name.
- By IBM Tivoli Storage Productivity Center Server: Shows filesystem free space aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Computer

Computer Name
Name of the computer on which a filesystem resides.

Total Free Space
Total amount of unused storage space in the filesystems on a computer.

Total Used Space
Total amount of used storage space in the filesystems on a computer. On some UNIX filesystems, this value does not include storage space reserved by the operating system (overhead).

Total Capacity
Total amount (capacity) of storage space in the filesystems on a computer. Note that the value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**, the difference representing the space wasted due to the system overhead.

Number of FS's/LV's
Number of filesystems or logical volumes on a computer.

By Computer Group

Computer Group
Name of a computer monitoring group or computer reporting group.

Total Free Space
Total amount of unused storage space on the filesystems in a group.

Total Used Space
Total amount of used storage space on the filesystems in a group. On some UNIX filesystems, this value does not include storage space reserved by the operating system (overhead).

Total Capacity
Total amount (capacity) of storage space on the filesystems in a group. Note that the value in this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**, the difference representing the space wasted due to the system overhead.

Number of FS's/LV's
Number of filesystems or logical volumes in a computer group.

By Cluster

Cluster
Name of a cluster.

Total Free Space
Total amount of unused storage space on the filesystems in a cluster.

Total Used Space
Total amount of used storage space on the filesystems in a cluster. On some UNIX filesystems, this value does not include storage space reserved by the operating system (overhead).

Total Capacity
Total storage capacity of the filesystems in a cluster. Note that the value in

this field might be greater than **Total Used Space** + **Total Free Space** for some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Total Capacity** field is greater than or equal to the value of **Total Used Space** + **Total Free Space**, the difference representing the space wasted due to the system overhead.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring filesystems in your environment.

Total Free Space
Total amount of unused storage space on the filesystems monitored by a subordinate server.

Total Used Space
Total amount of used storage space on the filesystems monitored by a subordinate server. On some UNIX filesystems, this value does not include storage space reserved by the operating system (overhead).

Total Capacity
Total storage capacity of all the filesystems monitored by a subordinate server.

Number of FS's/LV's
Number of filesystems or logical volumes monitored by a subordinate server.

Rollup Reports - Database Capacity

Use these reports to view storage capacity, used space, and free space at the computer, computer group, database, instance, and IBM Tivoli Storage Productivity Center server level for RDBMSs that are monitored by subordinate servers.

For a specific instance, computer, group of computers, or database that is monitored by a subordinate server, you can use these rollup reports to view:

- how much storage you have for an RDBMS and the objects it contains
- how much storage is being used for an RDBMS and the objects it contains
- what percentage of RDBMS storage is free for expansion

The information shown in these reports is gathered by the master server when it runs Tivoli Storage Productivity Center server probes against associated subordinate servers. If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

Rollup Reports - Database Capacity - All DBMSs

Use this report to view storage capacity information for the DBMSs (including Oracle, Sybase, SQL Server, and UDB/DB2) that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Instance: Shows database capacity aggregated at the computer level and sorted by instance name.

- By Computer: Shows database capacity aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows database capacity aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By TPC Server: Shows database capacity aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Instance

Instance
Name of an instance.

RDBMS Type
RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

Computer Name
Name of the computer on which an Instance resides.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Tablespace Count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

By Computer

Computer Name
Name of the computer on which an Instance resides.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Tablespace Count
Number of Tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

By Computer Group

Computer Group
Name of a computer monitoring group or reporting group.

Total Size
Amount of storage space allocated to instances in a group.

File Capacity
Amount of storage space allocated to the data files associated with instances in a group.

File Free Space
Amount of storage free space available on the data files associated with instances in a group.

Tablespace Count
Number of databases or tablespaces associated with instances in a group.

File Count
Number of data files associated with instances in a group.

Log File Count
Number of log files associated with instances in a group.

TPC Server
Name of a subordinate server that is monitoring the computer group.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring DBMSs in your environment.

Total Size
Amount of storage space allocated to instances monitored by a subordinate server.

Total File Capacity
Amount of storage space allocated to the data files associated with instances monitored by a subordinate server.

File Free Space
Amount of storage free space available on the data files associated with instances monitored by a subordinate server.

Tablespace Count
Number of tablespaces associated with instances monitored by a subordinate server.

File Count
Number of data files associated with instances monitored by a subordinate server.

Log File Count
Number of log files associated with instances monitored by a subordinate server.

Rollup Reports - Database Capacity - Oracle, Sybase, SQL Server, and UDB

Use this report to view storage capacity information for the Oracle, Sybase, SQL Server, and UDB/DB2 databases that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Database: Shows database capacity aggregated at the computer level and sorted by instance name.
- By Computer: Shows database capacity aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows database capacity aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By TPC Server: Shows database capacity aggregated at the subordinate server level and sorted by subordinate server name.

Note: The first column in each report remains displayed when scrolling through the other columns on the report.

By Database

Database Instance
Name of a database instance.

Computer Name
Name of the computer on which an instance resides.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Tablespace Count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

By Computer

Computer Name
Name of the computer on which an instance resides.

Total Size
Amount of storage space allocated to an instance.

File Capacity
Amount of storage space allocated to the data files associated with an instance.

File Free Space
Amount of storage free space available on the data files associated with an instance.

Tablespace Count
Number of tablespaces associated with an instance.

File Count
Number of data files associated with an instance.

Log File Count
Number of log files associated with an instance.

By Computer Group

Computer Group
Name of a computer monitoring group or reporting group.

Total Size
Amount of storage space allocated to instances in a group.

File Capacity
Amount of storage space allocated to the data files associated with instances in a group.

File Free Space
Amount of storage free space available on the data files associated with instances in a group.

Tablespace Count
Number of tablespaces associated with instances in a group.

File Count
Number of data files associated with instances in a group.

Log File Count
Number of log files associated with instances in a group.

TPC Server
Name of a subordinate server that is monitoring the group.

By TPC Server

TPC Server
Name of a subordinate server that is monitoring DBMSs in your environment.

Total Size
Amount of storage space allocated to instances monitored by a subordinate server.

Total File Capacity
Amount of storage space allocated to the data files associated with instances monitored by a subordinate server.

File Free Space
Amount of storage free space available on the data files associated with instances monitored by a subordinate server.

Tablespace Count
Number of tablespaces associated with instances monitored by a subordinate server.

File Count
Number of data files associated with instances monitored by a subordinate server.

Log File Count
Number of log files associated with instances monitored by a subordinate server.

Data source reports - overview

Use data source reports to view detailed information about the agents from which IBM Tivoli Storage Productivity Center collects information about storage resources. Data sources include CIMOM agents, Data agents, Fabric agents, VMWare, and subordinate Tivoli Storage Productivity Center servers.

Data source reports - CIMOM agents

Use these reports to view detailed information about the CIMOM agents that have been associated with IBM Tivoli Storage Productivity Center.

Related Topics

- Selecting resources upon which to report
- Filtering the resources that appear in a report
- Generating a report
- Drilling down a report row (By Managed Device report only)
- Navigating tabs

Data source report - By CIMOM Agent

Use this report to view detailed information about the CIMOM agents that are used as data sources for IBM Tivoli Storage Productivity Center.

Each CIMOM that has been associated with IBM Tivoli Storage Productivity Center appears as its own row in this report.

Service URL
The service URL of the CIMOM containing the IP address of the CIMOM, the port on which the CIMOM is listening, and the protocol used for communication. This URL has a protocol [http|https], an IP address (IPv4 or IPv6), and a port number. For example:

- **IPv4:** https://127.0.0.1:5989
- **IPv6:** https://[2001:DB8::1234:0000:0000:5678:ABCD]:5989

Display Name
The name of the CIMOM as specified by the CIMOM provider that will appear in the Tivoli Storage Productivity Center user interface.

Description
The optional description that was entered on the Add CIMOM panel.

Managed Device Category
The type of device that is managed by a CIMOM (for example, storage subsystem, switch, or tape library). The value **NA** appears in this column if you have not run a CIMOM discovery against a CIMOM or a CIMOM does not manage any devices.

Connection Status
The status of this CIMOM with respect to Tivoli Storage Productivity Center. Possible values are: SUCCESS, UNCONFIGURED, UNKNOWN,

INVALID_NAMESPACE, TIMEOUT, REFUSED, LOGIN_FAILED, SSL_HANDSHAKE_ERROR, SSL_REGISTRATION_INVALID, CIMCLIENT_ERROR

Status Timestamp
The date and time when the **Connection Status** information was last collected.

Number of Managed Devices
The number of devices managed by a CIMOM.

Interoperability Namespace
The interoperability namespace of a CIMOM. This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information. For a NetApp CIMOM, enter /interop.

Username
The CIMOM user name used for authentication.

Truststore Location
The location (path on this computer) of a certificate file for certificate based authentication in the https protocol.

User Interface Description
The name of the Human Interface Service (if any) supported by this CIMOM.

Software Level
The software version level of the CIMOM agent.

Protocol Version
The version of the cim-xml protocol.

Authentication Mechanism
The authentication mechanism supported by the CIMOM. This column can contain the following values: **Unknown**, **None**, **Other**, **Basic**, **Digest**.

Alias The alias of the CIMOM.

Protocol
The protocol used to communicate with the CIMOM. Possible values are **http** and **https**.

SLP Attributes
The standard set of attributes for this CIMOM. The attributes are retrieved via SLP.

Data source report - By Managed Device

Use this report to view information about the devices managed by CIMOMs in your environment. Only devices that are managed by the CIMOMs used as data sources for IBM Tivoli Storage Productivity Center are displayed.

You must run a CIMOM discovery job before you can view the devices managed by CIMOMs in this report. If a device is managed by multiple CIMOMs, the report displays a separate row for each of those CIMOMs. Therefore, the number of rows for a device in the report is equal to the number of CIMOMs managing that device.

Magnifying glass icon 
Click this button to view a more detailed report about a device.

Device Name
The display name (label) of the managed storage subsystem, switch, or tape library.

Device Category
Indicates whether a device that is managed by a CIMOM is a storage subsystem, switch, or tape library.

Device Type
Indicates the type of a managed device based on its category. For example, a device that is a storage subsystem might have a value of **DS8000** in this column.

Device Status
The consolidated status of the managed device.

Device Operational Status
The operational status of a storage subsystem.

Device Last Probe Time
The date and time when a probe was last run on a managed device.

Device Vendor
The manufacturer of the managed device.

Device Model
The model number of the managed device.

Device Serial Number
The serial number of the managed device.

Device Microcode Level
The microcode or firmware level of the managed device.

Device IP Address
The IP address of the managed device.

Device Element Manager
The URL or the fully-qualified name of the executable file that launches the element manager.

CIMOM Service URL
The service URL of the CIMOM. This URL includes the IP address or host name of the CIMOM, the port on which the CIMOM is listening, and the protocol used for communication (http|https). For example: https://127.0.0.1:9551.

CIMOM Display Name
The name of the CIMOM as specified by the CIMOM provider that will appear in the Tivoli Storage Productivity Center user interface.

CIMOM Description
The optional description that was entered on the Add CIMOM panel.

CIMOM Connection Status
The status of this CIMOM with respect to Tivoli Storage Productivity Center. Possible values are: SUCCESS, UNCONFIGURED, UNKNOWN, INVALID_NAMESPACE, TIMEOUT, REFUSED, LOGIN_FAILED, SSL_HANDSHAKE_ERROR, SSL_REGISTRATION_INVALID, CIMCLIENT_ERROR

CIMOM Status Timestamp
The date and time when the **Connection Status** information was last collected.

CIMOM Interoperability Namespace
The interoperability namespace of a CIMOM. This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information. For a NetApp CIMOM, enter /interop.

CIMOM Username
The CIMOM user name used for authentication.

CIMOM Trustore Location
The location (path on this computer) of a certificate file for certificate based authentication in the https protocol.

CIMOM User Interface Description
The name of the Human Interface Service (if any) supported by this CIMOM.

CIMOM Software Level
The software version level of the CIMOM agent.

CIMOM Protocol Version
The version of the cim-xml protocol.

CIMOM Authentication Mechanism
The authentication mechanism supported by the CIMOM. This column can contain the following values: **Unknown**, **None**, **Other**, **Basic**, **Digest**.

CIMOM Alias
The alias of the CIMOM.

CIMOM Protocol
The protocol used to communicate with the CIMOM. Possible values are **http** and **https**.

CIMOM SLP Attributes
The standard set of attributes for this CIMOM. The attributes are retrieved via SLP.

Appendix A. Supported Fabric Agent Types

The following table shows the supported agent types for Switch and Fabric functions.

Table 23. Supported agent types for Switch and Fabric Functions

Function > Switch	Brocade	McDATA	Cisco[1]	QLogic
Switch Performance Monitoring	CIMOM agent	CIMOM agent	CIMOM agent	Not supported
Zone Control	Recommended: CIMOM agent Also supported: Out-of-band fabric agent	Recommended: CIMOM agent Also supported: In-band fabric agent	In-band fabric agent required in each VSAN	In-band fabric agent
Zone Control with Zone Aliases	CIMOM agent	Not supported	Not supported	Not supported
Switch and switch port information collected	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Out-of-band fabric agent, In-band fabric agent
Topology connectivity information collected	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Out-of-band fabric agent, In-band fabric agent	Out-of-band fabric agent, In-band fabric agent
Zoning information collected	Recommended: CIMOM agent Also supported: Out-of-band fabric agent	Recommended: CIMOM agent Also supported: In-band fabric agent	In-band fabric agent required in each VSAN	In-band fabric agent
Tivoli Storage Productivity CenterAlerts	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Out-of-band fabric agent, In-band fabric agent	Out-of-band fabric agent, In-band fabric agent
Hosts, endpoint devices, device-centric and host-centric information collected	In-band fabric agent	In-band fabric agent	In-band fabric agent	In-band fabric agent
Switch "Sensors and Events"	Out-of-band fabric agent	Out-of-band fabric agent	Out-of-band fabric agent	Out-of-band fabric agent

Table 23. Supported agent types for Switch and Fabric Functions (continued)

Function > Switch	Brocade	McDATA	Cisco[1]	QLogic
Notes: • The In-band fabric agent gets the information on a VSAN basis. • Each VSAN is viewed as an individual SAN. • The Out-of-band fabric agent and CIMOM agent gets the physical fabric information and can correlate the VSAN information to a physical infrastructure. • The Out-of-band fabric agent also collects some VSAN information.				

Appendix B. Frequently Asked Questions

Use this section to view answers to common questions about IBM Tivoli Storage Productivity Center. These questions are organized according to the manager to which they apply.

General and Migration Information

Learn about the issues when working with IBM Tivoli Storage Productivity Center, upgrading to a new version of the product, or migrating from another product.

What Java version is officially supported by Tivoli Storage Productivity Center version 4.1 graphical user interface?

Java 1.5 is the level supported for the Tivoli Storage Productivity Center GUI. For information about how to start the Tivoli Storage Productivity Center graphical user interface, see the information center. Search for **Starting and stopping the console**.

You are entering namespaces manually for CIM agents and switches in the Tivoli Storage Productivity Center user interface. What are the current namespaces for the supported brands?

Table 24. Interop Namespaces for CIM Agents for switches and storage subsystems

Switch or subsystem	Namespace
IBM	/root/ibm
Brocade	/interop
Cisco	/root/cimv2 For version 3.2.1 or later: /root/pg_interop
Engenio	/interop
EMC	/root/emc
HDS	For HiCommand 5.6 or later, use: /root/smis/current For a HiCommand version that supports SMI-S 1.2, use following namespace to traverse the model as SMI-S 1.1: /root/smis/smis11 For HiCommand versions earlier than HiCommand 5.6, use: /root/hitachi/dm*xx*, where *xx* is represents the level of HiCommand.
HP	/root
McData	/interop
SUN	/root/sun3510 or /interop Note: This namespace is for a subsystem and not a switch.
XYRATEX	/root/PG_interop

Check with your switch vendor to ensure you use the most current namespaces.

What if you specify an incorrect namespace?

If you specify an incorrect namespace the following issues might occur:

- The connection test fails when the CIMOM is added.
- The discovery does not discover all information of the system managed by the CIMOM.
- The probe fails.
- The function you want to perform on the system might fail (for example, collecting performance data).

What are the required steps when upgrading the Tivoli Storage Productivity Center to the latest version?

See the Tivoli Storage Productivity Center information center for detailed information about how to upgrade the product:http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp. To access information about upgrading, expand the **Tivoli Storage Productivity Center** > **Upgrading** node in the information center.

You are using an older version of IBM Tivoli Storage Productivity Center for Data. How can I upgrade to this latest version?

See the Tivoli Storage Productivity Center information center for detailed information about how to upgrade the product: http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp. To access information about upgrading, expand the **Tivoli Storage Productivity Center** > **Upgrading** node in the navigation tree.

Data Manager

Use the questions in this section to answer questions about working with the Data Manager component of IBM Tivoli Storage Productivity Center.

Keep in mind the following information and considerations:

How do you know if your storage subsystem is supported?

A set of Tivoli Storage Productivity Center for Data reports are supported for any disk array storage subsystem that is Conformance Testing Program (CTP) Certified by the Storage Networking Industry Association (SNIA). See the following link to view which disk array storage subsystems are supported: http://www.snia.org/ctp/conformingproviders

Follow the appropriate link for the vendor of your disk array storage subsystem. Find the information for SMI-S version 1.0. The supported disk array subsystem models are listed. The Tested Software Version is the SMI agent (also known as CIMOM) version that you need to have in your environment. Be sure that "Array" is listed as a tested Profile.

Tivoli Storage Productivity Center supports additional reports, beyond those enabled through SMIS, for certain disk array subsystems. These reports correlate information gathered from the agents running on your hosts, and thus allow reports that show relationships between computers, filesystems, and subsystem volumes. The list of subsystems with this level of support can be found on the IBM Support web site: http://www.ibm.com/systems/support/supportsite.wss/supportresources?taskind=3&brandind=5000033&familyind=5329731. Select the **Products Supported List** link for the appropriate system version under **Integration/Interoperability**.

Where do you get the CIMOM (also known as SMI agent) for my subsystem

For DS6x00 and DS8x00: http://www-1.ibm.com/support/docview.wss?uid=ssg1S4000595&rs=555

For Tivoli Storage Enterprise Storage Server: http://www-1.ibm.com/support/docview.wss?uid=ssg1S4000595&rs=555

For FAStT or DS4xxxx: http://www.engenio.com/products/smi_provider.html

For other vendors: Some storage vendors allow you to download the CIMOMs for subsystems from their web site. Other vendors do not make their CIMOMs available for download over the Internet. Check with your storage vendor to find out how to obtain their SMI-S 1.0.2 - Certified SMIS Provider.

Fabric Manager

Use the questions in this section to answer questions about working with the Fabric Manager component of IBM Tivoli Storage Productivity Center.

Keep in mind the following information and considerations:

What type of fabric agents should you be using in Tivoli Storage Productivity Center?

The following table shows the current agent types supported for the Switch Performance Management functionality and the Fabric Zone Configuration functionality.

Table 25. Supported agent types for Switch and Fabric Functions

Function > Switch	Brocade	McDATA	Cisco[1]	QLogic
Switch Performance Monitoring	CIMOM agent	CIMOM agent	CIMOM agent	Not supported
Zone Control	Recommended: CIMOM agent Also supported: Out-of-band fabric agent	Recommended: CIMOM agent Also supported: In-band fabric agent	In-band fabric agent required in each VSAN	In-band fabric agent
Zone Control with Zone Aliases	CIMOM agent	Not supported	Not supported	Not supported
Switch and switch port information collected	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Out-of-band fabric agent, In-band fabric agent
Topology connectivity information collected	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Out-of-band fabric agent, In-band fabric agent	Out-of-band fabric agent, In-band fabric agent
Zoning information collected	Recommended: CIMOM agent Also supported: Out-of-band fabric agent	Recommended: CIMOM agent Also supported: In-band fabric agent	In-band fabric agent required in each VSAN	In-band fabric agent

Table 25. Supported agent types for Switch and Fabric Functions (continued)

Function > Switch	Brocade	McDATA	Cisco[1]	QLogic
Tivoli Storage Productivity CenterAlerts	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Recommended: CIMOM agent Also supported: Out-of-band fabric agent, In-band fabric agent	Out-of-band fabric agent, In-band fabric agent	Out-of-band fabric agent, In-band fabric agent
Hosts, endpoint devices, device-centric and host-centric information collected	In-band fabric agent	In-band fabric agent	In-band fabric agent	In-band fabric agent
Switch ″Sensors and Events″	Out-of-band fabric agent	Out-of-band fabric agent	Out-of-band fabric agent	Out-of-band fabric agent

Notes:

- The In-band fabric agent gets the information on a VSAN basis.
- Each VSAN is viewed as an individual SAN.
- The Out-of-band fabric agent and CIMOM agent gets the physical fabric information and can correlate the VSAN information to a physical infrastructure.
- The Out-of-band fabric agent also collects some VSAN information.

What are the prerequisites for using the Switch Performance functionalities in Tivoli Storage Productivity Center?

The prerequisite steps for using the Switch Performance functionalities are:

1. Add a CIMOM as a data source to the Tivoli Storage Productivity Center system that manages the switch.
2. Run a CIMOM discovery and wait for it to complete successfully for that CIMOM.
3. Run a probe against the fabric and wait for it to complete successfully.

After you complete these steps, create a Switch Performance monitor to trigger defined switch performance alert conditions or generate switch performance reports.

Why does an Out of Band Fabric (SNMP) agent show a state of ″not contacted″?

Perform the following tasks to help determine why an Out of Band Fabric (SNMP) agent show a state of ″not contacted″:

- Check that the IP Address or name is entered correctly in the Out of Band Fabric agent details panel.
- Ensure the FA MIB is enabled on the switch.
- Ensure that the Tivoli Storage Productivity Center server IP address is included in any SNMP Access Control List that exists on the switches.
- Ensure you are using SNMP v1.
- Ensure the string on the Tivoli Storage Productivity Center Out of Band Fabric agent details panel in the SNMP Community field matches one of the community strings on the switch.

Can you add a storage subsystem as an Out of Band Fabric (SNMP) agent?

No. Out of Band Fabric agents within Tivoli Storage Productivity Center are meant for switches only. Disk Manager uses CIMOMs to gather subsystem information.

Why aren't all the properties of a switch populated in the Topology Switch views?

A combination of Inband and CIMOM agents or a combination of Inband and Out of Band fabric agents might be required to gather all topology information for a switch. However, all of the properties are not normally needed.

How can you specify a different management application to launch for a switch?

Launch the Element Manager for a switch from the topology viewer. To do this, access the Switch view, right click a switch, and select **Launch Element Manager**. To change what is launched, right click a switch and select the **Launch Details Panel**. The element manager information is displayed at the bottom of the Details tab. To change this information, press **Set** and enter the HTTP, Telnet, or application information into the pop-up dialog and press **OK**.

How do you manually change a fabric or switch label that is displayed throughout the system?

Access the topology viewer in the Fabric or Switch view as appropriate. Right-click the entity you want to label to and select **Launch Details Panel**. Within the Details Panel, change the data in the **Label** field and click the save icon (or press Ctrl S).

Why do you get the error "Agent capable of configuration zoning could not be found on this Fabric" for a fabric?

For Brocade switches you must have a CIMOM agent, an Out of Band Fabric agent, or both configured for the fabric. A CIMOM agent is used by default if present. If a CIMOM for the fabric is present, you must run a CIMOM discovery and probe of the fabric before you can perform zone control. If you use an Out of Band Fabric agent for Brocade zone control, you must first enable Advanced Brocade Discovery. To do this, enter a user ID and password on the Out of Band Fabric agent panel. This enables Tivoli Storage Productivity Center to gather zone information during Out of Band discovery. Note that because all switches in a Brocade fabric contain the same zone information, you only need to have **Advanced Brocade Discovery** enabled for a minimal number of switches within the fabric. Typically, this is one or two (for redundancy).

For McDATA switches, a CIMOM agent is used by default (if available). If a CIMOM agent is not available, an Inband Fabric agent is used. If a CIMOM agent for the fabric is present, you must run a CIMOM discovery and probe of the fabric before you can perform zone control.

For all other supported switches, you must have a Tivoli Storage Productivity Center Inband Fabric agent in the fabric.

Why is it not possible to configure inactive zoneset information for a McDATA fabric?

If no CIMOM agent is configured for the McDATA fabric, and only Inband Fabric Agents are configured, then the zone control feature is limited to configuring only the active zoneset. Configure a CIMOM agent for the McDATA fabric to work with inactive zonesets for this fabric.

Disk Manager

Use the questions in this section to answer questions about working with the Disk Manager component of IBM Tivoli Storage Productivity Center.

Keep in mind the following information and considerations:

You encounter errors while collecting performance data on SAN Volume Controller. You fail to associate SAN Volume Controller performance data from non-configuration node with SAN Volume Controller performance data from configuration node. You encounter incomplete SAN Volume Controller performance data sample.

This issue is caused by a configuration issue with SAN Volume Controller (time zone). Reset the time zone on SAN Volume Controller by logging into the SAN Volume Controller through putty. Run this command first:

```
svctask settimezone -timezone 509
```

This forces the cluster into the Universal time zone. To get you to whatever time zone you want the cluster to be in, run this command:

```
svctask settimezone -timezone
```

Click on Launch Element Manager. The URL that appears is http://0.0.0.0/ica

This error is caused by wrong configuration of the SAN Volume Controller ICAT. Use the SAN Volume Controller GUI to remove and add the cluster again to the ICAT. This populates the correct value for the access point which is stored in the database repository during the next discovery. Use this cli command to set the IP address and port of the user interface.

For example: `svcluster -x modify -p 9.155.62.91:9080`

This is the same command used by the user interface and CIMOM.

Does the DS API server contains multiple IP interfaces?

When a CIMOM server has multiple IP adapters, the Tivoli Storage Enterprise Storage Server can send the performance data to the wrong IP interface on that CIMOM server. This will cause that data to be lost. The adapter to which the data is sent depends on the IP configuration of the host server that is running the DS API. The DS API enables you to configure the preferred IP address to which performance data is sent. This is the IP address of the adapter on the same network as the Tivoli Storage Enterprise Storage Server.

In some revisions of the DS API, this parameter is ignored. This is fixed at the 5.1.0.51 level. Note that due to a number of related issues with the DS API, Tivoli Storage Enterprise Storage ServerCLI, and Tivoli Storage Productivity Center, you should open a software PMR to address this problem as updates may be required for all three components.

Must the file system where the DS API and Tivoli Storage Enterprise Storage Server CLI are installed have freespace?

Yes. The Tivoli Storage Enterprise Storage Server performance statistics are written by the Tivoli Storage Enterprise Storage Server CLI into the file system on a CIMOM server. If there is no space in the file system, the log files will be 0 bytes in size and no performance data will be received by the CIM agent and sent to Tivoli Storage Productivity Center.

Must the Tivoli Storage Enterprise Storage Server userid for DS API configuration have Tivoli Storage Enterprise Storage Server administrator rights?

Yes. If the ID used in the SET DEVICE command when the DS API is

configured does not have administrator rights, the CIMOM is unable to execute the esscli LIST PERFSTATS command and unable to collect performance data.

Must the Tivoli Storage Enterprise Storage Server Web Specialist InfoServer (running inside the Tivoli Storage Enterprise Storage Server machine) be running?

Yes. If the Tivoli Storage Enterprise Storage Server Web Specialist InfoServer is not running, performance statistics cannot be collected. To check if the Tivoli Storage Enterprise Storage Server Web Specialist InfoServer is running, start the Tivoli Storage Enterprise Storage Server Web Specialist user interface and access the Storage Allocation panel. Check if volumes and hosts can be displayed. If no information is returned, restart the Tivoli Storage Enterprise Storage Server Web Specialist InfoServer.

Should any firewalls between the Tivoli Storage Enterprise Storage Server and DS API server be configured to allow LIST PERFSTATS traffic?

For the CIMOM machine, you must open all IP ports above 1023 to receive performance data from the Tivoli Storage Enterprise Storage Server.

How do you diagnose Tivoli Storage Enterprise Storage Server performance collection issues?

Use Tivoli Storage Productivity Center to collect performance data from an Tivoli Storage Enterprise Storage Server. If no data is shown after more than two collection intervals, you can check the providerTrace.log in the /cimagent directory on the server where the DS API is installed.

If the log file contains long entries of performance data, then the CIMOM is correctly collecting performance data. If the CIMOM log does not show entries containing performance data, the problem might be with the esscli LIST PERFSTATS command used to collect performance data. If the log does contain performance data, the problem is with the Tivoli Storage Productivity Center.

- If the CIMOM is collecting data correctly, the tracePerfMgr.log file in the /Tivoli Storage Productivity Center/device/logs directory on the Tivoli Storage Productivity Center server might contain more information about the failure.
- If the CIMOM is not collecting performance data, use the ESSCLI LIST PERFSTATS command to determine if the Tivoli Storage Enterprise Storage Server is able to collect performance statistics: `esscli list PerfStats -d "ess=2105.nnnnn" -s  -u username - p password`

If time-stamped log files containing performance data are written to the local file system, the Tivoli Storage Enterprise Storage Server is in a good state to collect performance data. Otherwise, check for any of the previously described issues.

You can use Tivoli Storage Productivity Center to collect volume and subsystem information and configure volumes and volume assignments by using the DS API CIM agent, but you are unable to collect performance data.

Here is a list of known issues, hints, and tips on diagnosing Tivoli Storage Enterprise Storage Server performance collection problems. All the issues are due to the setup of the DS API environment and configuration of the Tivoli Storage Enterprise Storage Server:

- When a CIMOM server has multiple IP adapters, the Tivoli Storage Enterprise Storage Server can send the performance data to the wrong IP interface on that CIMOM server. This will cause that data to be lost. The

adapter to which the data is sent depends on the IP configuration of the host server that is running the DS API. The DS API enables you to configure the preferred IP address to which performance data is sent. This is the IP address of the adapter on the same network as the Tivoli Storage Enterprise Storage Server.

In some revisions of the DS API, this parameter is ignored. This is fixed at the 5.1.0.51 level. Note that due to a number of related issues with the DS API, ESSCLI, and Tivoli Storage Productivity Center, you should open a software PMR to address this problem as updates may be required for all three components.

- The file system where the DS API and Tivoli Storage Enterprise Storage Server CLI are installed must have freespace. The Tivoli Storage Enterprise Storage Server performance statistics are written by the Tivoli Storage Enterprise Storage Server CLI into the file system on the CIMOM server. If there is no space in the file system, the log files will be 0 bytes in size and no performance data will be received by the CIM agent and sent to Tivoli Storage Productivity Center.
- The Tivoli Storage Enterprise Storage Server user ID for DS API configuration must have Tivoli Storage Enterprise Storage Server administrator rights. If the ID used in the SET DEVICE command when the DS API is configured does not have administrator rights, the CIMOM is unable to execute the esscli LIST PERFSTATS command and unable to collect performance data.
- If the Tivoli Storage Enterprise Storage Server Web Specialist InfoServer is not running, performance statistics cannot be collected. To check if the Tivoli Storage Enterprise Storage Server Web Specialist InfoServer is running, start the Tivoli Storage Enterprise Storage Server Web Specialist user interface and access the Storage Allocation panel. Check if volumes and hosts can be displayed. If no information is returned, restart the Tivoli Storage Enterprise Storage Server Web Specialist InfoServer.
- Configure any firewalls between the Tivoli Storage Enterprise Storage Server and DS API server to allow LIST PERFSTATS traffic. For the CIMOM machine, you must open all IP ports above 1023 to receive performance data from the Tivoli Storage Enterprise Storage Server.

Appendix C. Script parameters

Script parameters provide specific information on the alert that triggered the script to be run.

The parameters passed to a script depends on the type of alert that was triggered by Data Manager.

The following table describes all the script parameters:

Script Parameter	Description
amount	Threshold exceeded amount.
archive-file-count	The number of log files residing in the archived log directory.
archive-log-directory	Name of the archive log directory that has triggered the archive log directory Instance alert.
available-extents	The number of extents still available to the segment for growth. This value equals the maximum extents available to the object minus the extents currently allocated to the segment.
available-space	Available pool space after a change
blade	Name of a blade.
chained-row-count	The number of chained rows in a table that has triggered the `Chained Row` table alert.
computer	Computer name where the triggering condition occurred.
consecutive-failures	Number of consecutive failed attempts to ping the computer.
controller	Name of a back-end controller.
cluster-name	The name of an HACMP or MSCS cluster.
creator.name	Creator of the ping, probe, or scan schedule. Name of the schedule.
current-grown-defects	Current number of grown defects on the disk.
current-node-name	When an HACMP or MSCS cluster resource group has moved, this parameter identifies the cluster node that now hosts the cluster resource group.
current-RAM MB	Current value of the RAM in megabytes.
current-VM MB	Current value of the sum of the RAM and the swap space in megabytes.
database	The name of the database where the triggering condition occurred.
database-tablespace	The name of the database or tablespace where the triggering condition occurred.
datapath	Name of a data path.
device-name	Name of a device.
disk-array	Name/alias of a disk array.
dump-date	The date when the last dump was performed.
endpoint	Name of an end point device.
extent-count	The number of extents allocated to a segment, or the number of free extents in the tablespace (depends on Alert type).
failed-jobs	Number of failed jobs in the run. (Each job runs on a different computer).
file-of-violating-files	Temporary file containing a list of files violating the constraint. This is listed as one file per line.

Script Parameter	Description
file-of-violating-owners	Temporary file containing a list of owners who owns the violating files.
free-inodes	Maximum number of files available to be created on this file system.
freespace size-designator	Total amount of freespace, in KB, MB, or GB.
from-entity-type	Type of new fabric connection from an entity.
HBA-driver	HBA driver
HBA-firmware	HBA firmware
io-group	Name of the I/O group.
largest-extent-size size-designator	Total amount of the largest free extent in the tablespace, measured in KB, MB, or GB.
manufacturer/serial-number	Manufacturer of the disk. Serial number of the disk.
mdisk	Name of an MDisk.
mdisk-group	Name of an MDisk group.
mount-point	Path to the file system.
new-capacity	New capacity of a storage subsystem, volume, or pool.
new-version	New version of the HBA driver, firmware, or a subsystem.
node	Name of a node.
old-capacity	Previous capacity of a storage subsystem, volume, or pool.
old-grown-defects	Previous number of grown defects on the disk.
old-node-name	When an HACMP or MSCS cluster resource group has moved, this parameter identifies the cluster node that previously hosted the cluster resource group.
old-RAM MB	Previous value of the RAM in megabytes.
old-version	Previous version of the HBA driver, firmware, or a subsystem.
old VM MB	Previous value of the sum of the RAM and the swap space in megabytes.
path	Path to the directory.
percent-of-capacity %	Percentage of capacity of the file system, database, or tablespace.
percent-of-table-size	The percentage of space allocated to a segment that is empty and unused (i.e., the percentage of space above the “high-water mark”). Available on the Empty Used Segment Space table alert.
percent-of-total-rows %	The percentage of table rows that are chained.
pool	Name of a storage pool.
port	Name of a port.
rdbms-instance-name	Oracle SID, SQL Server name, Sybase Server name, UDB Instance name
rdbms-type	Oracle, SQL Server, or Sybase
run-number	Number of the run.
segment	The name of the table segment that has triggered the alert.

Script Parameter	Description
segment-type	The type of segment that has triggered the alert. The possible types of segments include: • TABLE • TABLE PARTITION • TABLE SUBPARTITION • NESTED TABLE • CLUSTER • INDEX • INDEX PARTITION • INDEX SUBPARTITION • LOBINDEX • LOBSEGMENT • LOB PARTITION • LOB SUBPARTITION
storage-volume	Name of a storage volume
subsystem	Name of a storage subsystem
switch	Name of a switch
table	The name of the table that triggered the alert condition.
tablespace	The name of the tablespace that triggered the alert condition.
threshold	Value that you set for the triggering condition. If the value unit was specified as a %, then a % will follow this value.
threshold thr-designator	Value of the triggering condition, in KB, MB, or GB, or % (value units).
to-entity-type	Type of new fabric connection to an entity.
total-jobs	Total number of jobs in a run.
totaled-file-size size designator	Total amount of storage consumed by the archive log directory, measured in KB, MB, or GB.
usage size-designator	Value of used disk space, in KB, MB, or GB.
violating-file-count	Number of files which met the conditions defined in the constraint.
virtual-server-name	The name of an HACMP or MSCS cluster resource group.
zone	Name of a zone.
zoneset	Name of a zone set.
zone-alias	Name of a zone alias
zone-member	Name of a zone member

Appendix D. Arguments and window locations available with the inbound launch in context feature

This section describes the arguments and window locations available with the inbound launch in context feature.

Launch in context parameters

You can use the launch in context parameters, definitions, and examples.

The command format for Java Web Start is unique in that a parameter and its value are separated by an equal sign (=). Parameter=value pairs are separated by an ampersand (&).

Not all parameters can be used with all launch types. Where this is true, it is noted in the parameter definition.

Parameters

-user The user ID. The value must be a valid user ID to log on to the IBM Tivoli Storage Productivity Center server, for example, `-user admin`.

-passwd
The password. The value must be a valid password to log on to the Tivoli Storage Productivity Center server, for example, `-passwd MyPassword`

Note: This parameter is not supported in a Web Browser's URL when using Java Web Start to launch the Tivoli Storage Productivity Center user interface.

-encpasswd
The password encrypted. This is used if you choose to encrypt the password instead of entering the text password, for example, `-encpasswd 48735`.

Tip: To encrypt a password using Tivoli Storage Productivity Center (the preferred method), enter the following command in the Command Prompt window:

TPC-HOME`/data/server/tools/tpctool encrypt` *`text password`*

The encrypted password is returned.

Note: This parameter is not supported in a Web Browser's URL when using Java Web Start to launch the Tivoli Storage Productivity Center user interface.

-encmethod
The encryption method used to encrypt the password. The password can be encrypted using Tivoli Storage Productivity Center or Tivoli Enterprise Portal. The preferred method is to use Tivoli Storage Productivity Center to encrypt the password. If **encpasswd** is selected, you must also select **encmethod** with a value. The default is TPC. For example, `-encpasswd 48735 -encmethod TPC`

Note: This parameter is not supported in a Web Browser's URL when using Java Web Start to launch the Tivoli Storage Productivity Center user interface.

Tip: The Tivoli Enterprise Portal graphical user interface (GUI) does not start the Tivoli Storage Productivity Center GUI properly with a Tivoli Enterprise Portal-encrypted password if that encryption creates a string with spaces in it. Such a string is handled as multiple values unless it is enclosed in double quotation marks (″).

Tip: Encryption using Tivoli Enterprise Portal creates a string pattern. There are multiple string patterns that can result in the same decrypted value, so it is possible that different string patterns used as an encrypted password might result in same decrypted password.

-function

Specify one of the following values:

- The task to perform on the remote data server. See "Launch in context task parameters" for tasks and their specific parameters. For example, `-function disk.create_volume -subsystem_id <ID of the subsystem>`.
- The window name to open on the target system. The window name is mapped to the name of the node on the navigation tree. See "Available windows" on page 788 for a list of available windows. For example, `-function topology.computer`.

-server

The Tivoli Storage Productivity Center server to log on to. This argument specifies a single server. For example, `-server tivoli11:9549`.

-servers

Not used with the Web Browser. The Tivoli Storage Productivity Center server or servers to log on to, in the format server:port. You can specify multiple server ports separated by commas. In that case, Tivoli Storage Productivity Center will connect to the first server. When a new connection from the top menu is created, (**Connection** → **New Connection**) the server list box includes the available servers in the **Connection** dialog. For example, `-servers ″tivoli11:9549,wanda:9549,invader:9549″`

-help **Not used with the Web Browser.** The help argument is used alone in the Arguments text field. It is entered as `-help functions`. This argument returns a list of the windows representing nodes in the navigation tree that can be launched. Typing `-help` with no arguments will return an error message.

You can also enter **help** as a command in the Command Prompt window as `tpc.bat -help functions` for Windows, and `TPCD.sh -help functions` for AIX. Linux, or UNIX.

Launch in context task parameters

Use the launch in context feature task parameters to perform tasks.

Task parameters associated with actions taken on the Data server also use the `-function` parameter. For example:`-function data.filesystem_capacity_by_computer -computer_id ″<computer_id>″`.

Create a volume (disk.create_volume)

-subsystem_id
The ID of the system on which the volume will be created. The subsystem_id is the name of the system as seen in the IBM Tivoli Storage Productivity Center graphical user interface (GUI). This parameter is required.

-name The name of the volume. The name is not valid for IBM Tivoli Storage Enterprise Storage Server and is ignored. If more than one volume is to be created, this value is treated as a prefix.

-size The size of the volume.

-unit The units to be used while creating the volume. Permissible values are KB, MB, and GB. The default unit is MB.

-number
The number of volumes to be created.

SAN Planner (disk.san_planner)

-computer_id
The ID of the computer for which the plan is to be created. The computer_id is the name of the computer as seen in the Tivoli Storage Productivity Center GUI.

-subsystem_id
The ID of the system for which the plan is to be created. The subsystem_id is the name of the system as seen in the Tivoli Storage Productivity Center GUI Disk Manager.

-plan_id
The ID of the previously saved plan. When the computer_id and subsystem_id are provided, plan_id is ignored. When you open a plan in the SAN Planner using launch in context, you must enter the exact name of the plan. The value for **-plan_id** is case-sensitive and must match the name of the plan when it was created. For example, if you created a plan named **UserPlan1**, you must use **-plan_id UserPlan1** in the launch in context command.

Capacity report (data.filesystem_capacity_by_computer)

-computer_id
The ID of the computer for which the report is to be created. If this value is not specified, the Tivoli Storage Productivity Center GUI opens the default report panel for **Data Manager > Reporting > Capacity > FileSystem Capacity > By Computer**.

You must enter a value for **-computer_id** that matches the ID for that computer that is stored in the Tivoli Storage Productivity Center database repository. You will receive a warning message if the ID for the computer you enter in the command line does not match the ID stored in the data repository. To ensure that this report is displayed properly when using launch in context, make sure to enter an ID for the computer that matches the ID that is displayed for it through the Topology Viewer or appropriate report in the product's user interface.

Wasted space report (TPC.reports_data.wasted_space)

This report has no parameters. The Tivoli Storage Productivity Center GUI opens the default report panel for **Data Manager > Reporting > System Reports > Data > Wasted Space**.

Available windows

You can use windows with the inbound launch in context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the `-function` *`window name`* argument.

Available window names for function argument

Table 26. Available windows for inbound launch in context

Tree Node	Window Name
Administrative Services->Data Sources->CIMOM Agents	datasource.cimom
Administrative Services->Data Sources->Data Agents	datasource.data
Administrative Services->Data Sources->Inband fabric Agents	datasource.inband_fabric
Administrative Services->Data Sources->Out of band fabric Agents	datasource.out_of_band_fabric
Administrative Services->Data Sources->IBM Tivoli Storage Productivity Center Servers	datasource.tpc_servers
Administrative Services->Data Sources->VMWare VI data source	datasource.vmware
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->All	alert.all
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Alerts Directed to user	alert.alerts_directed_to_user
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Storage Subsystem	alert.storage_subsystem
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Computer	alert.computer
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Disk	alert.disk
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Filesystem	alert.filesystem
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Directory	alert.directory
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->User	alert.user
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->OS User Group	alert.os_user_group
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Fabric	alert.fabric
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Switch	alert.switch
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Endpoint Device	alert.endpoint_device

Table 26. Available windows for inbound launch in context (continued)

Tree Node	Window Name
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->External	alert.external
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Tape Library	alert.tape_library
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Configuration Analysis	alert.configuration_analysis
IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Hypervisor	alert.hypervisor
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Disk Space Summary	data_report.disk_capacity
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Storage Access Times	data_report.storage_access_times
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Most Obsolete Files	data_report.most_obsolete_files
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Storage Availability	data_report.storage_availibility
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Storage Capacity	data_report.storage_capacity
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Total Freespace	data_report.total_freespace
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data-> User Space Usage	data_report.user_space_usage
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Disk->Port Performance	disk_report.port_performance
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Disk->Subsystem Performance	disk_report.subsystem_performance
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Disk->Top Volumes Data Rate Performance	disk_report.top_volumes_data_rate_performance
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Disk-> Top Volumes I/O Rate Perf	disk_report.top_volumes_io_rate_performance
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Fabric-> San Assets (ALL)	fabric_report.san_assets
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Fabric ->Port Connections	fabric_report.port_connections
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Fabric -> Switch Performance	fabric_report.switch_performance
IBM Tivoli Storage Productivity Center->My Reports->System Reports->Fabric ->Top Switch Ports Data Rate Performance	fabric_report.top_switch_ports_data_rate_performance
IBM Tivoli Storage Productivity Center->Topology->Computers	topology.computers
IBM Tivoli Storage Productivity Center->Topology->Fabrics	topology.fabrics
IBM Tivoli Storage Productivity Center->Topology->Switches	topology.switches

Table 26. Available windows for inbound launch in context (continued)

Tree Node	Window Name
IBM Tivoli Storage Productivity Center->Topology->Storage	topology.storage
IBM Tivoli Storage Productivity Center->Topology->Storage resource groups	topology.storage_resource_groups
IBM Tivoli Storage Productivity Center->Topology->Other	topology.other
IBM Tivoli Storage Productivity Center->Configuration Utility	config
Disk Manager->Storage Subsystems	disk.storage_subsystems
Fabric Manager->Fabrics	fabric.fabrics
Tape Manager->Tape Libraries	tape.tape_libraries
Windows that correspond to launch in context tasks	
Fabric Manager -> SAN Planner	disk.san_planner
Disk Manager -> Storage Subsystems	disk.create_volume
Data Manager-> Reporting -> Capacity -> Filesystem Capacity -> By Computer	data.filesystem_capacity_by_computer
IBM Tivoli Storage Productivity Center -> My Reports -> System Reports -> Data -> Wasted Space	TPC.reports.data.wasted_space

Appendix E. Triggering conditions for alerts

You can define alerts so that IBM Tivoli Storage Productivity Center notifies you when a specified event occurs. Such events are the *triggering conditions* for the alert. The specific triggering conditions that you can use to define an alert depend on the sort of storage resource that you are monitoring.

Some triggering conditions, like thresholds, require you to enter values. If input is not required for the triggering condition you select, the value fields will be greyed out. For information about thresholds, see "Performance thresholds" on page 821

Thresholds

Thresholds are a type of triggering condition with which you can monitor a component with user-defined values.

You can monitor the performance of your enterprise by creating alerts on performance thresholds for switches and storage subsystems. By creating alerts that are triggered by performance thresholds, you can be informed about performance issues in your enterprise. Threshold events tell you when a component has fallen outside of the user-defined values. For example, when a threshold value has reached critical stress.

Threshold boundaries

You can establish your boundaries for the normal expected subsystem performance when defining storage subsystem alerts for performance threshold events. When the collected performance data samples fall outside out of the range you have set, you are notified of this threshold violation so you are aware of the potential problem. The upper boundaries are **Critical Stress** and **Warning Stress**. The lower boundaries are **Warning Idle** and **Critical Idle**. Usually you will want the stress boundaries to be high numbers and the idle to be low numbers. The exception to this rule is **Cache Holding Time Threshold**, where you want the stress numbers to be low and the idle numbers to be high.

If you do not want to be notified of threshold violations for any boundaries, you can leave the boundary field blank and the performance data will not be checked against any value. For example, if the **Critical Idle** and **Warning Idle** fields are left blank, no alerts will be sent for any idle conditions.

The **Suppress alerts when sequential I/O exceeds** check button is active only for the triggering condition **Disk Utilization Percentage Threshold**. It is a filter condition. The default is 80%.

Setting the thresholds

Only a few thresholds have defaults and on the other thresholds you will have to determine the best values for stress, idle, critical, and warning values so you can derive the maximum benefit without generating too many false alerts. Since suitable stress thresholds are highly dependant on the type of workload you are running, your exact hardware configuration, the number of physical disks, exact model numbers, and so forth, there are no easy or standard default rules.

One of the best approaches is to monitor your performance for a number of weeks and, using this historical data, determine reasonable values for each threshold setting. Once that is done you can fine tune these settings to minimize number of false alerts.

Controller thresholds

The following table lists and describes the Controller thresholds.

Table 27. Controller thresholds

Threshold (Metric)	Device/Component Type	Description
Controller Thresholds		
Total I/O Rate (overall)	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 controller SAN Volume Controller I/O group	Sets threshold on the average number of I/O operations per second for read and write operations, for the subsystem controllers (clusters) or I/O groups. The Total I/O Rate metric for each controller or I/O group is checked against the threshold boundaries for each collection interval. These thresholds are disabled by default.
Total Data Rate	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 controller SAN Volume Controller I/O group	Sets threshold on the average number of MB per second for read and write operations for the subsystem controllers (clusters) or I/O groups. The Total Data Rate metric for each controller or I/O group is checked against the threshold boundaries for each collection interval. These thresholds are disabled by default.
NVS Full Percentage	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 controller	Sets thresholds on the percentage of time that NVS space constraints caused I/O operations to be delayed, for the subsystem controllers (clusters). The NVS Full Percentage metric for each controller is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries of 10, 3, -1, -1.
Cache Holding Time	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 controller	Sets thresholds on the average cache holding time, in seconds, for I/O data in the subsystem controllers (clusters). Shorter time periods indicate adverse performance. The Cache Holding Time metric for each controller is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries of 30, 60, -1, -1.

Port thresholds

The following table lists and describes the Port thresholds.

Table 28. Port thresholds

Threshold (Metric)	Device/Component Type	Description
Port Thresholds		

Table 28. Port thresholds (continued)

Threshold (Metric)	Device/Component Type	Description
Total Port IO Rate	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 port, switch port	Sets thresholds on the average number of I/O operations or packets per second for send and receive operations, for the ports. The Total I/O Rate metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Total Port Data Rate	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 port, switch port	Sets thresholds on the average number of MB per second for send and receive operations, for the ports. The Total Data Rate metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Overall Port Response Time	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 port	Sets thresholds on the average number of milliseconds that it took to service each I/O operation (send and receive) for ports. The Total Response Time metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Error Frame Rate	Switch port	Sets thresholds on the average number of frames per second received in error for the switch ports. The Error Frame Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default.
Link Failure Rate	Switch port	Sets thresholds on the average number of link errors per second experienced by the switch ports. The Link Failure Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default.
Total Port Data Rate	Switch port	Sets thresholds on the critical and warning data rates for stress and idle in MB per second. The Total Port Data Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default.
Total Port Packet Rate	Switch port	Sets thresholds on the critical and warning data rates for stress and idle in packets per second. The Total Port Packet Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default.

Array thresholds

The following table lists and describes the Array thresholds.

Table 29. Array thresholds

Threshold (Metric)	Device/Component Type	Description
Array Thresholds		

Table 29. Array thresholds (continued)

Threshold (Metric)	Device/Component Type	Description
Disk Utilization Percentage	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 array	Sets thresholds on the approximate utilization percentage of the arrays in a particular subsystem; for example, the average percentage of time that the disks associated with the array were busy. The Disk Utilization metric for each array is checked against the threshold boundaries for each collection interval. This threshold is enabled by default for Tivoli Storage Enterprise Storage Server systems and disabled by default for others. The default threshold boundaries are 80%, 50%, -1, -1. For DS6000 and DS8000 subsystems, this threshold applies only to those ranks which are the only ranks in their associated extent pool.
Total Back-end I/O Rate	SAN Volume Controller MDisk group	Sets thresholds on the average number of I/O operations per second for MDisk read and write operations for the MDisk groups. The Total I/O Rate metric for each MDisk group is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Total Back-end Data Rate	SAN Volume Controller MDisk group	Sets thresholds on the average number of MB per second that were transferred for MDisk read and write operations for the MDisk groups. The Total Data Rate metric for each MDisk group is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Overall Back-end Response Time	SAN Volume Controller MDisk	Sets thresholds on the average number of milliseconds that it took to service each MDisk I/O operation, measured at the MDisk level. The Total Response Time (external) metric for each MDisk is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.

Triggering conditions for computer alerts

This topic lists the events that can trigger computer alerts.

Table 30. Triggering conditions for computer alerts

Triggering Condition	Explanation
RAM Increased	The amount of RAM on a managed computer is increased.
RAM Decreased	The amount of RAM on a managed computer is decreased.
Virtual Memory Increased	The amount of virtual memory on a managed computer is increased.
Virtual Memory Decreased	The amount of virtual memory on a managed computer is decreased.
New Disk Detected	A new disk is discovered on a managed computer.

Table 30. Triggering conditions for computer alerts (continued)

Triggering Condition	Explanation
Disk Not Found	A disk is removed from a managed computer.
New Disk Defect Found	A new grown disk defect is detected on a managed disk.
Grown Disk Defects Exceed *n* value	The grown disk defect threshold is exceeded. Enter a value for *n*.
Disk Failure Predicted	A managed disk predicts that a disk failure is imminent.
New Filesystem Detected	A new file system is discovered on a managed computer.
Cluster Resource Group Added	A new cluster resource group is added.
Cluster Resource Group Removed	A cluster resource group is removed.
Cluster Resource Group Moved	A cluster resource group is moved.
Computer Unreachable	A managed computer cannot be reached.
Computer Discovered	A new unmanaged computer is discovered.
Computer Status Change Offline	A managed computer goes offline.
Computer Status Change Online	A managed computer comes online.
Computer Property Change	An HBA or HBA to computer node is newly discovered, missing, or rediscovered.
HBA Driver Version Change	The driver version of host bus adapter changes.
HBA Firmware Version Change	The firmware version of host bus adapter changes.

Triggering conditions for file system alerts

This topic lists the events that can trigger file system alerts.

Table 31. Triggering conditions for file system alerts

Triggering Condition	Explanation
Filesystem Not Found	A file system is removed or unmounted from a computer.
Filesystem Reconfigured	The physical space definition of a managed file system is reconfigured.
Filesystem Freespace Less Than	A file system is low on free space. Enter a value, usually as a percent.
Filesystem UNIX FS Inodes Less Than	A UNIX file system is low on free inodes. Enter a value, usually as a percent.

Triggering conditions for directory alerts

This topic lists the events that can trigger directory alerts.

Table 32. Triggering conditions for directory alerts

Triggering Condition	Explanation
Directory Not Found	A directory cannot be found.
Directory Consumes More Than	A user or directory storage quota is exceeded. Enter a value, usually in Megabytes.

Triggering conditions for instance alerts

This topic lists the events that can trigger instance alerts.

Table 33. Triggering conditions for instance alerts

Database Application	Triggering Condition	Explanation
All	New Database - Tablespace Discovered	A new database-tablespace is discovered.
Oracle	Archive Log Directory Contains More Than	An archived log directory quota is exceeded. Enter a value for quota.
	New Tablespace Discovered	A new tablespace is discovered.
SQL Server	New Database Discovered	A new database is discovered.
Sybase	Total Device Freespace Less Than	The total amount of device free space is less than a specified amount.
	Total Device Freespace Greater Than	The total amount of device free space is greater than a specified amount.
	New Device Discovered	A new device is discovered.
	Device Dropped	A device is dropped.
	New Database Discovered	A new database is discovered.
DB2	New Tablespace Discovered	A new database-tablespace is discovered.

Triggering conditions for database-tablespace alerts

This topic lists the events that can trigger database-tablespace alerts.

Table 34. Triggering conditions for database-tablespace alerts

Database application	Triggering Condition	Explanation
All	Database - Tablespace Freespace Less Than	The free space on a database-tablespace is less than a specified amount.
	Database - Tablespace Offline	A database-tablespace is taken offline.
	Database - Tablespace Dropped	A database-tablespace is dropped.
Oracle	Tablespace Freespace Less Than	The free space on a tablespace is less than a specified amount.
	Freespace Fragmented into More Than	The free space on a tablespace is fragmented into more than the specified number of extents.
	Largest Free Extent Less Than	The largest free extent is less than the specified size.
	Tablespace Offline	A tablespace is taken offline
	Tablespace Dropped	A tablespace is dropped.
SQL Server	Database Data Freespace Less Than	The free space on a database is less than a specified amount.
	Database Log Freespace Less Than	The free space on a database log is less than the specified percent.
	Database Offline	A database is taken offline.
	Database Dropped	A database is dropped.

Table 34. Triggering conditions for database-tablespace alerts (continued)

Database application	Triggering Condition	Explanation
Sybase	Database Data Freespace less Than	The free space on a database is less than a specified amount.
	Database Log Freespace Less Than	The free space on a database log is less than the specified percent.
	Last Dump Time Less Than	The last dump time was the less than the specified number of days.
	Database Offline	A database is taken offline.
	Database Dropped	A database is dropped
DB2	Tablespace Freespace Less Than	The free space on a tablespace is less than a specified amount.
	Tablespace Offline	A tablespace is taken offline.
	Tablespace Dropped	A tablespace is dropped

Triggering conditions for table alerts

This topic lists the events that can trigger table alerts.

Table 35. Triggering conditions for table alerts

Database application	Triggering Condition	Explanation
All	Total Table Size Greater Than	The total size of a table is greater than a specified size.
	Table Dropped	A table is dropped.
Oracle	Max Extents - Allocated Extents	A table segment is near its maximum allowed extents.
	Segment Has More Than	A table segment has taken too many extents.
	Chained Row Count Greater Than	A table has too many chained rows.
	Empty Used Segment Space Exceeds	An empty used segment space exceeds the specified size.
	Total Table Size Greater Than	The total size of a table is greater than a specified size.
	Table Dropped	A table is dropped.
SQL Server	Total Table Size Greater Than	The total size of a table is greater than a specified size.
	Table Dropped	A table is dropped.
Sybase	Forwarded Row Count Greater Than	A forwarded row count is greater than the specified percent.
	Total Table Size Greater Than	The total size of a table is greater than a specified size.
	Table Dropped	A table is dropped.
DB2	Overflow Row Count Greater Than	An overflow row count is greater than the specified percent.
	Total Table Size Greater Than	The total size of a table is greater than a specified size.
	Table Dropped	A table is dropped.

Triggering conditions for storage subsystem alerts

This topic lists the events that can trigger storage subsystem alerts.

Table 36. Triggering conditions for storage-subsystem alerts

Triggering Condition	Explanation
New Storage Subsystem discovered	A new storage subsystem is discovered.
Cache Increased	A subsystem cache increases.
Cache Decreased	A subsystem cache decreases
New Disk Detected	A new disk is detected.
Disk Not Found	A disk is not found.
Storage Subsystem Not Found	A storage subsystem is not found.
Subsystem Status Change Offline	A storage subsystem goes offline.
Subsystem Status Change Online	A storage subsystem comes online.
Subsystem Version Change	A subsystem version changes.
Subsystem Property Change	A port on a storage subsystem is newly discovered, missing, or rediscovered.
Subsystem Allocated Capacity Change	The allocated capacity of a subsystem changes.
Subsystem Available Capacity Change	The available capacity of a subsystem changes.
Subsystem Backend Capacity Change	The allocated capacity of a subsystem changes.
Backend Controller Status Change Offline	A back-end controller goes offline.
Backend Controller Status Change Online	A back-end controller comes online.
Volume Status Change Offline	A volume goes offline.
Volume Status Change Online	A volume comes online.
Volume Capacity Change	The capacity of a volume changes.
Pool Status Change	A pool is missing or rediscovered.
Pool Discovered	A pool is discovered.
Pool Status Change Offline	A pool goes offline.
Pool Status Change Online	A pool comes online.
Pool Capacity Change	The capacity of a pool changes.
Pool Available Space Change	The available space in a pool changes.
Node Status Change Offline	A node goes offline.
Node Status Change Online	A node comes online.
Node State Change	An IBM SAN Volume Controller node is missing or is rediscovered.
Node Discovered	An IBM SAN Volume Controller node is discovered.
Datapath State Change	The state of a data path changes.
Datapath Discovered	A data path is discovered.
Overall Port Response Time Threshold	An overall port response time threshold is reached.
Total Port Data Rate Threshold	A total port data rate threshold is reached.
Total Port IO Rate Threshold	A total port I/O rate threshold is reached.
Cache Holding Time Threshold	A cache holding time threshold is reached.

Table 36. Triggering conditions for storage-subsystem alerts (continued)

Triggering Condition	Explanation
NVS Full Percentage Threshold	An NVS full percentage threshold is reached.
Total Data Rate Threshold	A total data rate threshold is reached.
Total I/O Rate Threshold	A total I/O rate threshold is reached.
Overall Backend Response Time Threshold	An overall back-end response time threshold is reached.
Total Backend Data Rate Threshold	A total back-end data rate threshold is reached.
Total Backend I/O Rate Threshold	A total back-end I/O rate threshold is reached.
Disk Utilization Percentage Threshold	A disk utilization percentage threshold is reached.

Related tasks

"Creating a performance threshold within an alert" on page 98
You can monitor the performance of your enterprise by creating alerts with performance thresholds for storage subsystems and SAN switches. By defining alerts that are triggered by performance thresholds, you can get timely information about performance issues in your enterprise.

Related reference

"Performance thresholds" on page 821
Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

Triggering conditions for fabric alerts

This topic lists the events that can trigger fabric alerts.

Table 37. Triggering conditions for fabric alerts

Triggering Condition	Explanation
Fabric Discovered	A fabric is discovered.
Fabric State Change	A fabric is missing or is rediscovered.
Fabric Status Change Offline	A fabric goes offline.
Fabric Status Change Online	A fabric comes online.
Fabric to Switch Change	A switch to fabric association is discovered, rediscovered, or missing.
Zone State Change	A zone is missing or is rediscovered.
Zone Discovered	A zone is discovered.
Zone Set State Change	A zone set in a fabric was activated, deactivated, or has gone missing.
Zone Set Discovered	A zone set is discovered.
Zone to Zone Alias Change	A zone to zone alias association is discovered, missing, or rediscovered.
Zone to Zone Member Change	A zone to zone member association is discovered, is missing, or is rediscovered.
Zone Set to Zone Change	A zone set to zone association is discovered, is missing, or is rediscovered.
New Connection	A new connection is detected.
Connection State Change	A connection is missing or rediscovered.
Zone Alias to Member Change	A zone alias to member association is discovered, rediscovered, or is missing.

Triggering conditions for switch alerts

This topic lists the events that can trigger switch alerts

Note: The following triggering conditions are performance thresholds. They trigger alerts only if performance monitors are run on the switch:

- Total Port Data Rate Threshold
- Link Failure Rate Threshold
- Error Frame Rate Threshold
- Total Port Packet Rate Threshold

Table 38. Triggering conditions for switch alerts

Triggering Condition	Explanation
Switch Discovered	A switch is discovered.
Switch State Change	A switch is missing or is rediscovered.
Switch Property Changes	A switch port is discovered, is missing, or is rediscovered.
Switch Status Change Offline	The status of the switch has changed from OK to a lower status. This does not necessarily mean that the switch has gone missing.
Switch Status Change Online	The status of the switch status of the switch has been upgraded to OK.
Switch Version Change	The version of a switch changes.
Switch to Port Change	A switch to switch port association is discovered, is missing, or is rediscovered.
Switch Blade Change	A switch module is discovered, is missing, or is rediscovered.
Switch Blade Change Offline	A switch module goes offline.
Switch Blade Change Online	A switch module comes online.
Total Port Data Rate Threshold	A total port data rate threshold is reached.
Link Failure Rate Threshold	A link failure rate threshold is reached.
Error Frame Rate Threshold	An error frame rate threshold is reached.
Total Port Packet Rate Threshold	A total port packet rate threshold is reached.

Triggering conditions for endpoint device alerts

This topic lists the events that can trigger endpoint device alerts.

Table 39. Triggering conditions for endpoint device alerts

Triggering Condition	Explanation
Endpoint Discovered	An endpoint device is discovered.
Endpoint State Change	An endpoint device is missing or is rediscovered.
Endpoint to Node Change	A peripheral entity to node association is missing, newly discovered, or rediscovered.
Endpoint Version Change	The version for an endpoint device changes.

Appendix F. Triggered actions for alerts

You define triggered actions to occur as a result of the alert.

An alert is always entered in the error log file each time it occurs. In addition, you can define the following notifications or triggered actions:

SNMP Trap
Generates an SNMP trap message to any network-management system, console, or terminal. You must set up the SNMP trap ringer with the `fabric.mib` and `tivoliSRM.mib` files.

TEC Event
Sends an alert to the IBM Tivoli Enterprise Console. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze theIBM Tivoli Storage Productivity Center events. Tivoli Enterprise Console also performs responses such as sending further notifications, creating or updating trouble tickets, and running programs. You must load the `tivoliSRM.baroc` and `fabric.baroc` files into the active rulebase of the Tivoli Enterprise Console server.

Login Notification
Tivoli Storage Productivity Center lists the alerts for the specified users when they log in. Use the **Preferences → Edit General** option on the Tivoli Storage Productivity Center menu bar to control which alerts a user will view upon logging in to the product.

Windows Event Log, UNIX Syslog
Writes the alerts to the operating-system log. You can select from the following levels of operating-system events:

- UNIX: Informational, Notice, Warning, Error, Critical, or Alert
- Microsoft Windows: Error, Warning, and Information.

For UNIX, you also can select User or Local.

Run Script
Runs a script in response to the alert. This enables you to use third-party tools for such actions as recovery and provisioning.

1. Click **Define** to specify the script that is run when the triggering condition occurs. The **Specify Script** window opens.
2. Type the name of the script in the **Script Name** field.
3. Select where to run the script.

If parameters are given for the triggering condition, they will be listed in the **Specify Script** window. These parameters will be used when the alert is triggered.

Note: Scripts can only be run against systems on which a Data agent is deployed.

Archive/Backup
(Constraints only) Runs an IBM System Storage Productivity Center backup or archive job in response to an alert associated with a constraint. Click the Define button.

Email Sends an e-mail. Use this section to view, add, and delete e-mail addresses to send the e-mail to, as well as edit the text of the message itself.

For **Directory Alerts**, if the **Directory Consumes More Than** triggering condition is selected, additional fields are available to send e-mail to the quota violator. The e-mail address can be based on **Directory Owner** or **Directory Name**.

Appendix G. Alert display colors

Status symbols for alerts are shown in the colors listed for the alert action.

When you have active alerts shown within the subnodes under the **IBM Tivoli Storage Productivity Center-> Alering -> Alert Log** node, the indicator next to those subnodes will appear as green or red depending on their corresponding alerts. The following table shows the colors that are generated for the subnodes by their different alerts:

Table 40. Alert colors

Alert	Color
Probe failed	Red
RAM increase	Green
RAM decrease	Red
Memory increase	Green
Memory decrease	Red
New disk	Green
Missing disk	Red
Disk defect	Red
Disk failure predicted	Red
New filesystem	Green
Filesystem reconfigured	Red
Scan failed	Red
Missing filesystem	Red
Filesystem freespace	Red
Filesystem free Inodes	Red
Computer unreachable	Red
New directory	Red
Missing directory	Red
Directory quota violation	Red
Filesystem quota violation	Red
Computer quota violation	Red
Network quota violation	Red
Constraint violation	Red
Failed job	Red
Disk defect quota	Red
Ping failed	Red
Discovery failed	Red
Aggregator failed	Red
Filesystem quota failed	Red
Computer quota failed	Red

Table 40. Alert colors (continued)

Alert	Color
Network quota failed	Red
Computer discovered	Red
Report failed	Red
DB probe failed	Red
DB tablespace quota violation	Red
DB instance quota violation	Red
DB net quota violation	Red
DB tablespace quota failed	Red
DB instance quota failed	Red
DB net quota failed	Red
Scan failed	Red
Archive log overflow	Red
New tablespace	Red
Tablespace free space	Red
Tablespace offline	Red
Tablespace fragmented	Red
Tablespace extent size	Red
Dropped tablespace	Red
Table fragmented	Red
Table maximum extents	Red
Table size	Red
Table chained rows	Red
Table wasted space	Red
Dropped table	Red
NAS discovery failed	Red
NAS computer discovered	Green
Script failed	Red
NetApp quota job failed	Red
Device freespace less	Red
Device freespace greater	Red
Dropped device	Red
New device	Red
DB log freespace	Red
DB dump time	Red
Disk array discovered	Green
Disk array missing	Red
Filer missing	Red
Cache increase	Green
Cache decrease	Red
Automatic filesystem extension	Red

Table 40. Alert colors (continued)

Alert	Color
Stopped automatic filesystem extension	Red
Extend job failed	Red
Upgrade failed	Red
Added cluster resource group	Green
Removed cluster resource group	Green
Moved cluster resource group	Green
DB report failed	Red
File backup failed	Red
Hypervisor discovered	Green
Hypervisor missing	Red
Virtual machine discovered	Green
Virtual machine missing	Red

Appendix H. Reference information

This reference information describes configuration files, log files, silent installation files, and other files.

agent.config file

The **agent.config** file contains configuration parameters for the Data agent. These parameters are set when the Data agent is installed; they can also be changed manually by modifying the file.

The following table contains the parameters for the **agent.config** file. If the Data agent is installed in the default location, this file is located at either `/opt/IBM/TPC/ca/subagents/TPC/Data/config` or `C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\config`.

Table 41. Parameters for the **agent.config** *file*

Parameter	Description
agentPort	Port on which the Data agent listens for requests. By default, this is set to 9510.
serverHost	Fully qualified host name of the system on which the Data server is installed.
serverPort	Port on which the Data server listens for requests. By default, this is set to 9549.
logFilesKept	Maximum number of Data agent logs that are retained. When this number is reached, the oldest log file is overwritten. By default, this is set to five.
messagesPerLog	Maximum number of messages in a Data agent log file. When this number is reached, the a new log file is created. By default, this is set to 100,000.
maxBacklog	Maximum number of uncompleted jobs that are permitted. When this number is reached, if additional job requests are made, any error is generated. By default, this is set to 500.
sendFailWait	Number of seconds to wait before the Data agent attempts to resend a message to the Data server. By default, this is set to 30.
maxIdleThreads	Maximum number of idle threads to retain for use by future jobs. By default, this is set to 10.
uptimePoll	How often (in seconds) should agent check to ensure it is up (20).
hostAlias	This parameter appears if the `HOST_ALIAS` is not specific and represents the name of the server. This value for this parameter is used when multiple computers have the same name or the name cannot be determined.
honorSentScripts	If this parameter is set to '1', 't', 'T', 'y', or 'Y', the Data agent can run scripts sent from the Data server. Otherwise, only scripts that are exist in the **scripts** directory on the system where the Data agent is installed can be run.
TPCInstallLocation	Directory where the Data agent is installed.

Information gathered by the agents

This section provides information about what data is collected by the agents.

Table 42. Information gathered by IBM Tivoli Storage Productivity Center for Fabric

Event which triggers discovery	In-band fabric agent	Out-of-band fabric agents	CIMOM fabric agents (Brocade or McData)	CIMOM fabric agents (QLogic or Cisco)
Administrative Services->Discovery->Out-of-band fabric	—	1. Finds switches within an IP address range 2. Gets fabric, switch, and topology information 3. Gets Cisco VSANs 4. Gets Brocade zoning data[1]	—	—
Administrative Services->Discovery->CIMOM	—	—	1. Discovers CIMOMs through SLP 2. Gets switches 3. Gets fabrics	1. Discovers CIMOMs through SLP 2. Gets switches 3. Gets fabrics
IBM Tivoli Storage Productivity Center ->Probe (including fabric)	1. Gets zoning information, excluding zone aliases[3] 2. Gets fabric, switch, and topology information 3. Hosts, endpoint devices, device-centric, and host-centric information collected	1. Gets fabric, switch, and topology information 2. Gets Cisco VSANs 3. Gets Brocade zoning data[1]	1. Gets zoning data, including zone aliases for Brocade 2. Gets fabric, switch, and switch port information, including blades[2] 3. Gets fabric topology connectivity data 4. Subscribes to CIM indications	1. Gets fabric, switch, and switch port information

Table 42. Information gathered by IBM Tivoli Storage Productivity Center for Fabric (continued)

Event which triggers discovery	In-band fabric agent	Out-of-band fabric agents	CIMOM fabric agents (Brocade or McData)	CIMOM fabric agents (QLogic or Cisco)
In-band fabric agent starts, or in-band agent connectivity changes, or in-band fabric change event detected	1. Gets zoning information, excluding zone aliases[3] 2. Gets fabric, switch, and topology information 3. Hosts, endpoint devices, device-centric and host-centric information collected	1. Gets fabric, switch, and topology information 2. Gets Cisco VSANs 3. Gets Brocade zoning data[1]	—	—
SNMP trap received from switch	1. Gets zoning information, excluding zone aliases[3] 2. Gets fabric, switch, and topology information 3. Hosts, endpoint devices, device-centric and host-centric information collected	1. Gets switch and topology information 2. Gets Cisco VSANs 3. Gets Brocade zoning data[1]	—	—

Table 42. Information gathered by IBM Tivoli Storage Productivity Center for Fabric (continued)

Event which triggers discovery	In-band fabric agent	Out-of-band fabric agents	CIMOM fabric agents (Brocade or McData)	CIMOM fabric agents (QLogic or Cisco)
CIM indication received from fabric CIMOM	—	—	Performs ″mini-probe″ to collect information relevant to the indication received, or sets entities as missing. Can do one of the following: • Get switches and fabrics • Gets zoning data • Gets switch port status and connection to node • Gets switch blade status and all associated port connections • Sets fabric, switch, blades, connections, and/or nodes as missing	—

1. Brocade zoning data is only retrieved from out-of-band fabric agents if no CIMOM is configured for this fabric. If an out-of-band fabric agent is used, Zone Aliases changes cannot be made through Tivoli Storage Productivity Center
2. Physical switch and blades information is not collected for McDATA i10000 switches (and same switch models from OEM partners). Virtual switches are collected.
3. Zoning data is only retrieved from in-band fabric agents for non-Brocade switches. For McDATA, if a CIMOM is configured for this fabric, zoning data is collected from the CIMOM.

Interop namespaces for CIM agents for switches and storage subsystems

This section describes the namespaces for switches and storage subsystem CIM agents that are used in IBM Tivoli Storage Productivity Center.

If you specify an incorrect namespace the following issues might occur:

- The connection test fails when the CIMOM is added.

- The discovery does not discover all information of the system that is managed by the CIMOM.
- The probe fails.
- The function that you want to perform on the system might fail (for example, collecting performance data).

See the following table contains the interop namespaces.

Table 43. Interop Namespaces for CIM Agents for switches and storage subsystems

Switch or subsystem	Namespace
IBM	/root/ibm
Brocade	/interop
Cisco	/root/cimv2 For version 3.2.1 or later: /root/pg_interop
Engenio	/interop
EMC	/root/emc
HDS	For HiCommand 5.6 or later, use: /root/smis/current For a HiCommand version that supports SMI-S 1.2, use following namespace to traverse the model as SMI-S 1.1: /root/smis/smis11 For HiCommand versions earlier than HiCommand 5.6, use: /root/hitachi/dm*xx*, where *xx* is represents the level of HiCommand.
HP	/root
McData	/interop
SUN	/root/sun3510 or /interop Note: This namespace is for a subsystem and not a switch.
XYRATEX	/root/PG_interop

Log files

There are several product logs files to check when you have a problem.

Default log file locations

Check the following default log file locations when you have a problem.

Table 44. Default log file locations for IBM Tivoli Storage Productivity Center components

Component	Log file location

Data Server	For Windows:
	`c:\Program Files\IBM\TPC\data\log`
	For UNIX or Linux:
	`/opt/IBM/TPC/data/log`

Table 44. Default log file locations for IBM Tivoli Storage Productivity Center components (continued)

Component	Log file location

Device server	For Windows:
	`c:\Program Files\IBM\TPC\device\log`
	For UNIX or Linux:
	`/opt/IBM/TPC/device/log`

Common agent	For Windows:
	`c:\Program Files\IBM\TPC\ca\logs`
	For UNIX or Linux:
	`/opt/IBM/TPC/ca/logs`
	For agent on Virtual I/O Server (see note):
	`/home/padmin/agentInstall.log`

Data agent	For Windows:
	`c:\Program Files\IBM\TPC\ca\subagents\TPC\Data\log`
	For UNIX or Linux:
	`/opt/IBM/TPC/ca/subagents/TPC/Data/log`
	For agent on Virtual I/O Server (see note):
	`/home/padmin/agentInstall.log`

Fabric agent	For Windows:
	`c:\Program Files\IBM\TPC\ca\subagents\TPC\Fabric\log`
	For UNIX or Linux:
	`/opt/IBM/TPC/ca/subagents/TPC/Fabric/log`
	For agent on Virtual I/O Server (see note):
	`/home/padmin/agentInstall.log`

GUI	For Windows:
	`c:\Program Files\IBM\TPC\gui\log`
	For UNIX or Linux:
	`/opt/IBM/TPC/gui/log`

Database schema	For Windows:
	`c:\Program Files\IBM\TPC\dbschema\log`
	For UNIX or Linux:
	`/opt/IBM/TPC/dbschema/log`

Note: The `agentInstall.log` file is composed of these log files:

TPC.log
This log is created by InstallShield.

log.txt This log is created by InstallShield.

installStatus.log
Shows the Data agent installation status.

dataAgentInstall.log
The Data agent installation log.

fabricAgentInstallIS.log
The Fabric agent installation log.

install.status
Installation status of the Fabric agent installation.

agentInstall.log
The Common agent installation log.

epInstallStatus.log
The Common agent installation status.

msgAgent.log
The Common agent runtime log.

Agent Manager log files

Agent Manager log files give important information about an installation, uninstallation, and so forth.

Agent Manager log files can be found in the following locations:

Table 45. Agent Manager Log file locations

Agent Manager log files	Location
Installation	`<Agent_Manager_install_dir>\logs`
Uninstallation	`<Agent_Manager_install_dir>\logs`
Run-time	`<Agent_Manager_install_dir>\AppServer\agentmanager\logs\AgentManager`

Appendix I. Performance data

Use the IBM Tivoli Storage Productivity Center to view the performance data of devices in your environment.

General performance guidelines

This topic provides some general performance considerations for configuring the IBM Tivoli Storage Productivity Center environment.

- Do not overpopulate the SLP discovery panel with SLP agent hosts. Remember that Tivoli Storage Productivity Center includes a built-in SLP User Agent (UA) that will receive information about SLP Service Agents and Directory Agents (DA) that reside in the same subnet as the Tivoli Storage Productivity Center server.
- You should have not more than one DA per subnet.
- Misconfiguring the Tivoli Storage Productivity Center CIMOM discovery preferences may impact performance on auto discovery or on device presence checking. It may also result in application time-outs, as attempts are made to resolve and communicate with hosts that are not available.
- It should be considered mandatory to run the CIM Agent software on separate host from the Tivoli Storage Productivity Center server. Attempting to run a full Tivoli Storage Productivity Center implementation on the same host as the CIM agent, will result in dramatically increased wait times for data retrieval. You may also experience resource contention and port conflicts.

Performance metrics

This topic lists the metrics that you can use in performance reports.

Volume-based metrics

The following table contains information on volume-based metrics.

Note: IBM Tivoli Storage Productivity Center does not calculate volume-based metrics if there are multiple ranks in the extent pool where the space-efficient volumes are allocated. The columns for volume-based metrics display the value N/A for the reports in which they appear. However, if there is only a single rank in the extent pool, Tivoli Storage Productivity Center does calculate the volume-based metrics regardless of the space-efficient volumes.

Column	Description
Subsystem ID	ID of the storage subsystem
Time	Date and time that the data was collected.
Interval	Size of the sample interval (in seconds). For Tivoli Storage Enterprise Storage Server, DS6000, and DS8000 models, you can specify a minimum interval length of five minutes and a maximum interval length of one hour. For SAN Volume Controller models, you can specify a minimum interval length of 15 minutes and a maximum interval length of one hour.
I/O Rates	

Column	Description
Global Mirror Write I/O Rate	Average number of write operations per second issued to the Global Mirror secondary site.
Global Mirror Overlapping Write Percentage	Average percentage of write operations issued to the Global Mirror secondary site which were overlapping writes.
Global Mirror Overlapping Write I/O Rate	Average number of overlapping write operations per second encountered by the Global Mirror primary site.
Cache hit percentages	
Readahead Percentage of Cache Hits	Percentage of all read cache hits which occurred on prestaged data.
Dirty Write Percentage of Cache Hits	Percentage of all write cache hits which occurred on already dirty data in the cache.
Response times	
Peak Read Response Time	The peak (worst) response time among all read operations.
Peak Write Response Time	The peak (worst) response time among all write operations.
Global Mirror Write Secondary Lag	The number of additional milliseconds it took to service each secondary write operation for Global Mirror, over and above the time needed to service the primary writes.
Transfer sizes	
Read Transfer Size	Average number of KB per I/O for read operations.
Write Transfer Size	Average number of KB per I/O for write operations.
Overall Transfer Size	Average number of KB per I/O for read and write operations.
Write-cache constraints	
Write-cache Delay Percentage	Percentage of I/O operations that were delayed due to write-cache space constraints or other conditions (the ratio of delayed operations to total I/Os).
Write-cache Delayed I/O Rate	Average number of I/O operations per second that were delayed due to write-cache space constraints or other conditions.
Write-cache Overflow Percentage	Percentage of write operations that were delayed due to lack of write-cache space.
Write-cache Overflow I/O Rate	Average number of tracks per second that were delayed due to lack of write-cache space.
Write-cache Flush-through Percentage	Percentage of write operations that were processed in Flush-through write mode.
Write-cache Flush-through I/O Rate	Average number of tracks per second that were processed in Flush-through write mode.
Write-cache Write-through Percentage	Percentage of write operations that were processed in Write-through write mode.
Write-cache Write-through I/O Rate	Average number of tracks per second that were processed in Write-through write mode.
Record mode reads	
Record Mode Read I/O Rate	The record mode read I/O rate for the controller.

Column	Description
Record Mode Read Cache %	The record mode read I/O cache for the controller.
NVS constraints	
Write-cache Delay Percentage	Percentage of I/O operations that were delayed due to write-cache space constraints or other conditions. Write-cache Delay Percentage is the ratio of delayed I/O operations to total I/O operations.
NVS Delayed I/O Rate	Average number of I/O operations per second that were delayed because of nonvolatile storage (NVS) space constraints.
Cache holding time	
Cache Holding Time	Average cache holding time, in seconds, for I/O data in this subsystem controller (cluster). Shorter time periods indicate adverse performance.
CPU Utilization	Average utilization percentage of the CPUs.

Back-end-based metrics

The following table contains information on back-end-based metrics.

Column	Description
I/O rates	
Backend Read I/O Rate	Average number of I/O operations per second for read operations.
Backend Write I/O Rate	Average number of I/O operations per second for write operations.
Total Backend I/O Rate	Average number of I/O operations per second for read and write operations.
Data rates	
Backend Read Data Rate	Average number of Mbps that were transferred for read operations.
Backend Write Data Rate	Average number of Mbps that were transferred for write operations.
Total Backend Data Rate	Average number of Mbps that were transferred for read and write operations.
Response times	
Backend Read Response Time	Average number of milliseconds that it took to respond to each read operation. For SAN Volume Controller models, this is the external response time of the managed disks (MDisks).
Backend Write Response Time	Average number of milliseconds that it took to respond to each write operation. For SAN Volume Controller models, this is the external response time of the managed disks (MDisks).
Overall Backend Response Time	Average number of milliseconds that it took to respond to each I/O operation (read and write). For SAN Volume Controller models, this is the external response time of the managed disks (MDisks).
Read Queue Time	Average number of milliseconds that each read operation spent on the queue before being issued to the back-end device.

Column	Description
Write Queue Time	Average number of milliseconds that each write operation spent on the queue before being issued to the back-end device.
Overall Queue Time	Average number of milliseconds that read and write operations spent on the queue before being issued to the back-end device.
Transfer sizes	
Back-end Read Transfer Size	Average number of KB per I/O for read operations.
Back-end Write Transfer Size	Average number of KB per I/O for write operations.
Overall Back-end Transfer Size	Average number of KB per I/O for read and write operations.
Disk utilization	
Disk Utilization Percentage	The approximate utilization percentage of a particular rank over a time interval (the average percent of time that the disks associated with the array were busy). **Note:** Tivoli Storage Productivity Center does not calculate a value for this column if there are multiple ranks in the extent pool where the space-efficient volumes are allocated. This column displays value of **N/A** for the reports in which it appears. However, if there is only a single rank in the extent pool, Tivoli Storage Productivity Center does calculate the value for this column regardless of the space-efficient volumes.
Sequential I/O Percentage	Percentage of all I/O operations performed for a particular array over a time interval that were sequential operations.

Front-end- and fabric-based metrics

The following table contains information on front-end- and fabric-based metrics.

Column	Description
I/O or packet rates	
Port Send I/O Rate	Average number of I/O operations per second for send operations.
Port Receive I/O Rate	Average number of I/O operations per second for receive operations.
Total Port I/O Rate	Average number of I/O operations per second for send and receive operations.
Port to Host Send I/O Rate	Average number of exchanges (I/Os) per second sent to host computers.
Port to Host Receive I/O Rate	Average number of exchanges (I/Os) per second received from host computers.
Total Port to Host I/O Rate	Average number of exchanges (I/Os) per second transmitted between host computers.
Port to Disk Send I/O Rate A	Average number of exchanges (I/Os) per second sent to storage subsystems.
Port to Disk Receive I/O Rate	Average number of exchanges (I/Os) per second received from storage subsystems.
Total Port to Disk I/O Rate	Average number of exchanges (I/Os) per second transmitted between storage subsystems.

Column	Description
Port to Local Node Send I/O Rate	Average number of exchanges (I/Os) per second sent to other nodes in the local SAN Volume Controller cluster.
Port to Local Node Receive I/O Rate	Average number of exchanges (I/Os) per second received from other nodes in the local SAN Volume Controller cluster.
Total Port to Local Node I/O Rate	Average number of exchanges (I/Os) per second transmitted between other nodes in the local SAN Volume Controller cluster.
Port to Remote Node Send I/O Rate	Average number of exchanges (I/Os) per second sent to nodes in the remote SAN Volume Controller cluster.
Port to Remote Node Receive I/O Rate	Average number of exchanges (I/Os) per second received from nodes in the remote SAN Volume Controller cluster.
Total Port to Remote Node I/O Rate	Average number of exchanges (I/Os) per second transmitted between nodes in the remote SAN Volume Controller cluster.
Data rates	
Port Send Data Rate	Average number of MBps that were transferred for send (write) operations.
Port Receive Data Rate	Average number of MBps that were transferred for receive (read) operations.
Total Port Data Rate	Average number of MBps that were transferred for send and receive operations.
Port to Host Send Data Rate	Average number of megabytes (2^20 bytes) per second sent to host computers.
Port to Host Receive Data Rate	Average number of megabytes (2^20 bytes) per second received from host computers.
Total Port to Host Data Rate	Average number of megabytes (2^20 bytes) per second transmitted between host computers.
Port to Disk Send Data Rate	Average number of megabytes (2^20 bytes) per second sent to storage subsystems.
Port to Disk Receive Data Rate	Average number of megabytes (2^20 bytes) per second received from storage subsystems.
Total Port to Disk Data Rate	Average number of megabytes (2^20 bytes) per second transmitted between storage subsystems.
Port to Local Node Send Data Rate	Average number of megabytes (2^20 bytes) per second sent to other nodes in the local SAN Volume Controller cluster.
Port to Local Node Receive Data Rate	Average number of megabytes (2^20 bytes) per second received from other nodes in the local SAN Volume Controller cluster.
Total Port to Local Node Data Rate	Average number of megabytes (2^20 bytes) per second transmitted between other nodes in the local SAN Volume Controller cluster.
Port to Remote Node Send Data Rate	Average number of megabytes (2^20 bytes) per second sent to nodes in the remote SAN Volume Controller cluster.
Port to Remote Node Receive Data Rate	Average number of megabytes (2^20 bytes) per second received from nodes in the remote SAN Volume Controller cluster.
Total Port to Remote Node Data Rate	Average number of megabytes (2^20 bytes) per second transmitted between nodes in the remote SAN Volume Controller cluster.
Response times	
Port Send Response Time	Average number of milliseconds that it took to service each send (write) operation.

Column	Description
Port Receive Response Time	Average number of milliseconds that it took to service each receive (read) operation.
Overall Port Response Time	Average number of milliseconds that it took to service each operation (send and receive).
Port to Local Node Send Response Time	Average number of milliseconds it took to service each send operation. For SAN Volume Controller, this is the external response time of the transfers.
Port to Local Node Receive Response Time	Average number of milliseconds it took to service all receive operations. For SAN Volume Controller, this is the external response time of the transfers.
Total Port to Local Node Response Time	Average number of milliseconds it took to service all send and receive operations. For SAN Volume Controller, this is the external response time of the transfers.
Port to Local Node Send Queued Time	Average number of milliseconds that each send operation spent on the queue before being issued.
Port to Local Node Receive Queued Time	Average number of milliseconds that each receive operation spent on the queue before being issued.
Total Port to Local Node Queued Time	Average number of milliseconds that each operation spent on the queue before being issued.
Port to Remote Node Send Response Time	Average number of milliseconds it took to service each send operation. For SAN Volume Controller, this is the external response time of the transfers.
Port to Remote Node Receive Response Time	Average number of milliseconds it took to service all receive operations. For SAN Volume Controller, this is the external response time of the transfers.
Total Port to Remote Node Response Time	Average number of milliseconds it took to service all send and receive operations. For SAN Volume Controller, this is the external response time of the transfers.
Port to Remote Node Send Queued Time	Average number of milliseconds that each send operation spent on the queue before being issued.
Port to Remote Node Receive Queued Time	Average number of milliseconds that each receive operation spent on the queue before being issued.
Total Port to Remote Node Queued Time	Average number of milliseconds that each operation spent on the queue before being issued.
Transfer sizes	
Port Send Transfer Size	Average number of KB sent per I/O by a particular port over a time interval.
Port Receive Transfer Size	Average number of KB received per I/O by a particular port over a time interval.
Overall Port Transfer Size	Average number of KB transferred per I/O by a particular port over a time interval.
Port Send Packet Size	Average number of KB sent per packet by a particular port over a time interval.
Port Receive Packet Size	Average number of KB received per packet by a particular port over a time interval.
Overall Port Packet Size	Average number of KB transferred per packet by a particular port over a time interval.
Error rates	

Column	Description
Error Frame Rate	The number of frames per second that were received in error by a particular port over a time interval.
Dumped Frame Rate	The number of frames per second that were lost due to a lack of available host buffers.
Link Failure Rate	The number of link errors per second that were experienced by a particular port over a time interval.

Performance thresholds

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

You can monitor the performance of your enterprise by creating alerts on performance thresholds for switches and storage subsystems. By creating alerts that are triggered by performance thresholds, you can be informed about performance issues in your enterprise. Threshold events tell you when a component has fallen outside of the user-defined values. For example, when a threshold value has reached critical stress.

Threshold boundaries

You can establish your boundaries for the normal expected subsystem performance when defining storage subsystem alerts for performance threshold events. When the collected performance data samples fall outside out of the range you have set, you are notified of this threshold violation so you are aware of the potential problem. The upper boundaries are **Critical Stress** and **Warning Stress**. The lower boundaries are **Warning Idle** and **Critical Idle**. Usually you will want the stress boundaries to be high numbers and the idle to be low numbers. The exception to this rule is **Cache Holding Time Threshold**, where you want the stress numbers to be low and the idle numbers to be high.

If you do not want to be notified of threshold violations for any boundaries, you can leave the boundary field blank and the performance data will not be checked against any value. For example, if the **Critical Idle** and **Warning Idle** fields are left blank, no alerts will be sent for any idle conditions.

The **Suppress alerts when sequential I/O exceeds** check button is active only for the triggering condition **Disk Utilization Percentage Threshold**. It is a filter condition. The default is 80%.

Setting the thresholds

Only a few thresholds have defaults and on the other thresholds you will have to determine the best values for stress, idle, critical, and warning values so you can derive the maximum benefit without generating too many false alerts. Since suitable stress thresholds are highly dependant on the type of workload you are running, your exact hardware configuration, the number of physical disks, exact model numbers, and so forth, there are no easy or standard default rules.

One of the best approaches is to monitor your performance for a number of weeks and, using this historical data, determine reasonable values for each threshold setting. Once that is done you can fine tune these settings to minimize number of false alerts.

Related tasks

"Archive and delete files" on page 101
This scenario shows you how to reclaim storage space by archiving and deleting files that meet criteria you define.

Related reference

"Controller thresholds" on page 792
The following table lists and describes the Controller thresholds.

"Port thresholds" on page 792
The following table lists and describes the Port thresholds.

"Array thresholds" on page 793
The following table lists and describes the Array thresholds.

Array thresholds

The following table lists and describes the Array thresholds.

Table 46. Array thresholds

Threshold (Metric)	Device/Component Type	Description
Array Thresholds		
Disk Utilization Percentage	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 array	Sets thresholds on the approximate utilization percentage of the arrays in a particular subsystem; for example, the average percentage of time that the disks associated with the array were busy. The Disk Utilization metric for each array is checked against the threshold boundaries for each collection interval. This threshold is enabled by default for Tivoli Storage Enterprise Storage Server systems and disabled by default for others. The default threshold boundaries are 80%, 50%, -1, -1. For DS6000 and DS8000 subsystems, this threshold applies only to those ranks which are the only ranks in their associated extent pool.
Total Back-end I/O Rate	SAN Volume Controller MDisk group	Sets thresholds on the average number of I/O operations per second for MDisk read and write operations for the MDisk groups. The Total I/O Rate metric for each MDisk group is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Total Back-end Data Rate	SAN Volume Controller MDisk group	Sets thresholds on the average number of MB per second that were transferred for MDisk read and write operations for the MDisk groups. The Total Data Rate metric for each MDisk group is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.

Table 46. Array thresholds (continued)

Threshold (Metric)	Device/Component Type	Description
Overall Back-end Response Time	SAN Volume Controller MDisk	Sets thresholds on the average number of milliseconds that it took to service each MDisk I/O operation, measured at the MDisk level. The Total Response Time (external) metric for each MDisk is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.

Controller thresholds

The following table lists and describes the Controller thresholds.

Table 47. Controller thresholds

Threshold (Metric)	Device/Component Type	Description
Controller Thresholds		
Total I/O Rate (overall)	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 controller SAN Volume Controller I/O group	Sets threshold on the average number of I/O operations per second for read and write operations, for the subsystem controllers (clusters) or I/O groups. The Total I/O Rate metric for each controller or I/O group is checked against the threshold boundaries for each collection interval. These thresholds are disabled by default.
Total Data Rate	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 controller SAN Volume Controller I/O group	Sets threshold on the average number of MB per second for read and write operations for the subsystem controllers (clusters) or I/O groups. The Total Data Rate metric for each controller or I/O group is checked against the threshold boundaries for each collection interval. These thresholds are disabled by default.
NVS Full Percentage	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 controller	Sets thresholds on the percentage of time that NVS space constraints caused I/O operations to be delayed, for the subsystem controllers (clusters). The NVS Full Percentage metric for each controller is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries of 10, 3, -1, -1.
Cache Holding Time	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 controller	Sets thresholds on the average cache holding time, in seconds, for I/O data in the subsystem controllers (clusters). Shorter time periods indicate adverse performance. The Cache Holding Time metric for each controller is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries of 30, 60, -1, -1.

Port thresholds

The following table lists and describes the Port thresholds.

Table 48. Port thresholds

Threshold (Metric)	Device/Component Type	Description
Port Thresholds		
Total Port IO Rate	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 port, switch port	Sets thresholds on the average number of I/O operations or packets per second for send and receive operations, for the ports. The Total I/O Rate metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Total Port Data Rate	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 port, switch port	Sets thresholds on the average number of MB per second for send and receive operations, for the ports. The Total Data Rate metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Overall Port Response Time	Tivoli Storage Enterprise Storage Server, DS6000, DS8000 port	Sets thresholds on the average number of milliseconds that it took to service each I/O operation (send and receive) for ports. The Total Response Time metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.
Error Frame Rate	Switch port	Sets thresholds on the average number of frames per second received in error for the switch ports. The Error Frame Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default.
Link Failure Rate	Switch port	Sets thresholds on the average number of link errors per second experienced by the switch ports. The Link Failure Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default.
Total Port Data Rate	Switch port	Sets thresholds on the critical and warning data rates for stress and idle in MB per second. The Total Port Data Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default.
Total Port Packet Rate	Switch port	Sets thresholds on the critical and warning data rates for stress and idle in packets per second. The Total Port Packet Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default.

Appendix J. Protocols and standards

This section provides an overview of the protocols and standards that are used within IBM Tivoli Storage Productivity Center.

Web Based Enterprise Management

Web Based Enterprise Management (WBEM) is an initiative of the Distributed Management Task Force (DTMF) with the objective to enable the management of complex IT environments. It defines a set of management and internet standard technologies in order to unify the management of complex IT environments.

The WBEM initiative is composed of three main conceptual elements:

Common Interface Model (CIM)
CIM is a formal object-oriented modeling language that is used to describe the management aspects of systems.

xmlCIM
This is the grammar to describe CIM declarations and messages used by the CIM protocol.

Hypertext Transfer Protocol (HTTP)
HTTP is used as a way to enable communication between a management application and a device that both use CIM.

The WBEM architecture defines the following elements:

CIM Client
The CIM Client is a management application like IBM Tivoli Storage Productivity Center that uses CIM to manage devices. A CIM Client can reside anywhere in the network, because it uses HTTP to talk to CIM Object Managers and Agents.

CIM Managed Object
A Managed Object is a hardware or software component that can be managed by a management application by using CIM.

CIM Agent
A CIM Object Manager that includes the provider service for a limited set of resources. An agent may be embedded or hosted and can be an aggregator for multiple devices.

CIM Provider
A CIM Provider is the element that translates CIM calls to the device-specific commands. A provider is always closely linked to a CIM.

CIM Object Manager (CIMOM)
The central component of the CIM Server responsible for the communication between the CIM server components.

CIM Server
A server that receives and processes CIM Operation Message Requests and issues CIM Operation Message Responses.

Storage Management Initiative Specification

The Storage Networking Industry Association (SNIA) defines a standard that is used within IBM Tivoli Storage Productivity Center to create and develop a universal open interface for managing storage devices including storage networks.

Information on SMI-S can be found at this Web site:

`http://www.snia.org`

SNIA has fully adopted and enhanced the Common Information Model (CIM) standard for storage management in its Storage Management Initiative - Specification (SMI-S). SMI-S was launched to create and develop a universal open interface for managing storage devices including storage networks. SMI-S provides:

- A comprehensive specification for the management of heterogeneous storage and storage area networks (SANs).
- The information available to a WBEM client from an SMI-S compliant CIM server (provider).
- Profiles organized by:
  - Storage
  - Fabric
  - Host
  - Common profiles and subprofiles
- An object-oriented CIM and XML-based interface for managing SAN devices, services, and fabrics.
- An initial discovery, which is SLP based.

The idea behind SMI-S is to standardize the management interfaces so that management applications can utilize these and provide cross-device management. This means that a newly introduced device can be immediately managed as it will conform to the standards.

The models and protocols in the SMI-S implementation are platform-independent, enabling application development for any platform, and enabling them to run on different platforms. The SNIA will also provide interoperability tests which will help vendors test their applications and devices if they conform to the standard.

Service Location Protocol

The Service Location Protocol (SLP) is an Internet Engineering Task Force (IETF) standard. SLP provides a scalable framework for the discovery and selection of network services.

The Internet Engineering Task Force (IETF) is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. The IETF includes formal standards for SNMP and MIBs. More information on IETF can be found at this Web site:

`http://www.ietf.org`

SLP enables the discovery and selection of generic services, which could range in function from hardware services such as those for printers or fax machines, to software services such as those for file servers, e-mail servers, Web servers, databases, or any other possible services that are accessible through an IP network.

Traditionally, to use a particular service, a user or client application needed to supply the host name or network IP address of that service. With SLP, however, the user or client no longer needs to know individual host names or IP addresses. Instead, the user or client can search the network for the desired service type and an optional set of qualifying attributes.

For example, a user could specify to search for all available printers that support Postscript. Based on the given service type (printers), and the given attributes (Postscript). SLP searches the user's network for any matching services, and returns the discovered list to the user.

Simple Network Management Protocol

The Simple Network Management Protocol (SNMP) is an Internet Engineering Task Force (IETF) protocol for monitoring and managing systems and devices in a network. Functions supported by the SNMP protocol are the request and retrieval of data, the setting or writing of data, and traps that signal the occurrence of events.

SNMP is a method that enables a management application to query information from a managed device. The managed device has software running that sends and receives the SNMP information. This software module is usually called the SNMP agent.

An SNMP manager can read information from an SNMP agent in order to monitor a device. Therefore, the device needs to be polled on an interval bases. The SNMP manager can also change the configuration of a device, by setting certain values to corresponding variables. A device can also be set up to send a notification to the SNMP manager (this is called a trap) to asynchronously inform this SNMP manager of a status change.

Depending on the existing environment and organization it is very likely that your environment already has an SNMP management application in place. IBM Tivoli Storage Productivity Center can be set up to send traps.

SNMP uses a hierarchical structured Management Information Base (MIB) to define the meaning and the type of a particular value. A MIB defines managed objects that describe the behavior of the SNMP entity, which can be anything from an IP router to a storage subsystem. The information is organized in a tree structure.

For users planning to make use of the Tivoli Storage Productivity Center SNMP trap alert notification capabilities, an SNMP MIB is included in the server installation.

The MIB is provided for use by your SNMP management console software. Most SNMP management products provide a program called a MIB compiler that can be used to import MIBs. This will allow you to view Tivoli Storage Productivity Center generated SNMP traps from within your management console software. Refer to your management console software documentation for instructions on how to compile or import a third party MIB.

Tivoli Storage Productivity Center uses port 162 to listen for SNMP traps. This is the default port. For switches, you must configure the switch to send SNMP traps to the Device server IP address. If you need to change the default port number, use the **setdscfg** command. The attribute to set is **SNMPTrapPort**.

System administrators must set up their SNMP trap ringer with the provided MIB files in order to receive SNMP traps from Tivoli Storage Productivity Center. These files are located in the following directories on the product installation CD:

For Data server:

```
data\snmp\tivoliSRM.mib
```

For Device server:

```
device\snmp\fabric.mib
```

IBM Tivoli Storage Productivity Center Universal Agent

IBM Tivoli Storage Productivity Center Universal Agent collects information about the Tivoli Storage Productivity Center Health, Data Server, and Data Server services information, Device information, alert information, and job information.

Tivoli Storage Productivity Center Health, Data Server and Data Server services Information

This will include:

- Data server status
- Device server status
- Services status
- CIMOM connection status
- CIMOM information such as last Discovery, Managed devices
- Agent connection status
- Agent information such as last Discovery and Probe, Managed devices
- Equivalent information about other Tivoli Storage Productivity Center servers, VM Ware

Alert Information

- Amount of all alerts
- Alerts per component such as Computer, Data, Disk, Fabric etc.

Job Information

- Amount of Jobs, such as Discovery, Probe, Scans, PM
- Job Status and Details, For example: Start Time, Finish Time, Status, Log File Name
- Scheduled Jobs and Details, For example: Intervals, Creator, Name

For more information, in the Tivoli Storage Productivity Center installation media, you can view the README.txt file located in the TPCUA.zip file, in the Tools folder of Disk 1.

Fibre Channel Methodologies of Interconnects

IBM Tivoli Storage Productivity Center supports the ANSI T11 Fibre Channel FC-MI (Fibre Channel Methodologies of Interconnects) for the automated discovery of FC SAN assets and topology.

ANSI T11 Fibre Channel FC-MI includes the following for the automated discovery of FC SAN assets and topology:

- Hosts (HBAs)
- FC interconnects
- FC storage devices

The T11 FC-MI also includes the following:
- FC-GS-3/4 (discovery, zoning, and so forth)
- RNID (advanced device recognition)
- Platform registration (device recognition and launch)
- Common HBA API (fabric and storage views)
- Name server (connectivity)
- Management server (SAN connectivity and topology)
- RSCN (advanced event detection)
- SCSI queries (storage views, volume information, and so forth)
- SNMP Fabric Element (FE) MIB
- SNMP FC Management MIB (discovery, performance statistics, and so forth)
- SNMP alerts

Appendix K. Windows services used by IBM Tivoli Storage Productivity Center

This topic lists Windows services used by IBM Tivoli Storage Productivity Center.

To start, stop, or restart a component or related program, use the Windows Services panel. The following table provides a list of the services.

Table 49. List of Windows services used by Tivoli Storage Productivity Center

Program	Service name	Comment
IBM DB2	DB2 - DB2COPY1 - DB2 - 0 DB2 Governer (DB2COPY1) DB2 License Server (DB2COPY1) DB2 Management Service (DB2COPY1) DB2 Remote Command Server (DB2COPY1) DB2DAS - DB2DAS00 DB2TS - DB2COPY1 - DB2-0	The service account owner is **db2admin**. The account needs to be part of Administrators and DB2ADMNS.
Agent Manager	IBM WebSphere Application Server V6.1 - Tivoli Agent Manager	The service account owner is the local system. During Agent Manager installation, you have the option to set this service to start automatically.
Common agent	IBM Tivoli Common agent - 'C:\Program Files\IBM\TPC\ca'	The service account owner is **itcauser**. This ID is created by the installation program.
Tivoli Storage Productivity Center Data server	IBM Tivoli Storage Productivity Center - Data Server	Note that the Data server does not have a WebSphere application service name.
Tivoli Storage Productivity Center Device server	IBM WebSphere Application Server V6.1 - DeviceServer	
IBM Tivoli Storage Productivity Center for Replication	IBM WebSphere Application Server V6.1 - CSM	
IBM Tivoli Integrated Portal	Tivoli Integrated Portal - TIPProfile_Port_16310	
IBM Deployment Engine	IBM ADE Service	This service is installed withTivoli Integrated Portal,

Appendix L. Accessibility features for IBM Tivoli Storage Productivity Center

IBM strives to provide products with usable access for everyone, regardless of age or ability.

Accessibility features

The following list includes the major accessibility features in IBM Tivoli Storage Productivity Center:

- IBM 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. It also allows you to use screen readers and other assistive technologies.
  - All documentation is available in PDF format.
  - 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 that are supported through the native Windows manager.

Interface information

Use the options available in the **Preferences** > **Look and Feel** menu to select how to display the IBM Tivoli Storage Productivity Center user interface. To do this, complete the following steps:

1. Start the IBM Tivoli Storage Productivity Center user interface.
2. Select one of the following options from the **Preferences** > **Look and Feel** menu to change the visual appearance of the user interface to best suit your visual needs:
   - Windows Classic
   - Windows
   - CDE/Motif
   - Metal

Related accessibility information

You can view the publications for IBM Tivoli Storage Productivity Center in Adobe Portable Document Format (PDF) using the Adobe Acrobat Reader. The PDFs are provided on a CD that is packaged with the product, or you can access them from the Printable documentation topic in the information center at http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp.

IBM and accessibility

See the IBM Human Ability and Accessibility Center website for more information about the commitment that IBM has to accessibility.

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:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, 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.

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
Software Interoperability Coordinator, Department YBWA
3605 Highway 52 N
11400 Burnet Road
Rochester, MN 55901
U.S.A

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml.

Adobe and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

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/.

agent An entity that represents one or more managed objects by sending notifications regarding the objects and by handling requests from servers for management operations to modify or query the objects.

Basic HyperSwap™
In System z, a replication feature that performs the following actions:

- Monitoring for events that indicate a storage device has failed
- Determining whether the failing storage device is part of a Peer-to-Peer Remote Copy (PPRC) pair
- Determining from policy, the action to be taken
- Ensuring that data consistency is not violated
- Swapping the I/O between the primary logical devices in the consistency group with the secondary logical devices in the consistency group.
- Allowing only CKD volumes to be added to the HyperSwap session.

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 CIM-compliant devices.

CIM object manager (CIMOM)
The common conceptual framework for data management that receives, validates, and authenticates the CIM requests from the client application. It then directs the requests to the appropriate component or service provider.

CIMOM
See CIM object manager

CKD Count key data

cluster
1. In SAN Volume Controller, 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.

comma-separated value file
A text file, created in a spreadsheet program such as Microsoft Excel. A CSV file includes each of the copy sets you want to add to the session separated by a comma.

community name
The part of an SNMP message that represents a password-like name and that is used to authenticate the SNMP message.

consistency group
The set of target volumes in a session that have been updated to preserve write ordering and are therefore, recoverable.

copy set
A set of volumes that contain copies of the same data. All the volumes in a copy set are the same format (count key data [CKD] or fixed block) and size.

CSV See comma-separated value file

data collection
See *discovery*.

data exposure
The time between the point at which the data is written to primary storage, and when it is replicated to secondary storage. Data exposure includes factors such as:
- Requested consistency-group interval time
- Type of storage systems
- Physical distance between the storage systems
- Available bandwidth of the data link
- I/O load on the storage systems

discovery
The process of finding resources within an enterprise, including finding the new location of monitored resources that were moved. Discovery includes the detection of changes in network topology, such as new and deleted nodes or new and deleted interfaces. See also *discovery interval*.

discovery interval
The frequency at which topology and attribute information is gathered. The discovery interval is set by a schedule to occur either periodically or at specific times. Discovery can also occur at other times, such as when triggered by an event from a SAN switch.

discovery job
A job that enables you to find new Windows machines that have been introduced into your environment, identify the servers and volumes within NetWare trees (NDS trees), discover the file systems within NAS filers, and discover the CIMOMs in your environment and the storage subsystems managed by those CIM/OMs.

enterprise repository
A component of the Data server that records and stores all information about the monitored computers' storage assets and their usage over time. The repository is organized into relational database tables and is accessed by Data server using Java™ Database Connectivity (JDBC).

event Any significant change in the state of a system resource, network resource, or network application. An event can be generated for a problem, for the resolution of a problem, or for the successful completion of a task. Examples of events are: the normal starting and stopping of a process, the abnormal termination of a process, or the malfunctioning of a server.

fabric A complex network using hubs, switches, and gateways. Fibre channel uses a fabric to connect devices.

failover and failback
The implementation of a complex local or remote disaster-recovery solution with the capability of a two-way site switch.

fibre channel
A technology for transmitting data between computer devices. It is especially suited for attaching computer servers to shared storage devices and for interconnecting storage controllers and drives.

FlashCopy
An optional feature of the DS8000 series that can make an instant copy of data; that is, a point-in-time copy of a volume.

global copy
An optional capability of the DS8000 remote mirror and copy feature that maintains a fuzzy copy of a logical volume on the same DS8000 storage unit or on another DS8000 storage unit. In other words, all modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume at a later point in time. The original order of update is not strictly maintained. See also *remote mirror* and *copy* and *metro mirror*.

global mirror
An optional capability of the remote mirror and copy feature that provides a 2-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.

heat map
A color-coded data chart where colors are used to differentiate values in a data set.

host A computer that is connected to a network (such as the Internet or a SAN) and provides a point of access to that network. Also, depending on the environment, the host can provide centralized control of the network. The host can be a client, a server, both a client and a server, a manager, or a managed host.

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, and receives read, write, and update application I/Os, depending on the site that the application is writing to.

in-band discovery
The process of discovering information about the SAN, including topology and attribute data, through the fibre-channel data paths. Contrast with *out-of-band discovery*.

intermediate volume
The target of the remote copy relationship, and the source of a FlashCopy relationship in which the target of the FlashCopy is the H2 volume.

job scheduler
A component of the Data server that deploys all monitoring activities. The job scheduler controls when monitoring jobs are run by agents.

journal volume
A volume that functions like a journal and holds the required data to reconstruct consistent data at the Global Mirror remote site. When a session must be recovered at the remote site, the journal volume is used to restore data to the last consistency point.

logical unit number (LUN)
An identifier used on a SCSI bus to distinguish among devices (logical units) with the same SCSI ID. For a SCSI bus, a LUN represents a storage volume.

LUN See *logical unit number*.

managed disk (MDisk)
A SCSI logical unit that a Redundant Array of Independent Disks (RAID) controller provides and a cluster manages. The MDisk is not visible to host systems on the SAN.

managed host
A host that is managed by Tivoli Storage Productivity Center and one or more active in-band fabric agents. Install in-band fabric agents on host systems with host bus adapters (HBAs) that are connected to the SAN fabrics that you want to manage.

Management Servers
Increased availability of the replication management software with the implementation of a high-availability configuration such that one management workstation runs as standby, ready to take over in case of a failure of the active workstation.

Note: The takeover is not automatic and requires you to issue a takeover command.

metro mirror
A function of a storage server that maintains a consistent copy of a logical volume on the same storage server or on another storage server. All modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume. See also *Remote Mirror* and *Copy* and *Global Copy*.

Metro Global Mirror
The three-site remote mirroring solution.

out-of-band discovery
The process of discovering SAN information, including topology and device data, without using the fibre-channel data paths. A common mechanism for out-of-band discovery is the use of SNMP MIB queries, which are invoked over a TCP/IP network. Contrast with *in-band discovery*.

ping job
A job that tracks the availability of assets and that is performed by an agent. Several ping jobs can be used to monitor the availability of any computer or subset of computers in the network.

pool A named set of storage volumes that is the destination for storing client data.

primordial pool
Unallocated storage capacity on a storage device. Storage capacity can be allocated from primordial pools to create storage pools.

probe job
A job that itemizes and creates an inventory of assets, such as computers, controllers, disk drives, file systems, and logical units, and that is performed by an agent. Several probe jobs can be used on any computer or subset of computer

RAID See *Redundant Array of Independent Disks*.

Recovery point objective (RPO)
The maximum amount of data that you can tolerate losing in the case of a disaster.

remote console
A console that is installed on a machine other than the one on which the server is installed. A remote console lets you access Tivoli Storage Productivity Center from any location.

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.

role A function that a volume assumes is the copy set. The role is is composed of the intended use and, for Global Mirror and Metro Mirror, the volume's site location. Every volume in a copy set is assigned a role. A role can assume the functions of a host volume, journal volume, or target volume. For example, a host volume at the primary site might have the role of `Host1`, while a journal volume at the secondary site has the role of `Journal2`.

role pair
The association of two roles in a session that take part in a copy relationship. For example, in a metro mirror session, the role pair could be the association between the volume roles of `Host1` and `Host2`. In another example, a `Host1` volume could be a host volume on the primary site, and a `Host2` volume could be a host volume on the secondary site.

SAN See *storage area network*.

scan job
A job that monitors the usage and consumption of your storage and the constraints and that is performed by an agent. Several scan jobs can be used to monitor the file systems on any computer or subset of computers.

SCSI See *Small Computer Systems Interface*.

session
A collection of multiple copy sets that comprise a consistency group.

site switching
See also *failover and failback*.

SMI-S See *Storage Management Initiative - Specification*.

SMI-S agent
See *CIM Object Manager (CIMOM)*. See also *Storage Management Initiative - Specification (SMI-S)*.

SNIA See *Storage Networking Industry Association*.

source The site where production applications run while in normal operation. The meaning is extended to the disk subsystem that holds the data as well as to its components: volumes and LSS.

storage area network
A dedicated storage network tailored to a specific environment, combining servers, storage products, networking products, software, and services.

storage group
A collection of storage units that jointly contain all the data for a specified set of storage units, such as volumes. The storage units in a group must be from storage devices of the same type.

Storage Management Initiative - Specification (SMI-S)
The standard that defines the protocol used for communication with SMI-S agents.

Storage Networking Industry Association (SNIA)
An alliance of computer vendors and

universities that focus on developing and promoting industry standards for storage networks.

storage pool
An aggregation of storage resources on a SAN that have been set aside for a particular purpose.

System z Global Mirror
See also *Global Mirror*.

target The site to where the data is replicated, the copy of the application data. The meaning is extended to the disk subsystem that holds the data as well as to its components: volumes and logical subsystem (LSS).

target volume
A volume that receives data from a host volume or another intermediate volume. It is used only in FlashCopy sessions.

topology
The physical and logical arrangement of devices in a SAN. Topology can be displayed graphically, showing devices and their interconnections.

VDisk See *virtual disk*.

virtual disk (VDisk)
A device that host systems attached to the storage area network (SAN) recognize as a Small Computer System Interface (SCSI) disk.

virtualization
A concept in which a pool of storage is created that contains several disk subsystems. The subsystems can be from various vendors. The pool can be split into virtual disks that are visible to the host systems that use them.

virtual storage area network (VSAN)
A Cisco technology that allows independent logical fabrics to be defined from a set of one or more physical switches. A given switch port is assigned to only one VSAN. Each VSAN is completely isolated from the other VSANs and functions as a separate and independent fabric with its own set of fabric services (for example, Name Services, zoning, routing, and so on).

volume
The basic entity of data storage as defined by the SCSI protocol. A volume is a logical address space, having its data content stored on the systems disk drives.

VSAN See *virtual storage area network*.

zone A segment of a SAN fabric composed of selected storage devices nodes and server nodes. Only the members of a zone have access to one another.

zone set
A group of zones that function together on the fabric. Each zone set can accommodate up to 256 zones. All devices in a zone see only devices assigned to that zone, but any device in that zone can be a member of other zones in the zone set.

Index

D

E

F

G

H

I

J

L

M

N

O

P

Q

R

S

T

U

V

W

X

Z

Program Number: 5608-WC0, 5608-WC4, 5608-WB1

Printed in USA

SC27-2338-00