

IBM Tivoli Storage Productivity Center
Version 5.1

Installation and Configuration Guide



IBM Tivoli Storage Productivity Center
Version 5.1

Installation and Configuration Guide



Note:

Before using this information and the product it supports, read the information in "Notices" on page 399.

This edition applies to version 5, release 1, modification 0 of IBM Tivoli Storage Productivity Center (product numbers 5725-F93, 5725-G33, and 5725-F98) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2001, 2012.**

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

Contents

Figures	vii
--------------------------	------------

Tables	ix
-------------------------	-----------

Preface	xi
--------------------------	-----------

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

New for Tivoli Storage Productivity Center Version 5.1	xvii
---	-------------

Chapter 1. Planning for Tivoli Storage Productivity Center	1
---	----------

Planning for installation	1
Hardware requirements.	1
Software requirements	4
Product licenses for Tivoli Storage Productivity Center	16
Running Tivoli Storage Productivity Center in a 64-bit environment	23
Planning for Tivoli Storage Productivity Center for Replication	23
Planning for Open HyperSwap replication	23
Planning for Configuration	24
Configuration guidelines	24
Planning for Tivoli Storage Productivity Center authentication and authorization	26
User names and passwords	28
General requirements for user names and passwords.	28
Worksheet for user names and passwords	31
TCP/IP ports used by Tivoli Storage Productivity Center	33
Tivoli Integrated Portal port numbers.	35
Storage systems, devices, file systems, databases, and switches	37
Planning for multipath subsystem device drivers	42
Planning for Storage Resource agents.	43
Storage Resource agent protocol support	44
Planning for Internet Protocol Version 6	46
Planning for the native interface	48
Planning to use LDAP.	50
Software requirements for CIM agents	51
Planning for the single sign-on feature	52
Planning for enterprise-rollup reports.	52
Planning for the SAN configuration tools	54

SAN Planner Overview	56
Storage Optimizer overview	62
Planning for the Data Path Explorer	63
Planning for the Data Manager.	65
Planning for the Storage Resource agents on Oracle Solaris.	66
Planning for the Disk Manager	66
Planning for the IBM CIM agents	68
Planning for the TagmaStore CIM agents	69
Planning for Disk Manager configuration	70
Planning for TotalStorage Enterprise Storage Server FlashCopy reports.	73
Storage capacity of volumes.	73
Planning for the DS8000	73
Planning for DS8000 performance	74
Planning for the SAN Volume Controller	74
Planning for Storwize V7000.	77
Planning for Storwize V7000 Unified	79
Planning for the XIV system.	83
Planning for the Fabric Manager	85
Planning for Brocade management agents	85
Planning for QLogic or Cisco fabric management agents	89
Collecting data with the Fabric Manager.	91
Storage Resource agent	95
Out-of-band Fabric agent.	96
Configuration guidelines for agent placement	96
Planning for Fabric Manager configuration	97
Planning for Fabric Manager switches and directors	99
Planning for host bus adapters (HBAs).	100
Planning for private switch networks	101
Planning for performance management.	101
Performance metrics	102
Performance management configuration	104
Microsoft Cluster Server.	106
Microsoft Cluster Server environment	107
Microsoft Cluster Server support	107
Planning for PowerHA	109
PowerHA environment	111
NAS support	112
Network Attached Storage system requirements	112
Planning for NetApp device support	114
Planning for System Storage N Series Gateway servers	117
Planning for VMware.	117
VMware configuration	120
VMware capacity reports	120
Planning for the Virtual I/O Server	122
Planning for SQL access to Tivoli Storage Productivity Center's views.	124
Setting up a view-only user on Windows	125
Monitoring your DB2 environment	126
Planning for the Discovery Library Adapter	126
Tivoli Storage Productivity Center Monitoring Agent	127

Tivoli Agent Management Services	128
Planning to install Tivoli Storage Productivity Center	130
Tivoli Storage Productivity Center components	133
Tivoli Storage Productivity Center in an IPv4 or IPv6 environment	134

Chapter 2. Installing Tivoli Storage Productivity Center 135

Installing Tivoli Storage Productivity Center	135
Tivoli Storage Productivity Center installation	135
Starting the installation program	138
Installing Tivoli Storage Productivity Center in a single-server environment	141
Installing Tivoli Storage Productivity Center in a multiple-server environment	150
Installing Tivoli Storage Productivity Center client components on a separate computer	161
Installing Tivoli Common Reporting on a remote server	161
Verifying the installation	163
Reviewing the log files to resolve installation issues	165
Starting the stand-alone GUI by using a web browser	166
Adding an installation license	167
Installing the agents	168
Installing DB2	180
Preparing to install DB2	180
Installing DB2 on a Windows operating system	182
Installing DB2 on UNIX or Linux - GUI install	184
Installing DB2 on AIX using the command line	186
Installing DB2 on AIX 6.1 POWER7 Systems	189
Verifying that DB2 is installed correctly.	191
Reinstalling the software if a failure occurs	192

Chapter 3. Configuring Tivoli Storage Productivity Center 195

Starting Tivoli Storage Productivity Center	195
Starting the Tivoli Storage Productivity Center stand-alone GUI	195
Starting the Tivoli Storage Productivity Center web-based GUI.	198
Starting Tivoli Storage Productivity Center for Replication	199
Starting the Tivoli Storage Productivity Center family of products through Tivoli Integrated Portal.	200
Overview of required user IDs for initial log on to the Tivoli Storage Productivity Center family of products	201
User authentication considerations for multiple server installations	204
Changing the user authentication configuration	204
Adding an LDAP repository to the federated repositories	205
Removing an LDAP repository from the federated repositories.	213
Tivoli Integrated Portal and Tivoli Common Reporting Roles	219
Assigning Tivoli Integrated Portal roles to users	220

Configuration overview	220
Changing passwords using the password tool	225
Single server installation where components use the same logon credentials	225
Single server installation where components use different logon credentials	227
Multiple server installation where Tivoli Integrated Portal is remote	229
Multiple server installation where DB2 is remote	231
Configuration Utility	234
Configuration utility tasks	234
Configure Devices wizard	240
Configuring storage subsystems	241
Configure fabrics and switches	251
Configure computers	253
Configure tape libraries	256
Managing schedules, jobs, and runs	257
Job Management - Recommendations	264
Configuration nodes	265
Role-to-Group Mappings	265
License Keys	268
Alert disposition	273
Log File and Cached Batch Report Retention	275
Quota and Constraint e-mail Address Rules	276
Scan/Probe Agent Administration	276
Manual NAS Server Entry	277
History Aggregator	280
Storage Resource agent Deployments	281
Storage Resource Agent Upgrades	295
Migrate Data/Fabric Agent	296
Update Storage Subsystem Credentials	297
Resource History Retention.	300
Removed Resource Retention	301
Resource History Retention for Databases	302
Removed Resource Retention for Databases	303
Configuration History Settings	303
Configuration: Device server specific	304
Configuring switches.	304
Configuring the inactive zone sets for a fabric data source	306
Managing a SAN without agents	306
Setting timeouts values for the Device server	306
Service Location Protocol (SLP) configuration considerations	308
Configuring Tivoli Storage Productivity Center for a Windows domain	310
Configuring the Data server for a Windows domain	310
Granting local administrative privileges to a domain account	311
Adding Tivoli Storage Productivity Center to the Windows domain.	312
Configuring IP addressing	314
Configuring Tivoli Storage Productivity Center with multiple IP addresses	314
Changing the HOSTS file	315
Configuring tracing	316
Configuration guidelines for 500 or more agents	317
Configuring LUN provisioning for Oracle Solaris	317

Assigning TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUNs to Oracle Solaris HBAs	318
Modifying the HBA configuration file	318
Setting Persistent Name Binding for QLA2200 and QLA2300 HBAs by using the <code>scfx</code> command	318
Modifying the SCSI disk configuration file	319
Checking for TotalStorage Enterprise Storage Server, DS6000, or DS8000 multipaths in VxDMP	320
How to import authentication information for a Storage Resource agent	321
Configuring Tivoli Storage Productivity Center for DS8000 LDAP authentication	322
Configuring multiple Tivoli Storage Productivity Center servers with one DS8000 R4.2	326
Setting up dual Tivoli Storage Productivity Center servers for high availability.	327
Configuring Tivoli Storage Productivity Center and SAN Volume Controller or Storwize V7000 for single sign-on	329
Configuring and controlling the Tivoli Storage Productivity Center Monitoring Agent	332
Configuring the Tivoli Storage Productivity Center Monitoring Agent on Windows	332
Configuring the Tivoli Storage Productivity Center Monitoring Agent on AIX or Linux	333
Starting and stopping the Tivoli Storage Productivity Center Monitoring Agent	333
Installing and configuring the Tivoli Storage Productivity Center server with multiple NIC cards	334
Creating an SSH certificate for the root user ID	336
Configuring DB2, AIX, and Linux for IPv6-only environment.	337
Changing the timeout value for the web-based GUI session	338

Chapter 4. Upgrading and migrating Tivoli Storage Productivity Center . . . 341

Choosing the upgrade method	343
Upgrading Tivoli Storage Productivity Center for Replication	344
Fix packs and patches	344
Storage Subsystem Credential Migration tool	345
Upgrading DB2 from DB2 9.1 or DB2 9.5 to DB2 9.7	345
Preparing for an upgrade	348
Backups	349
Upgrade response file	349
Upgrading Tivoli Storage Productivity Center in a single-server environment	351
Upgrading Tivoli Storage Productivity Center by using the wizard	351
Upgrading Tivoli Storage Productivity Center by using silent mode	352
Upgrading Tivoli Storage Productivity Center in a multiple-server environment	354
Upgrading Tivoli Storage Productivity Center with a remote database by using the wizard	355

Upgrading Tivoli Integrated Portal and Tivoli Common Reporting by using the wizard	356
Upgrading Tivoli Storage Productivity Center with a remote database by using silent mode.	357
Upgrading Tivoli Integrated Portal and Tivoli Common Reporting by using silent mode	360
Deleting duplicate entries in a DB2 table when an upgrade to Tivoli Storage Productivity Center Version 5.1 fails.	363
Upgrading Storage Resource agents	363
Upgrading the Storage Resource agent by using the Data Sources node	364
Upgrading the Storage Resource agent by using the Configuration node	364
Upgrading Storage Resource agents by using a command	365
Migrating Data agents and Fabric agents to Storage Resource agents	365
Migrating the Data agent or Fabric agent using the user interface	366
Migrating Data agents and Fabric agents using a command	366
Migrating Data agents and Fabric agents by using the upgrade wizard	368
Upgrading CIM agents for storage systems	368
Upgrading CIM agents	368
Migrating BIRT reports	369
Upgrading System Storage N Series Gateway servers	370

Chapter 5. Uninstalling Tivoli Storage Productivity Center 375

Uninstalling Tivoli Storage Productivity Center in a single-server environment by using the wizard	375
Uninstalling Tivoli Storage Productivity Center in a multiple-server environment by using the wizard	376
Uninstalling Tivoli Storage Productivity Center by using silent mode	378
Uninstalling the agents	380
Deleting or uninstalling an agent using the GUI	380
Uninstalling the Storage Resource agent manually	381
Uninstalling the Tivoli Storage Productivity Center Monitoring Agent	382
Uninstalling the IBM Tivoli Storage Productivity Center Monitoring Agent on Windows	382
Uninstalling the IBM Tivoli Storage Productivity Center Monitoring Agent on UNIX or Linux	382
Uninstalling DB2	383
Uninstalling DB2 on UNIX or Linux.	383
Deleting Storage Resource agent registry entries after a failed installation or uninstallation	385
Manually uninstalling Tivoli Storage Productivity Center components on AIX or Linux operating systems	386
Manually uninstalling Tivoli Storage Productivity Center components on Windows	387

Chapter 6. Reference 391

agent.sh command	391
----------------------------	-----

**Appendix A. Accessibility features for
Tivoli Storage Productivity Center . . . 393**

**Appendix B. Accessibility features for
Tivoli Storage Productivity Center for
Replication 395**

Notices 399
Trademarks 401

Glossary 403

Index 409

Figures

1. Global security page, configure federated repositories 208
2. Federated repositories page, Manage repositories 208
3. Manage repositories page, Add a new repository. 209
4. Manage repositories > New page 210
5. Global security > Federated repositories page, Add base entry to realm 211
6. Global security > Federated repositories > Repository reference page, Add base entry to realm 212
7. Launch the WebSphere administrative console 215
8. Opening the Global security page. 216
9. Configuring the available realm definitions 216
10. Removing the LDAP repository 217
11. Managing the repositories 217
12. Deleting the LDAP repository 218
13. Update Storage Subsystem Credentials panel 297
14. Save Confirmation prompt 370
15. Add NAS Server panel 371
16. Verify that the N Series Gateway server is added 372
17. File system Agent Editor panel. 372
18. Scan/Probe Agent Administration panel 372
19. View information for N Series Gateway server 373
20. View information for a file system 374

Tables

1. Locations of publications for Tivoli Storage Productivity Center and related products . . .	xii
2. Hardware requirements for the Tivoli Storage Productivity Center servers	1
3. Storage system requirements for the Replication server	3
4. Hardware requirements for the Storage Resource agents	3
5. Hardware requirements for the Tivoli Storage Productivity Center command-line interface . . .	4
6. Operating system support for Tivoli Storage Productivity Center server, Tivoli Storage Productivity Center for Replication server, and GUI	5
7. Operating system support for the Storage Resource agents	8
8. Functions with each Tivoli Storage Productivity Center license	17
9. Available reports with each IBM Tivoli Storage Productivity Center license	20
10. Tivoli Storage Productivity Center licenses when you upgrade from Version 4 to Version 5	22
11. Operating system groups that are mapped to the Superuser role	27
12. DB2 groups that are required to install the database repository	27
13. Valid characters for user names and passwords	29
14. User names and passwords for Tivoli Storage Productivity Center installation	31
15. TCP/IP ports used by Tivoli Storage Productivity Center	33
16. Other TCP/IP ports used by Tivoli Storage Productivity Center	34
17. TCP/IP ports used by Tivoli Storage Productivity Center for Replication for incoming communication	35
18. TCP/IP ports used by Tivoli Storage Productivity Center for Replication for outgoing communication	35
19. Multipath subsystem device drivers supported by Tivoli Storage Productivity Center	42
20. IPv4 and IPv6 configurations in a multiple-servers environment	47
21. Main features of the Fabric Manager	85
22. Supported and preferred interfaces for fabric management	86
23. Supported and preferred interfaces for fabric management	89
24. Information collected by agent	92
25. Supported platforms and storage systems for VMware	118
26. Tivoli Storage Productivity Center configurations	133
27. Operating system for Storage Resource agents	138
28. Options for the CHOSEN_INSTALL_TYPE parameter.	144
29. Parameters required for each protocol	170
30. Storage Resource agent installation images	170
31. DB2 Setup wizard default user and group names	181
32. Results from a partial rollback	193
33. Required user ID for initial log on after installation on a single server when the common user and a separate user for Tivoli Integrated Portal and Tivoli Common Reporting are defined.	202
34. Required user ID for initial log on after installation on multiple servers with a remote database schema	203
35. Required user ID for initial log on after installation on multiple servers with remote Tivoli Integrated Portal and Tivoli Common Reporting.	203
36. Roles in Tivoli Integrated Portal and Tivoli Common Reporting	219
37. Roles and authorization levels	266
38. Database permissions	268
39. Location of MIB and Baroc files	274
40. Upgrading Tivoli Storage Productivity Center	342

Preface

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

Who should read this guide

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

Administrators and installers should be familiar with the following topics:

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

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

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

Publications

This section lists publications in the IBM Tivoli Storage Productivity Center library and other related publications. It also describes how to access publications online, how to order publications, and how to submit comments on publications.

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

Accessing publications online

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

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

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

Product	Online location (see Note)
IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication	<p>http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp. In the left-hand navigation pane, click IBM Tivoli Storage Productivity Center.</p> <p>To obtain PDF documents, click IBM Tivoli Storage Productivity Center > Printable documentation.</p> <p>To view previous versions of the IBM Tivoli Storage Productivity Center Information Center, from the main page click Community and Support > Previous Versions.</p>
IBM Tivoli Storage Productivity Center for Replication for System z®	<p>http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp. In the left-hand navigation pane, click IBM Tivoli Storage Productivity Center for Replication for System z.</p> <p>To obtain PDF documents, click IBM Tivoli Storage Productivity Center for Replication for System z > Printable documentation.</p> <p>To view previous versions of the IBM Tivoli Storage Productivity Center for Replication Information Center, from the main page, click Community and Support > Previous Versions.</p>
IBM WebSphere® Application Server	<ul style="list-style-type: none"> • http://publib.boulder.ibm.com/infocenter/wasinfo/v7r0/index.jsp • http://publib.boulder.ibm.com/infocenter/wasinfo/v8r0/index.jsp
IBM System Storage® Productivity Center	<p>http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp. In the left-hand navigation pane, click System Storage Productivity Center.</p> <p>To obtain PDF documents, click System Storage Productivity Center > Printable documentation.</p> <p>To view previous versions of the IBM System Storage Productivity Center Information Center, from the main page click Community and Support > Previous Versions.</p>
IBM System Storage DS3000, IBM System Storage DS4000®, or IBM System Storage DS5000	<p>http://www.ibm.com/support/entry/portal/</p>
IBM System Storage DS6000™	<p>http://publib.boulder.ibm.com/infocenter/dsichelp/ds6000ic/index.jsp</p>
IBM System Storage DS8000®	<p>http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp</p>
IBM System Storage DS® Open Application Programming Interface publications	<p>http://www.ibm.com/support/entry/portal/Troubleshooting/Hardware/System_Storage/Storage_software/Other_software_products/CIM_Agent_for_DS_Open_(API)/</p> <p>Use these publications for information about how to install, configure, and use the CIM agent.</p>
IBM System Storage SAN Volume Controller	<p>http://publib.boulder.ibm.com/infocenter/svc/ic/index.jsp</p>

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

Product	Online location (see Note)
IBM Storwize® V7000	http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp
IBM Storwize V7000 Unified	http://publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp
IBM XIV® Storage System	http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp
IBM DB2 Database for Linux, UNIX, and Windows	http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp
IBM System Storage N series	http://www.ibm.com/systems/storage/network/redbooks.html For more information about IBM System Storage N series, see http://www.ibm.com/systems/storage/network/software/ .
IBM Systems Director	http://www.ibm.com/systems/software/director/index.html
Note: If you print PDF documents on other than letter-sized paper, in the Print window select Fit to Printable Area in the Page Scaling field. This option is available when you click File > Print . Fit to Printable Area ensures that the full dimensions of a letter-sized page print on the paper that you are using.	

IBM Redbooks

The IBM Redbooks® are books on specialized topics.

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

Translation

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

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

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

Ordering publications

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

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

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

Providing feedback about publications

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

Go to the information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp>. From any topic page, click the **Feedback** link.

Contacting the IBM Support Center

You can contact IBM Support Center in several ways.

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

To receive future support notifications, sign in under **Notifications**. You are required to enter your IBM ID and password. After you are authenticated, you can configure your subscription for Tivoli Storage Productivity Center technical support website updates.

- Customers in the United States can call 1-800-IBM-SERV (1-800-426-7378).
- For international customers, go to the Tivoli Storage Productivity Center technical support website to find support by country. Expand **Contact support** and click **Directory of worldwide contacts**.

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

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

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

Reporting a problem

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

Have the following information ready when you report a problem:

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

Conventions used in this guide

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

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

The following typeface conventions are used in this publication:

Bold

- 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
- 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 notations, remember the following:

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

New for Tivoli Storage Productivity Center Version 5.1

Learn about new features and enhancements in IBM Tivoli Storage Productivity Center Version 5.1.

New Features

The following features are new in Tivoli Storage Productivity Center 5.1:

Graphical user interfaces

Tivoli Storage Productivity Center provides two graphical user interfaces for managing the storage infrastructure in an enterprise environment: the stand-alone GUI and the web-based GUI. Each GUI provides different functions for working with monitored resources.

Viewing reports from the web-based GUI

From the web-based GUI, you can now run the following types of report in Tivoli Common Reporting:

Resources reports

Show the relationships between storage resources in your environment. Resources reports also provide a snapshot of the capacity, available space, and used space on servers, storage virtualizers, storage systems, and related storage resources.

Performance reports

Show performance statistics for the most active storage resources in your environment over a specified period. Performance reports also provide detailed statistics for specific storage resources.

Historical space reports

Show how space is used over a specified period on storage virtualizers, storage systems, and related storage resources.

Storage tiering reports

Show volumes that might require migration from a pool in a storage resource group to a pool in another storage resource group. You also use the Tivoli Storage Productivity Center web-based GUI to view the workload activity reports for managed disk groups and virtual disks.

Creating and designing reports from the web-based GUI

From the web-based GUI, you can create custom reports and design custom reports. You can create and design the reports in report tools in Tivoli Common Reporting.

Custom reports can contain detailed information about the relationships between monitored resources, the properties of monitored resources, and detailed information about the performance of monitored resources. Unlike predefined reports, you can select the information that you want to include in custom reports.

You can create reports with basic formatting, or you can design more complex reports with advanced formatting features. The advanced formatting features include inserting charts that you can customize in reports, and inserting links that open related reports.

Installing Tivoli Storage Productivity Center

You can install Tivoli Storage Productivity Center in one of the following ways:

Installation wizard

The installation wizard provides a graphical user interface for installation.

Command-line interface by using silent mode

In silent mode, you run a command and provide the installation values in a response file. A silent-mode installation is useful if your system is running from a terminal that cannot display graphics.

Installation configurations

When you install Tivoli Storage Productivity Center, your system can be configured in one of the following ways:

Single server

All Tivoli Storage Productivity Center components are installed on one server.

Multiple servers

In the multiple-servers environment, you can have these configurations:

- DB2 and the Tivoli Storage Productivity Center database repository on a remote server and the rest of the Tivoli Storage Productivity Center components on a local server.
- DB2 and the Tivoli Storage Productivity Center servers, stand-alone GUI, command-line interface, and Storage Resource agent on one server, and IBM Tivoli Integrated Portal, Tivoli Common Reporting, and web-based GUI on a remote server.

Installation components

When you install Tivoli Storage Productivity Center, the following components are included:

- Database repository
- Data server
- Device server
- Replication server
- Stand-alone graphical user interface
- Web-based graphical user interface
- Command-line interface
- Storage Resource agent
- Tivoli Integrated Portal Version 2.2
- Tivoli Common Reporting Version 2.1.1

Lightweight Directory Access Protocol

You no longer have to configure Lightweight Directory Access Protocol (LDAP) as an authentication method during the installation. You configure LDAP as a post-installation task. The Tivoli Storage Productivity Center installation program establishes a default authentication configuration by using the federated repositories feature of the WebSphere Application Server.

For more information about LDAP and how to add an LDAP repository, see the Tivoli Storage Productivity Center information center and search for *Administering Tivoli Integrated Portal*.

Tivoli Storage Productivity Center licenses

The following licenses are available in Tivoli Storage Productivity Center 5.1:

IBM Tivoli Storage Productivity Center

This license contains the following functions:

- Data management. Data management includes enterprise-wide reporting and monitoring, and automated capacity provisioning for direct-attached storage (DAS), network-attached storage (NAS), and storage area networks (SAN) environments.
- Disk management. Disk management helps you manage SANs and heterogeneous storage from one console.
- Fabric management. Fabric management helps you manage the SAN fabric that connects the host systems and applications to the storage devices. This feature includes automatic resource and topology discovery, monitoring, alerts, and zone control.
- Tape management. Tape management helps you manage your tape libraries that must be SMI-S compliant.

Tivoli Storage Productivity Center for Replication

All of the Tivoli Storage Productivity Center for Replication functions, including the functions that were previously in the Two Site and Three Site Business Continuity licenses.

- Performance management. Disk and Fabric performance management functions.
- Device support. For the devices that you can use with Tivoli Storage Productivity Center, see <http://www.ibm.com/support/docview.wss?rs=40&uid=swg21386446>.

IBM Tivoli Storage Productivity Center Select Edition

This license contains all the functions of the IBM Tivoli Storage Productivity Center license.

IBM Tivoli Storage Productivity Center Advanced

This license contains all the functions in the Tivoli Storage Productivity Center license plus these additional analytical functions:

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

Tivoli Storage Productivity Center for data-related functions

- Scans
- Policy management
- Server probes
- Profiles
- Data Manager for Databases
- Data Manager for Chargeback

IBM Scale Out Network Attached Storage

IBM Scale Out Network Attached Storage (IBM SONAS) manages repositories of information in enterprise environments that require

large capacities, high levels of performance, and high availability. IBM SONAS uses up to 2 petabytes (PB) of storage in a large file system.

Policy-based storage tier reports

Use the Tivoli Storage Productivity Center command-line interface to issue commands that create user-defined storage tiering policies. Storage tiering policies set the conditions that must be met to determine whether a volume requires migration to another pool.

New for Tivoli Storage Productivity Center for Replication Version 5.1

Use this information to learn about new features and enhancements in IBM Tivoli Storage Productivity Center for Replication Version 5.1. This information highlights the changes since the last release of Tivoli Storage Productivity Center for Replication.

Failover operations that are managed by other applications

Applications such as the IBM Series i Toolkit, VMware Site Recovery Manager, and Veritas Cluster Server manage failover operations for certain session types and storage systems. If an application completes a failover operation for a session, the Severe status is displayed for the session. An error message is also generated for the role pairs for which the failover occurred.

For more information, go to the Tivoli Storage Productivity Center for Replication information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for *failover and failback*.

Additional support for space-efficient volumes in remote copy Tivoli Storage Productivity Center for Replication sessions

You can use extent space-efficient volumes as copy set volumes for the following IBM System Storage DS8000 session types:

- FlashCopy® (System Storage DS8000 6.2 or later)
- Metro Mirror (System Storage DS8000 6.3 or later)
- Global Mirror or Metro Global Mirror (System Storage DS8000 6.3 or later)

For more information, go to the Tivoli Storage Productivity Center for Replication information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for *copy sets*.

New features

The following features are new for Tivoli Storage Productivity Center for Replication version 5.1.

Reflash After Recover option for Global Mirror Failover/Failback with Practice sessions

You can use the **Reflash After Recover** option with System Storage DS8000 version 4.2 or later. Use this option to create a FlashCopy replication between the I2 and J2 volumes after the recovery of a Global Mirror Failover/Failback with Practice session. If you do not use this option, a FlashCopy replication is created only between the I2 and H2 volumes.

For more information, go to the Tivoli Storage Productivity Center for Replication information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for *session properties*.

No Copy option for Global Mirror Failover/Failback with Practice and Metro Global Mirror with Practice sessions

You can use the **No Copy** option with System Storage DS8000 version 4.2 or later. Use this option if you do not want the hardware to write the background copy until the source track is written to.

For more information, go to the Tivoli Storage Productivity Center for Replication information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for *session properties*.

Recovery Point Objective Alerts option for Global Mirror sessions

You can use the **Recovery Point Objective Alerts** option with IBM TotalStorage Enterprise Storage Server® Model 800, System Storage DS8000, and System Storage DS6000. Use this option to specify the length of time that you want to set for the recovery point objective (RPO) thresholds. The values determine whether a Warning or Severe alert is generated when the RPO threshold is exceeded for a role pair. The RPO represents the length of time in seconds of data exposure that is acceptable in the event of a disaster.

For more information, go to the Tivoli Storage Productivity Center for Replication information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for *session properties*.

StartGC H1->H2 command for Global Mirror sessions

You can use the StartGC H1->H2 command with TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000. This command establishes Global Copy relationships between site 1 and site 2, and begins asynchronous data replication from H1 to H2.

This command is available only for Global Mirror Failover/Failback and Global Mirror Failover/Failback with Practice sessions.

For more information about this command, go to the Tivoli Storage Productivity Center for Replication information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for *session commands*.

Export Global Mirror Data command for Global Mirror role pairs

You can use this option to export data for a Global Mirror role pair that is in a session to a comma-separated value (CSV) file. You can then use the data in the CSV file to analyze trends in your storage environment that affect your recovery point objective (RPO).

For more information, go to the Tivoli Storage Productivity Center for Replication information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for *exporting historical data*.

New chsess command parameters

The command-line interface command **chsess** has new parameters:

- **-reflash**
- **-dsRP0warning**
- **-dsRP0severe**
- **-enableHardenedFreeze**

For more information, go to the Tivoli Storage Productivity Center for Replication information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for *chsess*.

Chapter 1. Planning for Tivoli Storage Productivity Center

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

Planning for installation

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

The Tivoli Storage Productivity Center server must be a dedicated system for Tivoli Storage Productivity Center operations and not shared with other applications.

Hardware requirements

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

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

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

Item	Hardware requirements
Processor	For Tivoli Storage Productivity Center <ul style="list-style-type: none">• For Windows or Linux operating system: Intel Xeon or greater, at least 4 processors at 2.5 GHz CPUs each• For AIX operating system: IBM POWER5 or later, at least 4 processors at 2.3 GHz CPUs each
Memory	8 GB of RAM. If you have at least 4 GB but less than 8 GB of RAM, you can still install Tivoli Storage Productivity Center but you get a warning message.

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

Item	Hardware requirements
Disk space	<p>For a single server installation:</p> <ul style="list-style-type: none"> • For installations on Windows operating systems, 15 GB of free disk space is required. • For installations on Linux operating systems, 15 GB of free disk space is required. • The disk space distribution for AIX is: <ul style="list-style-type: none"> /usr - 220 MB /var - 1 MB /home - 2.0 GB /tmp - 150 MB /opt - 14.2 GB <p>Note: After you install Tivoli Storage Productivity Center and start collecting data, a large amount of disk space is required. The amount of data collected depends on many factors:</p> <ul style="list-style-type: none"> • How many devices you have • How long you keep the data • How frequently you collect data <p>Some users have experienced disk space usage of about 40 - 80 GB.</p>
Network interface card (NIC)	<p>If you have multiple NIC cards on the Tivoli Storage Productivity Center server, see "Installing and configuring the Tivoli Storage Productivity Center server with multiple NIC cards" on page 334 for special considerations.</p>
Console	<p>For AIX or Linux operating systems, a console or remote-connectivity application such as KDE, Remote Desktop, or VNC is required when you install Tivoli Storage Productivity Center.</p>
<p>Restriction: For AIX or Linux operating systems, when you use a Virtual Network Computing (VNC) client that is connected to a VNC server, you should not use the following key combinations:</p> <ul style="list-style-type: none"> • Alt + H (to start the Help page) • Alt + N (to proceed to next page) • Alt + P (to return to the previous page) • Alt + I (to start an installation action on the Summary page) <p>You must press Tab and then press Enter or the space bar after the button you want to press is selected. For example, to start the Help, instead of pressing Alt+H, press Tab until the Help button is highlighted and then press Enter or the space bar.</p>	

See Table 3 on page 3 for information about the storage system requirements for the Replication server.

Table 3. Storage system requirements for the Replication server

Storage system	Hardware requirements
IBM Storage System DS8000	<ul style="list-style-type: none"> • Release 2 minimum firmware level 6.2.410.30 or Release 3 minimum firmware level 63.1.32.0 • Advanced Copy services licenses • Optional Ethernet adapters for Tivoli Storage Productivity Center for Replication <ul style="list-style-type: none"> – Single image - feature code 1801 – Dual image - feature code 1802 and 1803
IBM System Storage DS6000	<ul style="list-style-type: none"> • Minimum firmware level 6.2.2.64 • Advanced Copy services licenses
IBM Enterprise Storage Server (ESS) 2105-800	<ul style="list-style-type: none"> • Minimum firmware level 2.4.4.72 • Feature codes 240 - PRC and 2240 - PRC
<ul style="list-style-type: none"> • IBM System Storage SAN Volume Controller 4.2.1, 4.3.0, 5.1.0, and 6.1 • IBM Storwize V7000 	Copy services licenses
Host adapter feature codes	One or more of the following host adapter feature codes: <ul style="list-style-type: none"> • 3021 - Fibre Channel/FICON Host Adapter • 3023 - Fibre Channel/FICON Host Adapter • 3024 - 2 Gb Fibre Channel/FICON Host Adapter • 3025 - 2 Gb Fibre Channel/FICON Host Adapter

See Table 4 for information about the minimum hardware requirements needed for the Storage Resource agents.

Table 4. Hardware requirements for the Storage Resource agents

Operating System	Item	Hardware requirements
<ul style="list-style-type: none"> • Windows and Linux on Intel • UNIX • AIX on IBM Power Systems™ 	Processor	Pentium 400 [®] MHz processor, or higher
	Disk space and RAM	Minimum memory and space requirements to run: <ul style="list-style-type: none"> • 256 MB of RAM. • Approximately 350 MB of hard disk space

Table 4. Hardware requirements for the Storage Resource agents (continued)

Operating System	Item	Hardware requirements
All	Network interface card (NIC)	<p>Tivoli Storage Productivity Center supports multiple NIC cards. When you install an agent on a machine locally or remotely and the machine has more than one NIC card, the Tivoli Storage Productivity Center agent installation program determines the NIC card to use. If there are no NIC cards that can be used for two-way communication, the installation program returns an error message.</p> <p>For more information about multiple NIC cards, see "Installing and configuring the Tivoli Storage Productivity Center server with multiple NIC cards" on page 334.</p>

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

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

Operating system	Item	Hardware requirements
Windows	Processor	Minimum is PII 500 MHz processor, suggest PII 1 GHz processor or higher
	Disk space and RAM	<ul style="list-style-type: none"> • 256 MB of RAM • Approximately 525 MB of hard disk space
	Console	1024 x 768 console or higher
UNIX or Linux	Disk space and RAM	<ul style="list-style-type: none"> • 256 MB of RAM • Approximately 525 MB of hard disk space
	Console	1024 x 768 console or higher

Software requirements

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

Software requirements for operating systems

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

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

Before installing Tivoli Storage Productivity Center, check the Tivoli Storage Productivity Center support site for the latest platform support. Go to the Tivoli Storage Productivity Center support website at <https://www.ibm.com/support/docview.wss?uid=swg21386446>.

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

Operating system or platform		Tivoli Storage Productivity Center server	Tivoli Storage Productivity Center for Replication server	Tivoli Storage Productivity Center stand-alone GUI and command-line interface (CLI)
AIX, see "Notes for AIX operating systems"	IBM AIX 6.1 (64-bit) <ul style="list-style-type: none"> • POWER5 • POWER6® 	Yes	Yes	Yes
	IBM AIX 7.1 (64-bit) <ul style="list-style-type: none"> • POWER5 • POWER6 • POWER7® 	Yes	Yes	Yes
Linux, see "Notes for Linux operating systems"	Red Hat Enterprise Linux Base Server and Advanced Platform 5 (x86-64)	Yes	Yes	Yes
	Red Hat Linux Base Server and Advanced Platform 6 (x86-64)	No (This operating system is supported in a multiple-server environment where DB2 and Tivoli Storage Productivity Center are installed on Red Hat Linux. You can install DB2 or DB2 and Tivoli Storage Productivity Center on this operating system.)	No (This operating system is supported in a multiple-server environment where DB2 and Tivoli Storage Productivity Center are installed on Red Hat Linux. You can install DB2 or DB2 and Tivoli Storage Productivity Center on this operating system.)	Red Hat Linux Base Server and Advanced Platform 6 (x86-64)
Windows, see "Notes for Windows operating systems"	Windows 2008 (x86, x86-64) <ul style="list-style-type: none"> • Standard Edition • Enterprise Edition 	Yes	Yes	Yes
	Windows 2008 R2 (x86-64) <ul style="list-style-type: none"> • Standard Edition • Enterprise Edition 	Yes	Yes	Yes
	Windows XP (x86)	No	No	Yes
	Windows Vista (x86)	No	No	Yes
	Windows Vista SP2 (x86)	No	No	Yes
	Windows 7 (x86)	No	No	Yes

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

Operating system or platform		Tivoli Storage Productivity Center server	Tivoli Storage Productivity Center for Replication server	Tivoli Storage Productivity Center stand-alone GUI and command-line interface (CLI)
VM on VMware	VMware ESX and ESXi 3.0.x, 3.5.x, 4.0.x, and 4.1.x <ul style="list-style-type: none"> • Red Hat Enterprise Linux 5 • Red Hat Enterprise Linux 6 • Windows 2008 • Windows 2008 R2 	Refer to the corresponding operating system above for Tivoli Storage Productivity Center support in that operating system.	Refer to the corresponding operating system above for Tivoli Storage Productivity Center support in that operating system.	Refer to the corresponding operating system above for Tivoli Storage Productivity Center support in that operating system.
z/OS® V1.10, V1.11, V1.12 (31-bit)	-	No	Yes	No

Note: For AIX operating systems

- For AIX 6.1, the minimum maintenance level required is Technology Level 6100-04 Service Pack 5. You can determine the Technology Level and Service Pack installed by running the following command:

```
oslevel -s
```

You can download the package from <http://www.ibm.com/eserver/support/fixes/fixcentral/pfixpacks/61>.
- AIX 6.1.0.4 or later comes with IBM Systems Director. IBM Systems Director has a common agent service that can use port 9510. The Storage Resource agent also uses port 9510. If you do not need IBM Systems Director, you can uninstall IBM Systems Director before installing Tivoli Storage Productivity Center. If you install Tivoli Storage Productivity Center and require IBM Systems Director, specify a different port for the Storage Resource agent. For information about the ports IBM Systems Director uses, go to http://publib.boulder.ibm.com/infocenter/director/v6r2x/index.jsp?topic=/com.ibm.director.plan.helps.doc/fqm0_r_ports_for_the_management_server.html.
- AIX 6.1 comes with Lightweight Infrastructure (LWI). If this package is enabled, the package can cause Tivoli Storage Productivity Center installation issues:
 - Create a port conflict with Tivoli Storage Productivity Center at port 9510.
 - Cause the Storage Resource agent installation to fail on hosts where only the agent is installed.
- Use the GNU tar program version 1.14 or later rather than the native AIX tar program to extract files from Tivoli Storage Productivity Center installation images. The AIX tar program truncates long file names, which can cause errors in Tivoli Storage Productivity Center and Tivoli Integrated Portal installation.

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

Note: For Windows operating systems

Use the Windows setting for "Windows Classic" when installing the Tivoli Storage Productivity Center GUI on Windows Vista.

For more information, see "Installing Tivoli Storage Productivity Center client components on a separate computer" on page 161.

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

Operating system or platform	Tivoli Storage Productivity Center server	Tivoli Storage Productivity Center for Replication server	Tivoli Storage Productivity Center stand-alone GUI and command-line interface (CLI)
<p>Note: For Linux operating systems</p> <ol style="list-style-type: none"> If you are planning to install Tivoli Storage Productivity Center Version 5.1 on 64-bit Red Hat Enterprise Version 5 or 6, you must ensure that you have the complete X Window System package group installed on your 64-bit Red Hat Enterprise Linux system. For 64-bit Red Hat Enterprise Linux Version 5, Tivoli Storage Productivity Center requires the following packages (libraries) for Tivoli Common Reporting: <ul style="list-style-type: none"> compat-libstdc++-33-3.2.3-61.i386 or later compat-glibc-2.3.4-2.26.i386 or later openmotif22-2.2.3-18.i386 or later <p>Because Tivoli Common Reporting installs 32-bit binary files, you must install the 32-bit versions of these prerequisite packages (libraries), even on a 64-bit Red Hat Linux Version 5 system. Otherwise, the installation fails.</p> <p>In addition to the listed packages (libraries), Tivoli Storage Productivity Center also requires the following packages (libraries) for Red Hat Enterprise Linux Version 6:</p> <ul style="list-style-type: none"> libXtst-1.0.99.2-3.e16.i686.rpm or later 32-bit version of pam_krb5 (required on 32-bit and 64-bit system) <p>To determine whether the 32-bit (i386) package is installed and which version is installed, use the <code>--queryformat</code> option for the <code>rpm</code> command:</p> <pre># rpm -q --queryformat='%{N}-%{V}-%{R}.%{arch}\n' compat-libstdc++-33-3.2.3-61 compat-glibc-2.3.4-2.26 openmotif22-2.2.3-18 grep -v x86_64</pre> <pre>compat-libstdc++-33-3.2.3-61.i386 compat-glibc-2.3.4-2.26.i386 openmotif22-2.2.3-18.i386 #</pre> <p>To install the required packages (libraries), search the Red Hat Enterprise Linux installation media for these .rpm files:</p> <pre># find <Red Hat Enterprise Linux install media location> -name compat-libstdc++-*.i386.rpm # find <Red Hat Enterprise Linux install media location> -name compat-glibc-*.i386.rpm # find <Red Hat Enterprise Linux install media location> -name openmotif22-*.i386.rpm</pre> <p>After you locate the directory that contains the .rpm files, install or update these files to the required version. Use the <code>-U</code> option for the <code>rpm</code> command, for example:</p> <pre># rpm -U compat-libstdc++-33-3.2.3-61.i386.rpm compat-glibc-2.3.4-2.26.i386.rpm openmotif22-2.2.3-18.i386.rpm</pre> <p>Tip: The .rpm file versions on your installation media might differ from the preceding example.</p> <p>To verify that you have the .rpm files installed, use the <code>--queryformat</code> option for the <code>rpm</code> command to see the architecture of the .rpm files that were installed:</p> <pre># rpm -q --queryformat='%{N}-%{V}-%{R}.%{arch}\n' compat-libstdc++-33-3.2.3-61 compat-glibc-2.3.4-2.26 openmotif22-2.2.3-18 grep -v x86_64</pre> <pre>compat-libstdc++-33-3.2.3-61.i386 compat-glibc-2.3.4-2.26.i386 openmotif22-2.2.3-18.i386 #</pre> Before installing Tivoli Storage Productivity Center on Red Hat Enterprise Linux version 5.2 or later, ensure that the shadow utility requirements are met. For information about the shadow utility, go to the Tivoli Storage Productivity Center information center and search for <i>shadow password</i>. 			

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

Table 7. Operating system support for the Storage Resource agents

Platform		Storage Resource agent	File systems	Volume managers
AIX, see "Notes for AIX operating systems"	AIX 6.1 (64-bit) • POWER5 • POWER6	Yes	<ul style="list-style-type: none"> • JFS • JFS2 • GPFS™ 3.2 • NFS • Veritas File System (VxFS) r4 or later and r5 or later 	Veritas Volume Manager (VxVM) r4 or later and r5 or later
	AIX 7.1 (64-bit) • POWER5 • POWER6 • POWER7	Yes		
	<ul style="list-style-type: none"> • AIX Virtual I/O server 1.5.2 or later • AIX Virtual I/O server 2.2.1.0 or later 	Yes		
HP-UX, see "Notes for HP-UX operating systems"	HP-UX 11i v3 • Itanium 64-bit	Yes	<ul style="list-style-type: none"> • HP_HFS • NFS • Veritas File System (VxFS) r4 or later and r5 or later 	<ul style="list-style-type: none"> • HP-UX Logical Volume Manager • Veritas Volume Manager (VxVM) r4 or later and r5 or later
Red Hat Enterprise Linux, see "Notes for Linux operating systems"	<ul style="list-style-type: none"> • POWER5 (32-bit, 64-bit) • POWER6 (32-bit, 64-bit) • zSeries® (31-bit, 64-bit) 	Yes	<ul style="list-style-type: none"> • EXT2 • EXT3 • NFS • TMPFS • Veritas File System (VxFS) r4 or later and r5 or later <p>yep</p>	Veritas Volume Manager (VxVM) r4 or later and r5 or later
	Red Hat Enterprise Linux Base Server and Advanced Platform 5 (x86, x86-64): <ul style="list-style-type: none"> • POWER5 (32-bit, 64-bit) • POWER6 (32-bit, 64-bit) • zSeries (64-bit) 	Yes		
	Red Hat Enterprise Linux Base Server and Advanced Platform 6 (x86, x86-64): <ul style="list-style-type: none"> • POWER5 (32-bit, 64-bit) • POWER6 (32-bit, 64-bit) • zSeries (64-bit) 	Yes		

Table 7. Operating system support for the Storage Resource agents (continued)

Platform		Storage Resource agent	File systems	Volume managers
SUSE Linux Enterprise Server, see “Notes® for Linux operating systems”	SUSE Linux Enterprise Server 9 (x86, x86-64) <ul style="list-style-type: none"> • POWER5 (32-bit, 64-bit) • POWER6 (32-bit, 64-bit) • zSeries (31-bit, 64-bit) 	Yes	<ul style="list-style-type: none"> • EXT2 • EXT3 • NFS • REISERFS • TMPFS • Veritas File System (VxFS) r4.x and 5.x 	Veritas Volume Manager (VxVM) 4.x and 5.x
	SUSE Linux Enterprise Server 10 (x86, x86-64) <ul style="list-style-type: none"> • POWER5 (32-bit, 64-bit) • POWER6 (32-bit, 64-bit) • zSeries (64-bit) 	Yes		
	SUSE Linux Enterprise Server 11 (x86, x86-64) <ul style="list-style-type: none"> • POWER5 (32-bit, 64-bit) • POWER6 (32-bit, 64-bit) • zSeries (64-bit) 	Yes		
Oracle Solaris	Solaris 9 <ul style="list-style-type: none"> • SPARC architecture • 32-bit, 64-bit 	Yes (patches are required, see 1 on page 12)	<ul style="list-style-type: none"> • NFS • TMPFS • UFS • Veritas File System (VxFS) r4 or later and r5 or later 	Veritas Volume Manager (VxVM) r4 or later and r5 or later
	Solaris 10 <ul style="list-style-type: none"> • SPARC architecture • 32-bit, 64-bit 	Yes see 2 on page 12		

Table 7. Operating system support for the Storage Resource agents (continued)

Platform		Storage Resource agent	File systems	Volume managers
Windows, see “Notes for Windows operating systems”	Windows 2003 (x86, x86-64) <ul style="list-style-type: none"> • Data Center Edition • Enterprise Edition • Standard Edition 	Yes	<ul style="list-style-type: none"> • FAT • FAT32 • NTFS 	
	Windows 2003 R2 (x86, x86-64) <ul style="list-style-type: none"> • Data Center Edition • Enterprise Edition • Standard Edition 	Yes		
	Windows 2008 (x86, x86-64) <ul style="list-style-type: none"> • Data Center Edition • Enterprise Edition • Standard Edition 	Yes		
	Windows 2008 R2 (x86, x86-64) <ul style="list-style-type: none"> • Data Center Edition • Enterprise Edition • Standard Edition 	Yes		
VM on VMware, see “Notes for VMware”	VMware ESX 3.0.x, 3.5.x, 4.0.x, and 4.1.x (guest operating system): <ul style="list-style-type: none"> • Red Hat Enterprise Linux 5 (x86) • SUSE Linux Enterprise Server 9 (x86) • SUSE Linux Enterprise Server 10 (x86) • SUSE Linux Enterprise Server 11 (x86) • Windows 2003 (x86) • Windows 2008 (x86) 			

Table 7. Operating system support for the Storage Resource agents (continued)

Platform	Storage Resource agent	File systems	Volume managers
<p>Note: For Windows operating systems</p> <ol style="list-style-type: none"> For information about IPv6 support in Windows Server 2003 Enterprise Server, see http://support.microsoft.com/kb/325449. <p>Note: For AIX operating systems</p> <ol style="list-style-type: none"> The following AIX fixes must be applied: For AIX 6.1, the minimum maintenance level required are: <ul style="list-style-type: none"> Technology Level 6100-04 Service Pack 5 Technology Level 6100-05 Service Pack 1 Technology Level 6100-06 You can determine the Technology Level and Service Pack installed by running the following command: <pre>oslevel -s</pre> You can download the package from http://www.ibm.com/eserver/support/fixes/fixcentral/pfixpacks/61. AIX 6.1.0.4 or later comes with IBM Systems Director. IBM Systems Director has a common agent service that can use port 9510. The Storage Resource agent also uses port 9510. If you do not need IBM Systems Director, you can uninstall IBM Systems Director before installing Tivoli Storage Productivity Center. If you install Tivoli Storage Productivity Center and require IBM Systems Director, specify a different port for the Storage Resource agent. For information about the ports IBM Systems Director uses, go to http://publib.boulder.ibm.com/infocenter/director/v6r2x/index.jsp?topic=/com.ibm.director.plan.helps.doc/fqm0_r_ports_for_the_management_server.html. AIX 6.1 comes with Lightweight Infrastructure (LWI). If this package is enabled, the package can cause Tivoli Storage Productivity Center installation issues: <ul style="list-style-type: none"> There can be a port conflict with Tivoli Storage Productivity Center at port 9510. Can cause the Storage Resource agent installation to fail on hosts where only the agent is installed. Tivoli Storage Productivity Center also supports PowerHA[®] software on AIX for the following versions: <ul style="list-style-type: none"> 5.5.x 5.4.0.1 (APAR IY87447) or later 5.3.0.4 (APAR IY87534) or later Also check the PowerHA version compatibility matrix on the support site at http://www.ibm.com/support/techdocs/atmstr.nsf/84279f6ed9ffde6f86256ccf00653ad3/c03127c863a313898625711100635c80?OpenDocument Your fabric probe or zone control job can fail on AIX with error message AGT0430I. AIX has a constraint wherein the size of the response buffer passed in for GS-3 commands made to the switch cannot exceed 4 K. The Storage Resource agent passes 4 K buffer sizes for AIX. If the zone configuration for a fabric probe or zone control job exceeds the 4 K limitation for AIX, Tivoli Storage Productivity Center flags this condition. The fabric probe or zone control operation fails under this condition with error message AGT0430I. Consider using Storage Resource agents deployed on another platform that is connected to the same fabric. Use this Storage Resource agent for collecting zone information and for performing zone changes. In this case, the fabric functions for the Storage Resource agent that is failing can be disabled. If the fabric probe fails with the error message, and you try to make zone changes for the fabric, then you might get a warning message that the zoning has changed. Also, the zone control job log displays the AGT0432E error message. Use the GNU tar program version 1.14 or later, not the native AIX tar program, to extract files from Tivoli Storage Productivity Center installation images. The AIX tar program truncates long file names, which can cause errors in Total Storage Productivity Center and Tivoli Integrated Portal installation. The GNU tar file archiving program can be downloaded from http://www.ibm.com/systems/power/software/aix/linux/toolbox/alpha.html. 			

Table 7. Operating system support for the Storage Resource agents (continued)

Platform	Storage Resource agent	File systems	Volume managers
<p>Note: For Linux operating systems</p> <p>1. If you install the Storage Resource Agent on Red Hat Enterprise Linux Version 4 update 5 or later, RedHat Enterprise Linux Version 5, or RedHat Enterprise Linux Version 6, if some RPM package managers (RPMs) are not on your system, you must install them.</p> <p>For Storage Resource agents installed on Red Hat Linux AS release 4, you must have Update 5 or later. You must have the following package installed: <code>compat-libstdc++-33</code></p> <p>For Storage Resource agents installed on Red Hat Enterprise Linux release 5, you must have the following package installed: <code>compat-libstdc++-296-2.96-138.i386.rpm</code></p> <p>For Storage Resource agents installed on Red Hat Enterprise Linux release 6, you must have the following package installed: <code>compat-libstdc++-296-2.96144.el6.i686.rpm</code></p> <p>2. For Storage Resource agent for fabrics connected to hosts running on Linux IBM Power Systems, or zSeries hardware:</p> <ul style="list-style-type: none"> • Does not allow zone configuration changes • Does not report on HBA, fabric topology, or zoning information <p>These limitations also apply to Storage Resource agents on all guest operating systems for VMware configurations.</p> <p>3. Tivoli Storage Productivity Center does not support software RAID on Linux systems. A Storage Resource agent does not install and run on Linux systems with software RAID.</p>			
<p>Note: For HP-UX operating systems</p> <p>1. Tivoli Storage Productivity Center agents running on HP-UX must have the HP libc cumulative patch PHC_34275 installed on those machines. You can download patch PHC_34275 from https://www1.itrc.hp.com/service/home/home.do.</p>			
<p>Note: For Oracle Solaris operating systems</p> <p>1. To install Storage Resource agents on Solaris 9, patches 111711 and 111712 are required. For more information about downloading patches, go to the Solaris web site at http://www.oracle.com. You can search for and download the patches that you want to install.</p> <p>2. For Solaris 10, only Global Zone is supported for the Tivoli Storage Productivity Center agents. Local or non-global zones are not supported.</p> <p>If you upgrade from Solaris 9 to Solaris 10, reinstall all of the agents.</p>			
<p>Note: For VMware</p> <p>1. For more information about VMware platforms, see “Planning for VMware” on page 117.</p> <p>2. The Storage Resource agent supports the VMware ESX environment in tolerance mode only. No fabric functions are available with the Storage Resource agent for any guest operating systems.</p>			

Software requirements for the database repository

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

Tivoli Storage Productivity Center supports DB2 as the database repository.

DB2 database

For the repository database, you can use the following products with Tivoli Storage Productivity Center:

- DB2 Enterprise Server Edition version 9.5 with fix pack 4.
- DB2 Enterprise Server Edition version 9.7 with fix pack 4.

For the most current information about the supported database versions, see the Tivoli Storage Productivity Center support matrix at <https://www.ibm.com/support/docview.wss?uid=swg21386446>.

Restriction:

- If DB2 is installed in the LDAP environment, you cannot use Tivoli Storage Productivity Center.
- You must use a local operating system user name for DB2.

DB2 9.5

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

Tip:

- For DB2 9.5 on AIX, a 64-bit AIX kernel is required.
- For DB2 9.5, do not use fix pack 6 (this fix pack causes issues with Tivoli Storage Productivity Center).

DB2 9.7

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

Note:

- For DB2 9.7, use the version of DB2 shipped with Tivoli Storage Productivity Center.
- For DB2 9.7 on AIX, a 64-bit AIX kernel is required.

The Windows operating system is enabled for both IPv6 and IPv4 addressing. If the Windows operating system is used in IPv6-only scenarios, all communication must be done over IPv6. In this scenario, the Windows operating system still has an IPv4 stack active, even if no IPv4 address is configured.

When you install Tivoli Storage Productivity Center, a single, shared database instance is created on DB2. The default database name is TPCDB.

Web browser support

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

Note: The Tivoli Storage Productivity Center software is bundled with its own IBM Java JRE. When you use the Java Web Start method of remotely accessing the stand-alone GUI, you must use the IBM Java.

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

If you do not have a web browser configured for use with Tivoli Storage Productivity Center when you access a web page from its user interface, the is displayed. Use this window to configure a web browser for use with the product.

Tip: For information about the web browsers that you can use with the DS8000, see the DS8000 Information Center at <http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp>. Search for **Internet browser support**.

Browser-specific information: To download Firefox 3 for Linux or Windows, go to <ftp://ftp.mozilla.org/pub/mozilla.org/firefox/releases/>.

The software to block pop-up windows must be turned off in Windows Internet Explorer for any application to start the graphical user interface (GUI) for the following storage systems:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- IBM SONAS

For example, if you are in Tivoli Storage Productivity Center and you want to start the Element Manager of the SAN Volume Controller GUI, you must configure Internet Explorer with the following settings.

1. Open Internet Explorer.
2. From the Internet Explorer toolbar, click **Tools > Pop-up Blocker > Turn Off Pop-up Blocker**.

If the Internet Explorer Information Bar displays a message that content is blocked because it was not signed by a valid security certificate, click the Information Bar and select **Show blocked content**.

Web browser support for 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 is displayed in a web browser. Depending on the configuration settings for the web browser on your computer, Tivoli Storage Productivity Center for Replication is displayed within a tab of an existing browser session or it is displayed within a new browser session.

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

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

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

Note: If you start Tivoli Storage Productivity Center for Replication from within Tivoli Storage Productivity Center, you might be logged out of the Tivoli Storage Productivity Center for Replication user interface unexpectedly. This action occurs when you use Tivoli Storage Productivity Center to start a session of Tivoli Storage Productivity Center for Replication in an Internet Explorer web browser and then open and close a wizard or secondary window in that Tivoli Storage Productivity Center for Replication session.

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

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

Web browser support for Help menu items

Tivoli Storage Productivity Center provides a number of items in its **Help** menu that you can click to learn more about how to use the product. When you click one of these items, Tivoli Storage Productivity Center starts a web browser and displays the appropriate web page. You must have an internet connection on the system where you are running the Tivoli Storage Productivity Center user interface. You must also have a default browser configured to successfully access these web pages.

Software requirements for LDAP servers

The LDAP repositories that can be used with Tivoli Storage Productivity Center depend on the operating systems on which you run Tivoli Integrated Portal. For more information about the LDAP repositories supported, see one of the following topics, and search for *LDAP Servers using Federated Repository Configuration*.

- For the Windows operating system: <http://www.ibm.com/support/docview.wss?rs=180&uid=swg27012421> .
- For the AIX operating system: <http://www.ibm.com/support/docview.wss?rs=180&uid=swg27012389>.
- For the Linux operating system: <http://www.ibm.com/support/docview.wss?rs=180&uid=swg27012415>.

Software requirements for CIM agents

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

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

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

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

For information about the LSI CIM agent level required for each device and the namespace, see <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

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

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

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

CIM agents for switches and directors

CIM agents (SMI-S Providers) are also required for switches and directors. For a complete list of switches and directors supported and the level of CIM agents required for each device, see <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

CIM agents for non-IBM storage systems

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

For a complete list of storage systems supported and the level of CIM agents required for each device, see <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

Product licenses for Tivoli Storage Productivity Center

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

IBM Tivoli Storage Productivity Center

This license contains the following functions:

Data management

A file and capacity management solution for heterogeneous storage environments. Data management includes enterprise-wide reporting and monitoring, and automated capacity provisioning for Direct Attached Storage (DAS), network-attached storage (NAS), and Storage Area Network (SAN) environments.

Disk management

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

Fabric management

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

Tape management

Tape management helps you manage your SMI-S compliant tape libraries.

Tivoli Storage Productivity Center for Replication

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

Performance management

Disk and Fabric performance management functions.

Device support

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

IBM Storage Productivity Center Select Edition

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

IBM Tivoli Storage Productivity Center Advanced

This license contains everything in IBM Tivoli Storage Productivity Center license and the following functions:

Analytical functions

- Configuration analysis
- Configuration history
- SAN Planner
- Storage Optimizer

Tivoli Storage Productivity Center for data-related functions

- Scans
- Policy management
- Server probes
- Scans
- Profiles
- Data Manager for databases
- Data Manager for Chargeback

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

Table 8. Functions with each Tivoli Storage Productivity Center license

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	Tivoli Storage Productivity Center Advanced
Administrative Services		

Table 8. Functions with each Tivoli Storage Productivity Center license (continued)

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	Tivoli Storage Productivity Center Advanced
Device discovery	Yes	Yes
¹ Data Sources ¹¹		
<ul style="list-style-type: none"> • CIMOM agents • Data/Storage Resource Agents • Inband Fabric Agents • Out of Band Fabric Agents • Storage Subsystems • TPC Servers • VMWare VI Data Source 	<ul style="list-style-type: none"> • Yes • Yes • Yes • Yes • Yes • No • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes • Yes • Yes • Yes • Yes
Configuration	Yes ²	Yes
IBM Tivoli Storage Productivity Center		
Configuration Utility	Yes	Yes
Configure Devices	Yes	Yes
Job Management	Yes	Yes
Reporting	Yes	Yes
<ul style="list-style-type: none"> • Batch Reports • User ID reports • System Reports • Rollup Reports • Data Source Reports 	See Table 9 on page 20 for a list of available reports.	See Table 9 on page 20 for a list of available reports.
Topology Viewer	Yes	Yes
Monitoring (data collection)		
<ul style="list-style-type: none"> • Probes • TPC Server Probes 	<ul style="list-style-type: none"> • Yes • No 	<ul style="list-style-type: none"> • Yes • Yes
Storage Resource Group Management	Yes ⁵	Yes
Analytics		
<ul style="list-style-type: none"> • Configuration History • Storage Optimizer • SAN Planner • Configuration Analysis 	<ul style="list-style-type: none"> • No • No • No • No 	<ul style="list-style-type: none"> • Yes • Yes • Yes • Yes
Alerting:		
<ul style="list-style-type: none"> • Alert Log • Authentication Configuration Alerts 	<ul style="list-style-type: none"> • Yes³ • Yes 	<ul style="list-style-type: none"> • Yes • Yes
Data Manager		
Monitoring: (data collection)		

Table 8. Functions with each Tivoli Storage Productivity Center license (continued)

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	Tivoli Storage Productivity Center Advanced
<ul style="list-style-type: none"> • Groups • Pings • Scans • Profiles 	<ul style="list-style-type: none"> • Yes • Yes • No • No 	<ul style="list-style-type: none"> • Yes • Yes • Yes • Yes
Alerting	Yes ⁴	Yes
Policy Management	No	Yes
Reporting	Yes Not all reports are available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license. See Table 9 on page 20 for a list of available reports.	Yes
Data Manager for Databases	No	Yes
Data Manager for Chargeback	No	Yes
Disk Manager		
Storage Subsystems	Yes	Yes
<ul style="list-style-type: none"> • Create Volumes • Create Virtual Disks • Launch Element Manager 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes
Storage Optimizer	No	Yes
SAN Planner	No	Yes
Monitoring (data collection)		
<ul style="list-style-type: none"> • Groups • Jobs • Storage Subsystem Performance 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes
Alerting	Yes	Yes
Profile Management	Yes	Yes
Reporting		
<ul style="list-style-type: none"> • Groups • Storage Subsystems • Storage Subsystem Performance 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes
Fabric Manager		
Fabrics	Yes	Yes
Monitoring (data collection)		
<ul style="list-style-type: none"> • Groups • Jobs • Switch Performance Monitors 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes

Table 8. Functions with each Tivoli Storage Productivity Center license (continued)

Function	Tivoli Storage Productivity Center or Storage Productivity Center Select Edition	Tivoli Storage Productivity Center Advanced
Alerting	Yes	Yes
Reporting Switch Performance	Yes	Yes
Tape Manager	Yes	Yes
Element Manager	Yes	Yes
Replication Manager	Yes	Yes
<p>Note:</p> <ol style="list-style-type: none"> You cannot associate TPC Servers with Tivoli Storage Productivity Center in the Tivoli Storage Productivity Center or Storage Productivity Center Select license. VMware Data Sources are associated with the Tivoli Storage Productivity Center or Storage Productivity Center Select license. TPC Servers are associated with the Tivoli Storage Productivity Center Advanced license. The following nodes under Configuration are not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: <ul style="list-style-type: none"> License Keys Quota and Constraint e-mail Address Rules Resource History Retention for Databases Removed Resource Retention for Databases Configuration History Settings The following nodes under Alert Log are not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: <ul style="list-style-type: none"> Directory User OS User Group Configuration Analysis The following node under Data Manager > Alerting is not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: Directory Alerts. Hypervisor Alerts and Other NAS Alerts are available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license. You cannot create a storage resource group for a type of capacity pool. 		

Table 9. Available reports with each IBM Tivoli Storage Productivity Center license

Reports	Tivoli Storage Productivity Center or Storage Productivity Center Select	Tivoli Storage Productivity Center Advanced
Batch reports	Yes ¹ , ²	Yes ¹
System Reports: <ul style="list-style-type: none"> Data Fabric Disk 	<ul style="list-style-type: none"> Yes⁴ Yes Yes 	<ul style="list-style-type: none"> Yes Yes Yes
Rollup Reports	No	Yes
Data Source Reports	Yes	Yes
Asset reports		

Table 9. Available reports with each IBM Tivoli Storage Productivity Center license (continued)

Reports	Tivoli Storage Productivity Center or Storage Productivity Center Select	Tivoli Storage Productivity Center Advanced
<ul style="list-style-type: none"> • By Cluster • By Computer • By Hypervisor • By IBM Storwize V7000 Unified/IBM SONAS • By OS Type • By Storage Subsystem • System-wide 	<ul style="list-style-type: none"> • Yes • Yes • Yes • Yes • Yes • Yes • Yes³ 	<ul style="list-style-type: none"> • Yes • Yes • Yes • Yes • Yes • Yes • Yes
Availability reports	Yes	Yes
TPC-wide Storage Space reports: <ul style="list-style-type: none"> • Disk Space • File System Space • Consumed File System Space • Available File System Space 	<ul style="list-style-type: none"> • Yes • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes • Yes
Usage reports	No	Yes (these reports are a result of a scan)
Usage Violation reports	No	Yes (these reports are a result of a scan)
Backup reports:	No	Yes (these reports are a result of a scan)
Monitored Computer Storage Space reports	Yes	Yes
Storage Subsystem reports	Yes	Yes
Storage Subsystem Performance reports	Yes	Yes
Switch Performance reports	Yes	Yes
Tape Library report	Yes	Yes

Table 9. Available reports with each IBM Tivoli Storage Productivity Center license (continued)

	Tivoli Storage Productivity Center or Storage Productivity Center Select	Tivoli Storage Productivity Center Advanced
<p>Reports</p> <p>Note:</p> <ol style="list-style-type: none"> 1. A Storage Resource agent must be deployed on the computer where you want to run batch reports. 2. The following batch reports are not available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: <ul style="list-style-type: none"> • Rollup • Usage • Usage Violations • Backup 3. The following Asset > System-wide reports are available with the Tivoli Storage Productivity Center Advanced license: <ul style="list-style-type: none"> • Monitored Directories • Unmanaged Computers • Users • OS User Groups 4. The following Data system reports are available with the Tivoli Storage Productivity Center or Storage Productivity Center Select license: <ul style="list-style-type: none"> • Disk Space Summary • Disk Defects • Computer Storage Availability • Computer Disk Space • Available File System Space 		

If you purchased a license for IBM Tivoli Storage Productivity Center Version 4, when you upgrade to Version 5, you can upgrade to the appropriate license. Table 10 shows the parallel functions in Version 5.

Table 10. Tivoli Storage Productivity Center licenses when you upgrade from Version 4 to Version 5

If you have this Tivoli Storage Productivity Center Version 4 license	You are entitled to this Tivoli Storage Productivity Center Version 5 license
Basic	There is no equivalent license in Tivoli Storage Productivity Center Version 5.1. You can upgrade by purchasing the Tivoli Storage Productivity Center Version 5.1 license.
Disk Select	IBM Storage Productivity Center Select Edition
Data	IBM Tivoli Storage Productivity Center Advanced
Disk	IBM Tivoli Storage Productivity Center
Standard Select	IBM Tivoli Storage Productivity Center Advanced
Standard	IBM Tivoli Storage Productivity Center Advanced

Running Tivoli Storage Productivity Center in a 64-bit environment

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

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

- All Tivoli Storage Productivity Center programs that run in a 64-bit environment are run in 32-bit compatibility mode for Tivoli Storage Productivity Center.
- The Tivoli Storage Productivity Center agents must be running in a 32-bit native mode or 32-bit compatibility mode environment.
- The databases to be monitored by Tivoli Storage Productivity Center must be in a 32-bit native mode or 32-bit compatibility mode environment.
- Tivoli Storage Productivity Center can monitor DB2 in 64-bit native mode on Windows and AIX as long as the DB instance is created in 32-bit mode.

Planning for Tivoli Storage Productivity Center for Replication

When planning for disaster recovery, the Tivoli Storage Productivity Center for Replication server must be accessible if a disaster occurs. Even if you have a Tivoli Storage Productivity Center for Replication standby server, this server still requires authentication.

For more information about planning for replication, see “Planning for Open HyperSwap replication.”

Planning for Open HyperSwap replication

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

Overview

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

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

Software and hardware requirements

There are several requirements for Open HyperSwap support:

AIX requirements

Open HyperSwap support requires AIX version 5.3 or 6.1. (You can find the supported AIX version for each Tivoli Storage Productivity Center for Replication release in the support matrix at http://www.ibm.com/support/docview.wss?rs=40&context=SSBSEX&context=SSMN28&context=SSMMUP&context=SS8JB5&context=SS8JFM&uid=swg21386446&loc=en_US&cs=utf-8&lang=en. Click the link for the applicable release under Agents, Servers and GUI.)

You must have the following AIX modules installed:

- Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
- Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)

DS8000 hardware requirements

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

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

General tasks

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

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

Planning for Configuration

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

Configuration guidelines

Tivoli Storage Productivity Center must be configured properly for monitoring.

For a list of the capacity guidelines for Tivoli Storage Productivity Center, see <http://www.ibm.com/support/docview.wss?uid=swg21424912>.

Consider the following guidelines when you determine your capacity requirements:

Database repository

For best performance, install the database repository across multiple physical disks, either through operating system striping or hardware RAID. You must have a minimum of three disks that are separate from the Tivoli Storage Productivity Center host server operating system location and Tivoli Storage Productivity Center product installation directory location.

Tivoli Storage Productivity Center server

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

- The Tivoli Storage Productivity Center server must have dual 3.2-GHz processors with a minimum of 8 GB of RAM.
- The Tivoli Storage Productivity Center server must be a dedicated computer for Tivoli Storage Productivity Center operations and not shared with other applications.
- Paging space: On the AIX operating system, the paging space must be at least half or, if possible, the same amount as the physical memory. You must periodically monitor paging space usage by running the `lpsps -s` command and increasing the space if the percentage used is higher than 50%.

Stand-alone GUI

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

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

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

CIM agents

CIM agents are not required for storage systems that Tivoli Storage Productivity Center communicates with using the native interface, such as DS8000, IBM XIV Storage System, SAN Volume Controller, and IBM Storwize V7000. For those devices that require CIM agents to be set up, review the CIM agent documentation understand how many devices the CIM agent can be configured to manage. The memory consumed by the CIM agent includes the memory required for Tivoli Storage Productivity Center to probe a device through the CIM agent. The CIM agent is typically affected by the number of devices that the CIM agent manages. If no guidelines are available in the disk array subsystem CIM agent documentation, you should limit it to three subsystems per CIM agent.

Storage Resource agents

Storage Resource agents are designed to reduce the load of management calls on the fabric. Consider some improvements that Storage Resource agents offer: :

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

Storage subsystems

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

LSI SMI-S Provider

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

Planning for Tivoli Storage Productivity Center authentication and authorization

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

User requirements for installing Tivoli Storage Productivity Center

The user name that you use when you run the Tivoli Storage Productivity Center installation program must belong to the following operating system groups. The user that you define as the common user during the Tivoli Storage Productivity Center installation must also belong to these groups. The Tivoli Storage Productivity Center installation program automatically maps these groups to the Tivoli Storage Productivity Center Superuser role.

Table 11. Operating system groups that are mapped to the Superuser role

Operating System	Group
Windows	Administrators
AIX	system
Linux	root

The user name that is used to install the database repository must belong to one of the following DB2 groups. The user name that is used to install the database repository depends on your installation configuration. For example, if you installed Tivoli Storage Productivity Center on a single server and you did not enter separate user information for DB2, the database repository is installed by using the common user name. Therefore, the common user must have these DB2 privileges.

Table 12. DB2 groups that are required to install the database repository

Operating System	Group
Windows	DB2ADMNS
AIX and Linux	db2iadm1

The Tivoli Storage Productivity Center installation program establishes a default authentication configuration by using the federated repositories feature of the WebSphere Application Server. If you want to change the user authentication configuration, go to the Tivoli Storage Productivity Center information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp>. Search for *Changing the user authentication configuration*.

To authenticate user names on Tivoli Storage Productivity Center by using an instance of Microsoft Active Directory, see “Configuring Tivoli Storage Productivity Center for a Windows domain” on page 310.

Defining roles for Tivoli Storage Productivity Center

After you install Tivoli Storage Productivity Center, you can assign roles for Tivoli Storage Productivity Center users. These roles are predefined in Tivoli Storage Productivity Center and are assigned at the group level. The role determines the authorization level for the users who are in the group. The authorization level determines the functions in Tivoli Storage Productivity Center that the user can access.

For example, if a user is in a group that is assigned to the Superuser role, the user has full access to all Tivoli Storage Productivity Center functions. However, if a user is in a group that is assigned to the Tape Administrator role, the user has access to only tape-related functions.

For more information about roles and how to assign a role to a group, see “Role-to-Group Mappings” on page 265.

Defining roles for Tivoli Storage Productivity Center for Replication

You can assign roles for Tivoli Storage Productivity Center for Replication users at the user or group level. The roles are predefined in Tivoli Storage Productivity Center for Replication and determine the authorization level for individual users or all users who are in a group.

For example, if a user or a user group is assigned to the Administrator role, the user has full access to all Tivoli Storage Productivity Center for Replication functions. However, if a user or group is assigned to the Monitor role, the user can view the health and status of the Tivoli Storage Productivity Center for Replication components, but cannot complete any commands or actions.

For more information about roles and how to assign a role to a group, see the topic about Tivoli Storage Productivity Center for Replication security in the IBM Tivoli Storage Productivity Center User's Guide.

User names and passwords

Several user names and passwords are required to install, configure, and use IBM Tivoli Storage Productivity Center. There are also some requirements and limitations that you must understand before you install Tivoli Storage Productivity Center.

A worksheet is also provided to help you plan and document the user names and passwords that you use with Tivoli Storage Productivity Center. Keep the completed worksheet in a safe, secure place. Before you install Tivoli Storage Productivity Center, you must understand your local security policies and requirements. These policies can impose certain standards that you must consider when you create user names and passwords to comply with requirements.

Related reference:

“General requirements for user names and passwords”

This section provides information about the requirements for user names and passwords for Tivoli Storage Productivity Center.

“Worksheet for user names and passwords” on page 31

Use this worksheet to document the user names and passwords that you create when you install and administer IBM Tivoli Storage Productivity Center. Space is provided for you to list additional user names and passwords you create when you configure additional devices in Tivoli Storage Productivity Center.

General requirements for user names and passwords

This section provides information about the requirements for user names and passwords for Tivoli Storage Productivity Center.

Important: Do not create passwords that begin with a dash character.

DB2 user names and passwords

This topic lists rules for DB2 user names and passwords.

DB2 user names and passwords must follow these rules:

- UNIX user names and passwords cannot be more than eight characters long. They cannot begin with a numeric digit or end with \$.
- Windows 32-bit user names and passwords can contain 1 to 20 characters.
- Group and instance names can contain 1 to 8 characters.
- Names cannot be any of the following:
 - USERS
 - ADMINS
 - GUESTS
 - PUBLIC
 - LOCAL
- Names cannot begin with:
 - IBM
 - SQL
 - SYS
- Names cannot include accented characters.
- UNIX users, groups, and instance names must be lowercase.
- Windows 32-bit users, groups, or instance names can be any case.

User name and password requirements

This information helps you create valid user names and passwords for Tivoli Storage Productivity Center.

Table 1 lists the user names and passwords that are used with Tivoli Storage Productivity Center and specifies the following criteria:

- User names must be local or an LDAP account (a Windows Active Directory or a domain)

Restriction: If you use LDAP, Tivoli Storage Productivity Center must be installed and configured for LDAP authentication.

- Valid user name characters and minimum/maximum length requirements
- Valid password characters and minimum/maximum length requirements

For more information about setting passwords on the Windows operating system, see [http://msdn.microsoft.com/en-us/library/windows/desktop/ms679635\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/ms679635(v=vs.85).aspx)

Table 13. Valid characters for user names and passwords

User name and password for	Windows domain OK?	LDAP OK?	Valid characters
Tivoli Storage Productivity Center installation	Yes	Yes	Not applicable

Table 13. Valid characters for user names and passwords (continued)

User name and password for	Windows domain OK?	LDAP OK?	Valid characters
DB2 administrator, DB2 user	Yes Restriction: You must follow these naming conventions: <ul style="list-style-type: none"> Windows domain user name: <i>domain_name/user_name format</i> Local user name: <i>machine_name/username format</i> 	Yes	<ul style="list-style-type: none"> A through Z (uppercase characters) a through z (lowercase characters) 0 through 9 (numeric characters) Special characters: ~ @ # % ^ & () - _ { }. Restriction: User names and passwords cannot contain spaces and must have at least 1 character.
WebSphere administrator or host authentication password	Yes Restriction: You must follow these naming conventions: <ul style="list-style-type: none"> Windows domain user name: <i>domain_name/user_name format</i> Local user name: <i>machine_name/username format</i> 	Yes	<ul style="list-style-type: none"> A through Z (uppercase characters) a through z (lowercase characters) 0 through 9 (numeric characters) Special characters: - _.
NAS Filer user name	No	No	<ul style="list-style-type: none"> A through Z (uppercase characters) a through z (lowercase characters) 0 through 9 (numeric characters) Special characters: ` ~ # % ^ & () - _ { } ' . Restriction: User names and passwords cannot contain spaces and must have at least 1 character.

Table 13. Valid characters for user names and passwords (continued)

User name and password for	Windows domain OK?	LDAP OK?	Valid characters
NAS Filer password	No	No	<ul style="list-style-type: none"> A through Z (uppercase characters) a through z (lowercase characters) 0 through 9 (numeric characters) Special characters: ` ~ @ # % ^ & * () - _ = + [] { } \ ; : ' " , . < > / ? <p>Restriction: User names and passwords cannot contain spaces and must have at least 1 character.</p>
Tivoli Integrated Portal user name and password	Yes Restriction: You must follow these naming conventions: <ul style="list-style-type: none"> Windows domain user name: <i>domain_name/user_name format</i> Local user name: <i>machine_name/username format</i> 	Yes	<ul style="list-style-type: none"> A through Z (uppercase characters) a through z (lowercase characters) 0 through 9 (numeric characters) Special characters: ` ~ @ # % ^ & * () - _ = + [] { } \ ; : ' " , . < > / ? <p>Restriction: User names and passwords cannot contain spaces and must have at least 1 character.</p>

For information about the user names and passwords that are required to install and use Tivoli Integrated Portal, go to the Tivoli Storage Productivity Center information center and search for *User names and passwords for installing Tivoli Storage Productivity Center*.

Worksheet for user names and passwords

Use this worksheet to document the user names and passwords that you create when you install and administer IBM Tivoli Storage Productivity Center. Space is provided for you to list additional user names and passwords you create when you configure additional devices in Tivoli Storage Productivity Center.

Table 14. User names and passwords for Tivoli Storage Productivity Center installation

Item	Description	Your input
Tivoli Storage Productivity Center and DB2 installation user <ul style="list-style-type: none"> Windows: local administrator UNIX: root 		

Table 14. User names and passwords for Tivoli Storage Productivity Center installation (continued)

Item	Description	Your input
DB2 administrator user name and password ¹	This user name and password are required to install Tivoli Storage Productivity Center. On Windows, this user name must be a member of the DB2ADMNS group and Administrators group. On UNIX, the user name must be the instance owner of the instance you want to use. This user name and password are created when you install DB2.	
Tivoli Storage Productivity Center Device server host authentication password ²		
WebSphere administrator user name and password ³	This user name and password are required by the Device server to communicate with embedded WebSphere. Embedded WebSphere runs as a Windows Service. The Windows Service runs under the authority of this user name and password. This user name and password are only used when you install the Device server.	If this user name and password do not exist, they are created at installation time.
Tivoli Integrated Portal user name ⁴		
Tivoli Integrated Portal admin user name ⁵		
LDAP bind user name ⁶		
Storage subsystem user name for access from Tivoli Storage Productivity Center ⁷		
CIMOM / SMI provider login name ⁸		
NAS filer user name and password ⁹		

Table 14. User names and passwords for Tivoli Storage Productivity Center installation (continued)

Item	Description	Your input
Note:		
1.	DB2 admin user name - this user name and password is created when you install DB2 and is required to install Tivoli Storage Productivity Center.	
2.	Tivoli Storage Productivity Center device server host authentication password – this password is used by services that must connect to the Tivoli Storage Productivity Center Device server, such as the Storage Resource agent for SAN fabric functions.	
3.	WebSphere administrator - This user name and password are required to communicate with embedded WebSphere. This user name and password is created by any component that has an embedded WebSphere instance. DEFAULT: if not explicitly created, this defaults to the DB2 admin account.	
4.	Tivoli Integrated Portal user name – this is the user name that is used to login to TIP, which initially defaults to the user name that is used to install Tivoli Storage Productivity Center.	
5.	Tivoli Integrated Portal admin user name – This user name has administrator authority in Tivoli Integrated Portal. The default is 'tipadmin'.	
6.	LDAP bind user name – This user name has the authority to log in and connect to the LDAP server for authentication. It is not needed if the LDAP server allows anonymous bind. This user name and password must be provided to you by your LDAP administrator.	
7.	Storage subsystem user name for access from Tivoli Storage Productivity Center – this is a user name that Tivoli Storage Productivity Center uses to login to a storage device.	
8.	CIMOM / SMI provider login name – This user name that has log in access to the CIMOM or SMI provider software that interfaces with a storage device or a switch in a fabric.	
9.	9. NAS filer user name and password – Tivoli Storage Productivity Center uses this user name to log in to a NAS device.	

For information about valid user names and passwords, see “User name and password requirements” on page 29

TCP/IP ports used by Tivoli Storage Productivity Center

When you install Tivoli Storage Productivity Center, default ports must be opened through the firewall. You must disable the firewall program or open the ports to allow incoming requests to the Tivoli Storage Productivity Center ports. Review these ports before you install the Tivoli Storage Productivity Center.

TCP/IP ports used by Tivoli Storage Productivity Center

Table 15 lists the Tivoli Storage Productivity Center default ports.

Table 15. TCP/IP ports used by Tivoli Storage Productivity Center

Component	Default ports
Data server	9549
Device server	9550 - 9557
Replication server	9558 - 9566
Storage Resource agent	9567

Other TCP/IP ports used by Tivoli Storage Productivity Center

Table 16. Other TCP/IP ports used by Tivoli Storage Productivity Center

Port	Default port
DB2	50000
Tivoli Integrated Portal	The default port numbers for Tivoli Integrated Portal can vary. To view information about Tivoli Integrated Portal port numbers, see “Tivoli Integrated Portal port numbers” on page 35.
Tivoli Common Reporting	9581
SAN Volume Controller or IBM Storwize V7000	For native interface for SAN Volume Controller or IBM Storwize V7000 cluster: <ul style="list-style-type: none"> • 22 for SSH (outbound). This port is used to upload the SSH keys. This action is done once at setup time. • 5989 (outbound). Optionally, an SSH key can be uploaded once at setup time using this port.
CIM agent for IBM TotalStorage Enterprise Storage Server (ESS)	5989
DS8000	For native interface: Inbound for events: 1750 Outbound (on HMC): 1750 Inbound and outbound on HMC for offloading logs: 1755
IBM XIV Storage System	For native interface: Inbound: none Outbound (on XIV storage system): 7777, 7778
IBM Tivoli Storage Productivity Center with DS8000 GUI	8451 8452
Tivoli Storage Productivity Center	162 (default SNMP listening port)

Restriction: The Tivoli Storage Productivity Center 4.x ports are reused when you migrate from Tivoli Storage Productivity Center 4.x to 5.x.

The Tivoli Storage Productivity Center 5.x ports listed here apply only to the Tivoli Storage Productivity Center 5.x installations that were not migrated or upgraded from Tivoli Storage Productivity Center 4.x. For more information about Tivoli Storage Productivity Center 4.x ports, see http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=%2Fcom.ibm.sspc_v15.doc%2Ffqz0_r_sspc_planning_tcpip_ports.html

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

Table 17 on page 35 show the ports used for incoming communication. Table 18 on page 35 show the ports used for outgoing communication.

Table 17. TCP/IP ports used by Tivoli Storage Productivity Center for Replication for incoming communication

Port	Description
9558 and 9559	A web browser typically communicates with the Tivoli Storage Productivity Center for Replication GUI using HTTP ports 9558 and 9559.
9560	The Tivoli Storage Productivity Center for Replication command line and GUI typically communicate with the Tivoli Storage Productivity Center for Replication server using port 9560.
9561	Tivoli Storage Productivity Center for Replication uses port 9561 for communication with other Tivoli Storage Productivity Center for Replication servers for high-availability purposes.
Note: If you changed ports 9558 or 9559 during the installation or changed other port settings, note the values to which these settings were changed.	

Table 18. TCP/IP ports used by Tivoli Storage Productivity Center for Replication for outgoing communication

Port	Description
2433	For communication with ESS/DS (direct connection).
1750	For communication with DS8000 (HMC connection).
443 and 22	For communication with SAN Volume Controller or Storwize V7000 clusters.
9930	For communication with the AIX host for the Open HyperSwap.
<p>Note:</p> <ul style="list-style-type: none"> • If you changed the port configuration of your storage controller, the ports would be different. • Your network configuration must allow for Tivoli Storage Productivity Center for Replication to send outgoing TCP/IP packets to the storage controllers. It is possible when adding the storage controllers to Tivoli Storage Productivity Center for Replication to set a specific port number for your storage controller. • Because there are typically multiple applications running on a server, it is possible that port conflicts might arise if other applications attempt to use the same ports that Tivoli Storage Productivity Center for Replication is using. You can find the port numbers used on your system by running the following command: <pre>netstat -an</pre> • If firewalls are being used in your configuration, make sure that none of these ports are being blocked. Ensure that not only is the Tivoli Storage Productivity Center for Replication server granted access to reach the other components, but that the other components are granted access to reach the Tivoli Storage Productivity Center for Replication server. • If you are running Windows Server 2008, you must to configure the inbound and outbound rules for Tivoli Storage Productivity Center for Replication. To create an outbound rule, launch the New Outbound Rule wizard from the Windows Firewall with Advanced Security menu. • If you are installing or deploying Storage Resource agents on a system, you must first disable the firewall on that system before installing or deploying an agent. 	

Tivoli Integrated Portal port numbers

Depending on how you installed Tivoli Storage Productivity Center, the default port numbers for Tivoli Integrated Portal can vary.

To find the default port numbers that are used in Tivoli Integrated Portal URLs, do the following steps:

For a new installation of Tivoli Storage Productivity Center Version 5.1:

1. Go to the following directory:

On Windows operating systems:

C:\Program Files\IBM\tipv2\properties\

On Linux or UNIX operating systems:

/opt/IBM/tipv2/properties/

2. Open the TIPPortDef.properties file.
3. Review the following entry:

For http:// protocols:

WC_defaulthost_XXXX

For https:// protocols:

WC_defaulthost_secure_XXXX

For an installation of Tivoli Storage Productivity Center Version 5.1 that was upgraded from an earlier version:

1. Go to the following directory:

On Windows operating systems:

C:\Program Files\IBM\Tivoli\tipv2\properties\

On Linux or UNIX operating systems:

/opt/IBM/Tivoli/tipv2/properties/

2. Open the TIPPortDef.properties file.
3. Review the following entry:

For http:// protocols:

WC_defaulthost_XXXX

For https:// protocols:

WC_defaulthost_secure_XXXX

Port numbers for a new installation of Tivoli Storage Productivity Center Version 5.1

- If you installed Tivoli Storage Productivity Center Version 5.1 and used default Tivoli Integrated Portal port numbers, the port numbers are as follows:

For http:// protocols:

9568

For example:

http://hostname:9568/ibm/console/logon.jsp

For https:// protocols:

9569

For example:

https://hostname:9569/ibm/console/logon.jsp

- If you installed Tivoli Storage Productivity Center version 5.1 and used customized port numbers for Tivoli Integrated Portal, the URLs for a Tivoli Integrated Portal logon might appear as follows:

For http:// protocols:

For example:

http://hostname:WC_defaulthost/ibm/console/logon.jsp

For **https://** protocols:

For example:

```
https://hostname:WC_defaulthost_secure/ibm/console/  
logon.jsp
```

Port numbers for an upgraded installation of Tivoli Storage Productivity Center Version 4.2.2

- If you upgraded from Tivoli Storage Productivity Center Version 4.2.2 to Version 5.1, the default Tivoli Integrated Portal port numbers are as follows:

For **http://** protocols:

16310

For example:

```
http://hostname:16310/ibm/console/logon.jsp
```

For **https://** protocols:

16311

For example:

```
https://hostname:16311/ibm/console/logon.jsp
```

- If you used customized port numbers when you installed Tivoli Storage Productivity Center Version 4.2.2, upgraded to Version 5.1, and then used customized port numbers for Tivoli Integrated Portal, the URLs for a Tivoli Integrated Portal logon might appear as follows:

For **http://** protocols:

```
http://hostname:12345/ibm/console/logon.jsp
```

For **https://** protocols:

```
https://hostname:12346/ibm/console/logon.jsp
```

Storage systems, devices, file systems, databases, and switches

The information about the storage systems, tape libraries, file system formats, and databases supported by Tivoli Storage Productivity Center allows you to effectively set up and configure your environment.

Storage systems

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

You can use the following IBM storage systems:

- SAN Volume Controller (with native interface)
- IBM TotalStorage Enterprise Storage Server
- IBM storage systems (DS3000, DS4000, DS5000, DS6000 series)
- DS8000 series (with native interface)
- FAS*t*T
- The XIV system (with native interface)
- Storwize V7000 (with native interface)
- Storwize V7000 Unified (with native interface)
- IBM SONAS (with native interface)

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

For the Storage Optimizer, the DS4000 supports up to four Fibre Channel host connections and up to 112 SATA drives. SMI-S agents are not required for DS8000, the XIV system, SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, and IBM SONAS storage systems.

Restriction:

- A specific SAN Volume Controller must not be managed by more than one Tivoli Storage Productivity Center server at the same time.

For the most current information about storage systems supported, firmware and provider levels, see the Tivoli Storage Productivity Center support website at <https://www.ibm.com/support/docview.wss?uid=swg21386446>.

Switches

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

Tivoli Storage Productivity Center supports the following switches:

- Brocade

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

- Cisco

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

- IBM
- QLogic

For the most current information about switches and directors supported, see the Tivoli Storage Productivity Center support website at <https://www.ibm.com/support/docview.wss?uid=swg21386446>.

Tape libraries

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

For the most current information about tape libraries supported, see the Tivoli Storage Productivity Center support website at <https://www.ibm.com/support/docview.wss?uid=swg21386446>.

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

Note:

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

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

Databases for monitoring

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

The Data Manager can monitor these databases:

- DB2 version 9.1 (32-bit and 64-bit)
- DB2 version 9.5 (32-bit and 64-bit)
- DB2 version 9.7 (32-bit and 64-bit, fix pack 1 and fix pack 2 are not supported)
- Microsoft SQL Server 8
- Microsoft SQL Server 2005 (32-bit and 64-bit)

Note: There are some configuration steps required before you can monitor the Microsoft SQL Server 2005 database. For information about configuration, see “Configuring Microsoft SQL Server 2005 or 2008” on page 272.

- Microsoft SQL Server 2008 (32-bit and 64-bit)

Note: There are some configuration steps required before you can monitor the Microsoft SQL Server 2008 database. For information about configuration, see “Configuring Microsoft SQL Server 2005 or 2008” on page 272.

- Oracle 10g (32-bit and 64-bit)
- Oracle 10.1 and 10.2 (32-bit and 64-bit)

- Oracle 11g (32-bit and 64-bit)
- Oracle 11.1 and 11.2 (32-bit and 64-bit)
- Sybase Adaptive Server version 15 (32-bit and 64-bit)

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

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

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

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

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

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

You can do these steps by expanding **Administrative Services > Configuration > License Keys** in the navigation tree. For more information about this topic, see the online help or “License Keys” on page 268.

For information about the operating systems supported for the agent, see “Software requirements for operating systems” on page 4.

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

File systems

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

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

- AIX JFS, JFS2
- Data ONTAP V7 for Network Appliance, including flexible volumes (FlexVol). Flexvol enables you to create multiple flexible volumes on a large pool of disks.
- EXT2, EXT3, EXT4
- FAT, FAT32
- GPFS 3.2, 3.3
- HP_HFS
- NTFS4, NTFS5
- REISERFS
- SANFS

- TMPFS
- UFS
- VMFS
- Veritas File System (VxFS) r4.5 and 5.x on the following operating systems:
 - AIX
 - HP-UX
 - Linux
 - Oracle Solaris

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

- WAFL
- Tivoli Storage Productivity Center SAN File System

Networked file systems

This section describes support for IBM General Parallel File System.

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

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

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

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

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

Volume managers

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

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

- Veritas Volume Manager is supported on the following operating systems:
 - AIX
 - HP-UX
 - Linux
 - Oracle Solaris

- AIX Logical Volume Manager (LVM)
- HP-UX Logical Volume Manager

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

Planning for multipath subsystem device drivers

The subsystem device driver (SDD) is a software solution for multiple configuration environments in supported storage devices.

The subsystem device driver is installed on a host system with the native disk-device driver and provides the following functions:

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

For the most current support for multipath subsystem device drivers, see the following website: Supported Products and Platforms. Select the appropriate version for Agents, Servers, and GUI.

With Tivoli Storage Productivity Center, you can use the subsystem device drivers (SDD) shown in Table 19.

Table 19. Multipath subsystem device drivers supported by Tivoli Storage Productivity Center

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

Full support means that Tivoli Storage Productivity Center can:

1. Report on the disks accessed using the multipath subsystem device driver (probe).
2. Display multiple paths in the Data Manager asset tree (Data Manager Asset Tree).

3. Report on the multiple paths in the Data Path Explorer.
4. Configure paths as part of the SAN Configuration Planner.

Note:

1. The AIX SDD cannot coexist with SDDPCM on the same system.
2. The Linux SDD driver is no longer available. Starting with Red Hat Enterprise Linux Version 5 and SUSE Linux Enterprise Server Version 10, only the DM_Multipath is available. With DM_Multipath, you can use the round-robin policy only for the SAN Planner.
3. The SDD driver is no longer supported on HP, starting with HP-UX 11i v3 with Itanium®. HP-UX 11i v3 provides a native multipathing feature which is supported by SAN Planner.

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

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

Upgrading subsystem device drivers

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

EMC PowerPath Multipathing

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

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

Planning for Storage Resource agents

Storage Resource agents collect asset information, file and file system attributes, and other information that is needed from the computer system.

Storage Resource agents can also gather information about database managers installed on the server, and NAS device information. You can create ping, probe, and scan jobs to run against the servers that have Storage Resource agents installed.

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

You can deploy the Storage Resource agent through the stand-alone GUI user interface rather than an installation wizard. You must have administrative privileges to deploy the Storage Resource agent. The first time you deploy a Storage Resource agent to a computer system, the information collected by that agent is added to the logs on that computer system. That information also displays in the Tivoli Storage Productivity Center GUI after the next default system scan completes.

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

For more information about Storage Resource agents, see “Deployment guidelines and limitations for Storage Resource agents” on page 282.

Storage Resource agent protocol support

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

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

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

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

On-Demand service (non-daemon service)

In this case, connectivity between the server and agent is established using the same protocols as for installation of the agent: SSH, SMB, REXEC, or RSH.

run as a service (daemon service)

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

Note: Running the Storage Resource agent in daemon mode requires only one socket connection to be open and thus simplifies any firewall rules that might be in place. However, the daemon mode always consumes resources (although a very small amount) when idle. In non-daemon mode, the Storage Resource agent only consumes resources when actively working for the Tivoli Storage Productivity Center server. Because the non-daemon mode uses the RXA protocol, additional firewall ports and security configuration are required.

The information required for these protocols are as follows:

SSH protocol

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

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

For information about how to create a certificate for SSH protocol connectivity, see “Creating a certificate for SSH protocol” on page 286. This information uses Cygwin for the Storage Resource agent.

Windows SMB protocol

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

REXEC protocol

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

RSH protocol

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

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

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

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

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

For more information about setting up the environment for installing Storage Resource agents, see “Deployment guidelines and limitations for Storage Resource agents” on page 282.

Planning for Internet Protocol Version 6

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

Overview

You can install and run Tivoli Storage Productivity Center on systems that are enabled for IPv4, IPv6, or dual stack. *Dual stack* indicates that a machine has both the IPv4 and IPv6 stacks enabled and both addresses configured.

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

- SMI-S CIM agents
- SLP directory agents
- Out of band SNMP agents
- Storage devices
- SMTP server

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

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

IPv6 when installing Tivoli Storage Productivity Center

Tivoli Storage Productivity Center accepts both IPv4 and IPv6 addresses. If you have a system that is configured for dual stack networking, Tivoli Storage Productivity Center defaults to IPv4 addressing.

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. You can also specify IPv6 addresses by using shortened formats that omit leading zeros or use double colons in place of a series of zeros. You can specify only one double colon (::) in an IPv6 address.

You can use IPv6 addresses if one of the following conditions are met:

- The system where Tivoli Storage Productivity Center is installed is IPv6.
- Dual stack (IPv4 and IPv6) is enabled.

The following examples show an IPv6 address in the long form and in the short form:

Long form: `2001:DB8:0:0:0:0:0:0`

Short form: 2001:DB8::

If you are installing Tivoli Storage Productivity Center in an IPv6-only environment, the following requirements apply:

- For Windows operating systems, the IPv4 loopback interface is enabled.
- For AIX or Linux operating systems, the localhost option resolves to the IPv4 address 127.0.0.1, instead of the IPv6 address ::1. This value is specified in the /etc/hosts file.

For Windows operating systems, the localhost option is specified in the C:\Windows\system32\drivers\etc\hosts directory and file.

Table 20. IPv4 and IPv6 configurations in a multiple-servers environment

If the database repository server is configured for...	The Tivoli Storage Productivity Center server is configured for...
IPv4	IPv4
IPv4	Dual stack
Dual stack	IPv4
Dual stack	Dual stack
Dual stack	IPv6
IPv6	IPv6

Using IPv6 addresses in the user interface

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

Enter the host name (or IPv6 address) and port number of an Tivoli Storage Productivity Center server to which you want to connect.

Administrative Services > Data Sources >

Specify IPv4 or IPv6 addresses for the CIMOMs that are located on an IPv6 or dual stack machine.

Administrative Services > Discovery >

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

Administrative Services > Data Sources >

Specify IPv6 addresses for the Out of Band Fabric agents that are located on an IPv6 or dual stack machine.

Administrative Services > Data Sources >

Specify the names or IPv6 addresses of other Tivoli Storage Productivity Center servers that you want to add as subordinate servers for roll up probes and reports.

Administrative Services > Discovery >

Specify a range of IPv6 addresses for an SNMP discovery on an IPv6-only or dual stack machine against IPv6-enabled target machines.

Administrative Services > Configuration >

If the Tivoli Storage Productivity Center servers are running on an IPv6-only or dual stack machine enter the following information:

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

IBM Tivoli Storage Productivity Center > Configuration Utility >

Use an IPv4 or IPv6 address when you specify the URL for an external tool that you want to start from Tivoli Storage Productivity Center.

IPv6 considerations in Tivoli Storage Productivity Center

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

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

Starting the Tivoli Storage Productivity Center user interface using Java Web

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

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

Using IPv6 communication between master and subordinate servers (for roll up reporting)

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

Defining Tivoli Enterprise Console® events as triggered actions

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

Selecting files for archive and backup

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

Gathering data from XIV system

Tivoli Storage Productivity Center does not support the version of IBM XIV Storage System that supports IPV6.

Planning for the native interface

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

Support for the native interfaces

For the DS8000, Tivoli Storage Productivity Center communication with the storage system is through the ESSNI interface. Through this interface, Tivoli Storage Productivity Center collects device information, performance information, and can

detect changes and complete provisioning operations. This support simplifies the Tivoli Storage Productivity Center configuration and the launch in context configuration required to manage the storage system. Tivoli Storage Productivity Center also supports dual HMC environments.

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

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

Use the Configure Devices wizard to define native interface connection information for these storage systems as described in “Configure Devices wizard” on page 240.

You can access the wizard in one of the following ways:

- In the navigation tree, click **Administrative Services > Data Sources > Storage Subsystems** and click **Add** on the Storage Subsystems panel.
- In the navigation tree, click **Disk Manager > Storage Subsystems** and click **Add Storage Subsystem** on the Storage Subsystems page.
- Click the **Configure Devices** wizard icon in the toolbar.

Upgrading Tivoli Storage Productivity Center

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

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

- Use the migration tool that is shipped with Tivoli Storage Productivity Center
- Use the Tivoli Storage Productivity Center GUI to migrate the credentials. For more information, see “Update Storage Subsystem Credentials” on page 297.
- Run the Tivoli Storage Productivity Center upgrade installation program and specify that you want to use the migration tool. Run the migration tool when you upgrade Tivoli Storage Productivity Center.

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

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

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

Planning to use LDAP

The Lightweight Directory Access Protocol (LDAP) is an application protocol that you can use to query and modify directory services running over TCP/IP. The Tivoli Storage Productivity Center installation program establishes a default authentication configuration using the federated repositories feature of the WebSphere Application Server. You can add an LDAP repository to this authentication configuration as a post-install activity.

A directory is a set of objects with similar attributes that are organized in a logical and hierarchical manner. An LDAP directory tree often reflects various political, geographic, and organizational boundaries, depending on the model chosen. The directory might contain entries representing people, organizational units, printers, documents, groups of people, or anything else that represents a tree entry (or multiple entries).

In the federated repositories framework, the Tivoli Storage Productivity Center installation program creates the following repositories:

File-based user repository

This repository contains the user ID "tpcFileRegistryUser". This user password is the same as the common user password that you enter during the Tivoli Storage Productivity Center installation.

Operating system repository

This repository is referred to as the localOS repository. This repository contains the users and groups managed by the local operating system.

With this default authentication configuration, you cannot use the Tivoli Storage Productivity Center single sign-on feature. Storage system element managers do not support the localOS repository or the file-based repository for single sign-on, even if the element manager is installed on the same system as Tivoli Storage Productivity Center.

To change the user authentication configuration, add or remove an LDAP repository in the federated repositories framework and complete these tasks:

1. Back up the existing Tivoli Integrated Portal WebSphere configuration files.
The WebSphere backup files provide a restoration point if a failure occurs when you complete these procedures.
2. Back up the WebSphere configuration files after you add an LDAP repository to the user authentication configuration and want to later change the LDAP authentication settings.

To authenticate user names on the Tivoli Storage Productivity Center by using an instance of Microsoft Active Directory, complete one of the following options:

Authenticate User names: Method one

1. Install Tivoli Storage Productivity Center on an AIX, Linux, or Windows operating system.
2. Configure Tivoli Storage Productivity Center for LDAP authentication.

The target LDAP repository is an instance of Active Directory running on Windows 2003 or later. This method provides the following benefits:

- The Tivoli Storage Productivity Center target system does not have to be a Windows system.
- The Tivoli Storage Productivity Center target system does not have to be a member of a Windows domain.
- There are fewer configuration steps for Tivoli Storage Productivity Center or Windows.
- The Tivoli Storage Productivity Center user names and group names do not have to include the Windows Domain Name.
- You can use the single-sign on feature between Tivoli Storage Productivity Center and storage system element managers.

Authenticate User names: Method two

1. Install Tivoli Storage Productivity Center on a Windows system that is a member of a Windows domain or on the Windows Domain Controller.
2. Configure Tivoli Storage Productivity Center for operating system authentication.

Software requirements for CIM agents

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

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

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

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

For information about the LSI CIM agent level required for each device and the namespace, see <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

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

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

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

CIM agents for switches and directors

CIM agents (SMI-S Providers) are also required for switches and directors. For a complete list of switches and directors supported and the level of CIM agents required for each device, see <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

CIM agents for non-IBM storage systems

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

For a complete list of storage systems supported and the level of CIM agents required for each device, see <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

Planning for the single sign-on feature

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

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

The Tivoli Storage Productivity Center installation program establishes a default authentication configuration using the federated repositories feature of the WebSphere Application Server.

For information about changing the user authentication configuration, see the Tivoli Storage Productivity Center information center and search for *Administering Tivoli Integrated Portal*.

Additional considerations

- The single sign-on feature is not available for the IBM XIV Storage System element manager, XIV Storage Manager.
- The single sign-on feature is not supported by the Tivoli Storage Productivity Center command-line interface (CLI).

Planning for enterprise-rollup reports

Tivoli Storage Productivity Center enables you to use a single interface to generate reports based on data collected by multiple Tivoli Storage Productivity Center servers.

You can perform the following actions:

- View the storage information that is gathered by other servers from a single Tivoli Storage Productivity Center user interface using enterprise-wide rollup reports. These rollup reports enable you to have a network-wide perspective of storage usage in an environment where multiple Tivoli Storage Productivity Center servers are deployed and collecting storage metrics.

Note: Performance-related reports are not rolled up.

- Partition many agents across multiple Data servers. This action alleviates performance issues when running aggregation reports and scan jobs for multiple agents on one Data server.

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

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

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

- The master and subordinate servers can change over time.
- A master server can also manage entities just like any other server and report on these entities.
- We recommend that the master server monitor no more than 500 unique data sources. This number includes subordinate servers, Storage Resource agents, CIMOM agents, and VM servers (VMWare).
- We recommend that each subordinate server monitors no more than 1200 unique data sources. This number includes Storage Resource agents, CIMOM agents, and VM servers (VMWare). After this threshold is met for a server, a new server is deployed and all new agents point to it.
- You must have Tivoli Storage Productivity Center superuser or Tivoli Storage Productivity Center administrator authority to perform administrative operations such as adding a subordinate server to the master server list to monitor.
- After upgrading to Tivoli Storage Productivity Center 5.1, you must run probes against the monitored storage assets for the master server to have information about those assets displayed in **Tivoli Storage Productivity Center > Reporting > Rollup Reports**.
- When you add a subordinate server to the master server list, you must provide the host authentication password of the subordinate server to the master server. The host authentication password is checked each time the master server is probed. This host authentication password is entered when installing Tivoli Storage Productivity Center. Contact your system administrator or the person who installed Tivoli Storage Productivity Center to determine this password.

Planning for enterprise roll up reports configuration

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

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

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

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

Planning for the SAN configuration tools

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

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

Configuration Analysis

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

Configuration History

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

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

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

Planning for configuration analysis

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

The configuration analysis feature allows a system administrator to select up to 16 policies and to specify the name of a target SAN (or zone set) for the analysis. The analysis tool reads the specified SAN data and the policy information for the specified policies from the IBM Tivoli Storage Productivity Center database. The policy violations are displayed as Tivoli Storage Productivity Center alerts.

System administrators can use configuration analysis to learn of any best-practices violations their actions might have caused. For example, in provisioning a new storage subsystem, an administrator might inadvertently violate a zoning-related

policy. If the zoning-related configuration changes persist in the database, configuration analysis can detect this policy violation.

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

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

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

To use the configuration analysis tool, follow these general steps:

1. Ensure that you have previously run discovery and probe jobs for the computers, fabrics, switches, storage, and other objects of interest.
2. Ensure that the Device server and Data server are up and running.
3. Ensure that the Tivoli Storage Productivity Center agents are installed and configured.
4. Go to **IBM Tivoli Storage Productivity Center > Analytics**. Right-click **Configuration Analysis**. Use the Create Analyzer window to select the scope and policies for the configuration analysis.
5. Schedule how often you want the job to run and how to handle time zones.
6. Save and run the job.
7. After the configuration analysis job has run, you can view the alerts generated by IBM Tivoli Storage Productivity Center. In the navigation tree, expand **IBM Tivoli Storage Productivity Center > Alerting > Alert Log > Configuration Analysis**.

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

Planning for the Configuration History tool

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

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

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

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

A Configuration History topology view is displayed of the item you selected.

5. Select the time range for the snapshots and click **Apply**.
6. You can view the SAN configuration changes in the configuration history view for the two time periods you selected.

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

SAN Planner Overview

SAN Planner assists the user in end-to-end planning for fabrics, hosts, storage controllers, storage pools, volumes, paths, ports, zones, zone sets, storage resource groups (SRGs), and replication. You can implement a plan by using SAN Planner.

Supported Systems

You can use the following products with SAN Planner:

- TotalStorage Enterprise Storage Server
- System Storage DS8000
- DS6000

- SAN Volume Controller
- Storwize V7000
- IBM Storwize V7000 Unified block level support
- The XIV system

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

Requirements for SAN Planner

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

You can select host systems that do not have a Storage Resource agent installed, if the initiator ports for that host system are known to Tivoli Storage Productivity Center. The provisioning profile selected for these host systems can include only capacity and zone provisioning information. Path settings are not allowed for host systems with no Storage Resource agent, and the path planning option is unavailable in the SAN Planner wizard.

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

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

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

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

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

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

Note:

- The Linux SDD driver is no longer available. Starting with Red Hat Enterprise Linux Version 5 and SUSE Linux Enterprise Server Version 10, only the DM_Multipath is available. The DM_Multipath is supported by SAN Planner
- The SDD driver is no longer supported on HP, starting with HP-UX 11i v3 with Itanium. HP-UX 11i v3 provides a native multipathing feature which is supported by SAN Planner.
- If you are using the SAN planning tools which collect data from Oracle Solaris systems, the HBAs must be configured to use Solaris systems. This note applies to Storage Resource agents.

Run the following command on the Solaris system:

```
cfgadm -al
```

If the system displays this message:

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

run this command:

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

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

You must reboot the Solaris system.

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

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

Creating a SAN plan

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

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

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

Provision storage only, including optional multipath and zone considerations

Click to create a plan for provisioning new storage. The plan does not include provisioning for storage replication.

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

Click to create a plan for provisioning new storage and storage replication.

Provision replication to existing storage

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

Configure multipathing or zoning only

Click to create a plan for provisioning multipathing and zoning for existing storage.

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

Provision storage only, including optional multipath and zone considerations

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

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

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

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

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

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

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

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

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

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

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

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

Provision replication to existing storage

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

Configure multipathing or zoning only

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

Note:

- The Linux SDD driver is no longer available. Starting with Red Hat Enterprise Linux Version 5 and SUSE Linux Enterprise Server Version 10, only the DM_Multipath is available.
- The SDD driver is no longer supported on HP, starting with HP-UX 11i v3 with Itanium. HP-UX 11i v3 provides a native multipathing feature which is supported by SAN Planner.

SAN Planner enables system administrators and storage area network (SAN) administrators to plan and implement storage provisioning for hosts and storage systems with multipath support in fabrics managed by Tivoli Storage Productivity Center. Planning the provisioning of storage to hosts with multipath drivers

requires knowing which storage systems are supported by the host multipath driver and the multipath modes that are supported by the driver and storage subsystem. Planning the paths between the hosts and storage controller requires designing paths between hosts and storage subsystems which are implemented through zones in the fabric.

SAN Planner is used for specifying multiple paths options between selected hosts and storage systems. There should be at least two paths available to a storage system and at least four paths available to SAN Volume Controller. Although not required, it is beneficial to have at least two host ports.

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

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

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

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

The guidance policies used for zone planning are as follows:

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

The validation policies used for zone planning are as follows:

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

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

Planner and user inputs interdependencies

The planner options depend on the inputs that you select. For example, selecting a host that does not have an agent installed limits the planning that can be done. The planning process reveals inconsistencies in your preferences. For example, it might not be possible to have your specified number of paths and have a different

zone for each host. In such cases, the planner makes a best effort to satisfy your inputs. If the planner cannot satisfy conditions, it issues warnings. Errors discovered before plan execution display in the planner. You can adjust the inputs to correct the problem. Errors discovered during plan execution are in the job execution log.

If the selected Storage Resource group has a provisioning profile associated with it, the profile is selected in the Effective Provisioning Profile field. You can select a different profile from the list of available profiles.

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 modify 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. Follow migration recommendations with care to ensure continuity of all replication relationships.

To use the Storage Optimizer, you must have an IBM Tivoli Storage Productivity Center Advanced license.

The following IBM storage subsystems or applications are supported:

- DS8000
- DS6000
- DS5000
- DS4000
- TotalStorage Enterprise Storage Server
- SAN Volume Controller
- Storwize V7000
- IBM Storwize V7000 Unified block level support
- The XIV system
- IBM SONAS

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

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

You must provide the Storage Optimizer with the following input:

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

The Storage Optimizer produces the following output:

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

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

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

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

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

Planning for the Data Path Explorer

The Data Path Explorer is a type of view in the topology viewer. Data path explorer combines several of the typical topology views into a visualization that allows users to follow connectivity from one end point in the topology view to another. Storage administrators can use this view to debug connectivity and performance problems in the storage environment.

Data Path Explorer displays the different data paths (SAN access paths or I/O paths) for each host disk. This action allows you to use the Data Path Explorer to check the health and performance statistics of the storage components along the I/O path.

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

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

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

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

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

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

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

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

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

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

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

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

1. Install and configure the SMI-S agents (CIM agents) for the subsystems to collect subsystem data. CIM agents are not required for DS8000, the XIV system, SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 to collect subsystem data. Install and configure the CIM agents for the fabric to

- collect fabric data. Contact your vendor for information about installing and configuring the CIM agent for the subsystem or fabric.
2. Install and configure the out-of-band Fabric agent to monitor a fabric.
 3. Install the Storage Resource agent on the hosts.
 4. Run a discovery job for the CIM agents, Storage Resource agents, and storage systems.
 5. Run a probe job for the subsystem, fabric, and Storage Resource agents.
 6. Run a performance monitoring job for the subsystem and switches if you want to see performance information.
 7. You can now see the data paths from a host to a storage subsystem in the Data Path Explorer view.

Planning for the Data Manager

This section provides planning information for the Data Manager.

Data Manager is a comprehensive file and capacity management solution for heterogeneous storage environments. Data Manager includes enterprise-wide reporting and monitoring, policy-based management and automated capacity provisioning for Direct Attached Storage (DAS), Network Attached Storage (NAS), and SAN environments.

Data Manager helps you improve storage utilization, plan for future capacity, and ensure availability by providing storage on demand for file systems. Use Data Manager to perform the following functions:

- Discover and monitor disks, partitions, shared directories, and servers.
- Monitor and report on capacity and utilization across platforms to help you identify trends and prevent problems.
- Provides a wide variety of standardized reports about file systems and storage infrastructure to track usage and availability.
- Provide file analysis across platforms to help you identify and reclaim space used by non-essential files.
- Provide policy-based management and automated capacity provisioning for file systems when user-defined thresholds are reached.

Using these functions, Data Manager helps you reduce storage costs by:

- Improving storage utilization.
- Enabling intelligent capacity planning.
- Helping you manage more storage with the same staff.
- Supporting application availability through computer uptime reporting and application database monitoring.

Information collected by Data Manager helps you understand what is really going on with data on your servers and in your storage environment. You can view when files are created, accessed, and modified, and by what group or user. This type of information helps system administrators map storage resources to the consumers of the resource.

In addition to understanding the current consumption and usage of data within the enterprise, Data Manager tracks the information over time. Not only does this historical view of storage consumption and utilization show usage trends over time, the system administrator can also see a projected use of storage in the future.

System administrators can prepare for the need to purchase additional capacity in a planned proactive manner rather than reacting to out-of-space emergencies.

Use the Data Manager policy functions to help you evaluate and control the usage and status of your enterprise storage management assets. Because Data Manager is policy-based, it has autonomic self-healing capabilities that can detect potential problems and automatically adjust parameters based on the policies and actions you have established. Use the capability to provision storage based upon storage policies to expand a file system, and to allocate storage to a volume (LUN).

Planning for the Storage Resource agents on Oracle Solaris

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

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

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

TPC_Data_Server_install_directory/config

The default directory for Windows is:

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

The default directory for UNIX or Linux is:

/usr_or_opt/IBM/TPC/data/config

3. Modify the TPCD.config file with the **skipAutoFS=1** parameter added:

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

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

Save the file.

4. Restart the Data server.

Planning for the Disk Manager

To effectively plan your Disk Manager, you must understand the terminology.

Tivoli Storage Productivity Center uses abstract terminology such as pools, volumes, and so on to help you view and manage your heterogeneous storage systems. Most of these terms are derived from SMI-S, which already provides a common model for storage subsystems.

storage pool

A collection of storage capacity that provides the capacity requirements for

a volume. A pool has certain storage capabilities, which indicate the range of quality of service requirements that can be applied to objects created from the pool.

primordial pool

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

storage volumes

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

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

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

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

Note:

- For DS6000 and DS8000: The DS6000 and DS8000 tracks the number of spares still available after a spare disk has been used to replace a failing disk drive in the device by marking the failing drive with an operational status of "Predictive Failure." If there are not enough spare disks remaining for a rank, then the operational status of the disk also shows a status of "Error."
If the DS6000 or DS8000 marks a disk drive with a status of "Predictive Failure" but not "Error", then Tivoli Storage Productivity Center shows a green icon with a consolidated status of "OK" in the health overlay, and does not create an alert. The topology viewer does show the real status of the disk drive as "Predictive Error" in the Operational Status column of the tabular view. The topology viewer lets you see the changed status.
- For EMC storage systems: The data that Tivoli Storage Productivity Center collects about EMC storage systems is based on SMI-S standards. Within SMI-S model different types of StoragePools are defined and several subsystem attributes can be modeled as special pool types. Beside the pools that are available for volume creation, additional storage space from the EMC storage system might be represented as a pool within the SMI-S model. Because of this, Tivoli Storage Productivity Center reports and CLI commands might display more pools for EMC storage systems than what is reported through the EMC Navisphere agent.

Planning for the IBM CIM agents

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

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

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

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

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

Note:

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

using that Engenio or LSI SMI-S provider. This controlled environment is required to ensure that the Engenio provider receives synchronized CIM Client requests.

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

Before installing Tivoli Storage Productivity Center check the Tivoli Storage Productivity Center support site for the latest CIM agent support. Go to the Tivoli Storage Productivity Center support website at <http://www.ibm.com/support/docview.wss?rs=40&uid=swg21386446>.

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

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

Planning for the TagmaStore CIM agents

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

IBM Tivoli Storage Productivity Center supports the Hitachi Data Systems TagmaStore CIM Agent 5.8. This version of the CIM Agent supports only the Array Profile and not the Storage Virtualizer Profile. However, Tivoli Storage Productivity Center supports the TagmaStore as a Storage Virtualizer. Tivoli Storage Productivity Center displays information for virtual disks and local disks.

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

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

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

For LUN and volume correlation, the host machine must have the Storage Resource agent installed and the TagmaStore device must be in the same SAN fabric. You must also have a zone configured in the active zone set between the

ports of the host machine and the ports of the TagmaStore device. The Fabric agent or Storage Resource agent needs to be configured for the fabric to which the host is connected.

For back-end correlation, the TagmaStore device ports and back-end subsystem ports must be in the same zone and the back-end subsystem has assigned storage volumes to all ports of the TagmaStore device.

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

The events returned by the CIM agent are:

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

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

General procedure

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

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

Planning for Disk Manager configuration

This section provides information for planning Disk Manager configuration.

Disk Manager data can only be collected from devices managed by a CIM agent or through the native storage system interface.

To get Disk Manager data, complete the following steps:

1. Prepare for collecting data from Disk Manager.

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

- The version of CIM agent and firmware for the device is supported. For information about the CIM agent and firmware supported, see <http://www.ibm.com/support/docview.wss?uid=swg21386446>. Select the appropriate version for **Storage**.
- A CIM agent is installed on a different server from IBM Tivoli Storage Productivity Center. For some vendors, the device comes with an embedded CIM agent, and therefore, no separate CIM agent installation is required.

- The CIM agent is configured to manage the intended device.

For information about how to install and configure the CIM agent, contact your CIM agent provider. Also see *IBM Tivoli Storage Productivity Center Hints and Tips* for information about installing and configuring CIM agents. See <http://www.ibm.com/support/docview.wss?uid=swg27008254>.

- For devices that use CIM agents, follow these steps:
 - Add the CIM agent connection to Tivoli Storage Productivity Center. For information about adding a CIM agent, use the Configure Devices wizard. For information about the Configure Devices wizard, see “Configure Devices wizard” on page 240.
 - Run a CIMOM discovery job for the storage system.
 - Run a probe job for the storage system. The probe job gets all component level information into the database. The probe job also includes getting all the instances of controllers, arrays, extents, volumes, LUN to host mappings, FC ports, and so on.
- For devices that use the native interface, add the device to Tivoli Storage Productivity Center. Use the Configure Devices wizard to add the device. For information about the Configure Devices wizard, see “Configure Devices wizard” on page 240. On the Select Device Type panel, select **Storage Subsystem**. After you add the storage system to Tivoli Storage Productivity Center, an automatic discovery and probe jobs are run.
- Once the device is added to Tivoli Storage Productivity Center, automatic discovery and probe operations are done for the storage system.
- Create alerts. You can specify which alerts you want to be notified of when monitoring the storage systems.
 For the storage systems that use the native interfaces, Tivoli Storage Productivity Center gets events directly from the storage system. For the DS8000, Tivoli Storage Productivity Center receives events from the DS8000 enterprise storage server network interface (ESSNI). For the SAN Volume Controller, Storwize V7000, and the XIV system, Tivoli Storage Productivity Center polls the device every minute for events.
 For storage systems that use the CIM agent, Tivoli Storage Productivity Center receives CIM indications.
- You can configure Tivoli Storage Productivity Center to send SNMP traps. if you configure Tivoli Storage Productivity Center to send SNMP traps or Omnibus events, you must first import the Tivoli Storage Productivity Center MIB and Baroc files into your SNMP manager or Omnibus manager. For more information about configuring SNMP traps, see “Configuring Tivoli Storage Productivity Center to send SNMP traps and IBM Tivoli Enterprise Console events” on page 273.
- Run a probe job for storage systems other than DS8000, the XIV system, SAN Volume Controller, or Storwize V7000. Probe jobs collect detailed information about the storage systems known to Tivoli Storage Productivity Center. In contrast to a CIMOM discovery job, the probe interacts with the CIM agent to get information about just the storage devices configured to be probed. You can set up a probe job to gather data for just one storage system, or you can set up probe jobs to collect data from more than one storage system. A probe job must be run before volumes can be created on a storage system or a storage system can be monitored for performance management.

Running a probe job puts additional workload on the Tivoli Storage Productivity Center server, the repository database, and the CIM agents. For probe jobs, the additional workload is especially important if multiple devices

are managed by the same CIM agent and are probed concurrently. As a result, the probe completion time is impacted by the number of devices probed in parallel and the number of CIM agents involved in the probe job. This fact is especially important if the CIM agent machines are lower-end machines of the hardware prerequisites.

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

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

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

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

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

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

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

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

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

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

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

For information about how to configure Tivoli Storage Productivity Center for performance monitoring, see “Performance management configuration” on page 104.

Planning for TotalStorage Enterprise Storage Server FlashCopy reports

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

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

1. Add the following types to the CIM agent:

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

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

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

Storage capacity of volumes

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

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

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

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

Planning for the DS8000

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

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

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

Tivoli Storage Productivity Center support

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

General procedure

To view the reports and topology for the DS8000, complete the following steps:

1. Open the Tivoli Storage Productivity Center user interface and add the DS8000 to Tivoli Storage Productivity Center. Click **Administrative Services > Data Sources > Storage Subsystems**. Click **Add**. This action opens the Configure Devices wizard. For information about adding the storage system, see "Configure storage subsystem connections page" on page 244.
2. After you add the DS8000, an automatic discovery and probe are performed.
3. View the reports and topology for the DS8000 in Tivoli Storage Productivity Center.

Planning for DS8000 performance

This section provides information about DS8000 performance.

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

Planning for the SAN Volume Controller

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

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

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

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

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

When you have hosts connected to storage subsystems which have multi-pathing enabled, you must have the multi-pathing subsystem device drivers (SDD)

installed on the hosts. For more information, see the SAN Volume Controller Information Center and search for **multipath support**.

SAN Volume Controller Version 6.1 or later

SAN Volume Controller includes IBM System Storage Easy Tier, a function that responds to the presence of solid-state drives (SSDs) in a storage pool that also contains hard disk drives (HDDs). The system automatically and nondisruptively moves frequently accessed data from HDD MDisks to SSD MDisks, thus placing such data in a faster tier of storage.

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

SAN Volume Controller supports these tiers:

Generic SSD tier

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

Generic HDD tier

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

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

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

Note:

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

However, the terms have not been changed in Tivoli Storage Productivity Center reports and views.

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

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

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

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

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

Upgrading SAN Volume Controller Version 5.1

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

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

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

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

General procedure

To view the reports and topology for the SAN Volume Controller environment, complete the following steps:

1. Log in to the SAN Volume Controller console and create a user ID and password for use with Tivoli Storage Productivity Center.
2. Open the Tivoli Storage Productivity Center user interface and add the SAN Volume Controller to Tivoli Storage Productivity Center. Click **Administrative Services > Data Sources > Storage Subsystems**. Click **Add**. This opens the Configure Devices wizard. For information about adding the storage system, see “Configure storage subsystem connections page” on page 244.
3. Tivoli Storage Productivity Center supports both OpenSSH format or putty key (.ppk) format that is not password protected. You only need to provide the private SSH key.

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

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

This action saves the key in OpenSSH format.

Important: If you upload a new key for the existing user ID using Tivoli Storage Productivity Center, that new key overrides the existing key. If you have another application using that SAN Volume Controller user ID (for

example Tivoli Storage Productivity Center for Replication), the application cannot log into SAN Volume Controller because the key associated with the user ID is overwritten.

If you have this situation, you should create a user ID specifically for Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication on the SVC cluster. Add the SAN Volume Controller to Tivoli Storage Productivity Center for Replication using the new user ID and password (or change the user ID and password while in the disconnected state).

Note:

- IBM Tivoli Storage Productivity Center ships default public and private keys in OpenSSH format. The public and private keys can be found in `tpc_installation_directory/device/conf`. The file names are `tpc_svc.pub` (public key) and `tpc_svc.pem` (private key).
4. After you add the SAN Volume Controller, an automatic discovery and probe are performed.
 5. Before you can view performance reports for the SAN Volume Controller, you must define and run a performance monitor job.
 6. View the reports and topology for the SAN Volume Controller in Tivoli Storage Productivity Center.

Planning for Storwize V7000

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

General concepts

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

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

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

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

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

focal point for configuration and monitoring requests. This node is called the configuration node and is the only node that activates the management IP addresses.

Each Storwize V7000 can have zero to four management IP addresses. You can assign up to two IPv4 addresses and up to two IPv6 addresses.

All configuration, monitoring, and service tasks are performed at the cluster level (a cluster is a Storwize V7000 system that consists of two nodes).

Storwize V7000 includes IBM System Storage Easy Tier, a function that supports solid-state drives (SSDs) in a storage pool that also contains hard disk drives (HDDs). The system automatically and nondisruptively moves frequently accessed data from HDD MDisks to SSD MDisks, thus placing such data in a faster tier of storage.

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

Using Tivoli Storage Productivity Center to monitor Storwize V7000

Tivoli Storage Productivity Center communicates with Storwize V7000 storage systems through a native interface. No CIM agents are required.

Use Tivoli Storage Productivity Center to monitor Storwize V7000 systems, including discovery, data collection (probes), performance monitoring, trending, storage optimization, and SAN planning. Tivoli Storage Productivity Center reports and topology views also provide information about the tier, tier capacity, and tier free capacity for Storwize V7000.

Note: For this release of Tivoli Storage Productivity Center, the performance monitor, SAN Planner, and Storage Optimizer do not support Easy Tier.

Similar to other systems, Storwize V7000 appears as an entity within data sources, reports, data collection schedules, the topology viewer, and so on. Tivoli Storage Productivity Center also supports launch in context and single sign-on for Storwize V7000.

General procedure

To view the reports and topology for the Storwize V7000 environment, complete the following steps:

1. Log in to the Storwize V7000 console and create a user ID and password for use with Tivoli Storage Productivity Center.
2. Open the Tivoli Storage Productivity Center user interface and add the Storwize V7000 to Tivoli Storage Productivity Center. Click **Administrative Services > Data Sources > Storage Subsystems**. Click **Add**. This opens the Configure Devices wizard. For information about adding the storage system, see “Configure storage subsystem connections page” on page 244.

3. Tivoli Storage Productivity Center supports both OpenSSH format or putty key (.ppk) format that is not password protected. You only need to provide the private SSH key.

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

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

This action saves the key in OpenSSH format.

Important: If you upload a new key for the existing user ID using Tivoli Storage Productivity Center, that new key overrides the existing key. If you have another application using that Storwize V7000 user ID (for example Tivoli Storage Productivity Center for Replication), the application cannot log into Storwize V7000 because the key associated with the user ID is overwritten.

If you have this situation, you should create a user ID specifically for Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication on the Storwize V7000 cluster. Add the Storwize V7000 to Tivoli Storage Productivity Center for Replication using the new user ID and password (or change the user ID and password while in the disconnected state).

Note:

- IBM Tivoli Storage Productivity Center ships default public and private keys in OpenSSH format. The public and private keys can be found in `tpc_installation_directory/device/conf`. The file names are `tpc_svc.pub` (public key) and `tpc_svc.pem` (private key).
4. After you add the Storwize V7000, an automatic discovery and probe are performed.
 5. Before you can view performance reports for the Storwize V7000, you must define and run a performance monitor.
 6. View the reports and topology for the Storwize V7000 in Tivoli Storage Productivity Center.

Planning for Storwize V7000 Unified

The IBM Storwize V7000 Unified system is a virtualizing redundant array of independent disks (RAID) storage system that supports both block protocols and file protocols. This unified system includes Storwize V7000 File Module and the Storwize V7000 storage system.

General concepts

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

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

Storwize V7000 File Module is a clustered system that consists of two units that provide file systems for use by network-attached storage. The file module uses the Storwize V7000 storage system to provide the file module with volumes. Other volumes, which are block volumes, are provided on the SAN to be presented to hosts.

The Storwize V7000 File Module software within the Storwize V7000 Unified system contains the interface node and management node functions. The management node function provides for configuring, administering, and monitoring a system. The interface node function connects a system to an Internet Protocol (IP) network. You can connect by using any of the following protocols:

- Common Internet File System (CIFS)
- Network File System (NFS)
- File Transfer Protocol (FTP)
- Hypertext Transfer Protocol Secure (HTTPS)
- Secure Copy Protocol (SCP)

The Storwize V7000 Unified software also performs the following functions for the host systems that attach to the Storwize V7000 Unified system by using block protocols:

- Creates a single pool of storage
- Provides logical unit virtualization
- Manages logical volumes
- Mirrors logical volumes
- Large scalable cache
- Copy Services
 - IBM FlashCopy (point-in-time copy) function, including thin-provisioned FlashCopy to make multiple targets affordable
 - Metro Mirror (synchronous copy)
 - Global Mirror (asynchronous copy)
 - Data migration
- Space management
 - IBM System Storage® Easy Tier™ to migrate the most frequently used data to higher performing storage
 - Metering of service quality when combined with IBM Tivoli Storage Productivity Center
 - Thin-provisioned logical volumes

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

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

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

focal point for configuration and monitoring requests. This node is called the configuration node and is the only node that activates the management IP addresses.

Each Storwize V7000 can have zero to four management IP addresses. You can assign up to two IPv4 addresses and up to two IPv6 addresses.

All configuration, monitoring, and service tasks are performed at the cluster level (a cluster is a Storwize V7000 system that consists of two nodes).

Storwize V7000 includes IBM System Storage Easy Tier, a function that supports solid-state drives (SSDs) in a storage pool that also contains hard disk drives (HDDs). The system automatically and nondisruptively moves frequently accessed data from HDD MDisks to SSD MDisks, thus placing such data in a faster tier of storage.

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

Using Tivoli Storage Productivity Center to monitor Storwize V7000 Unified

Tivoli Storage Productivity Center communicates with Storwize V7000 Unified storage systems through a native interface. No CIM agents are required.

Use Tivoli Storage Productivity Center to monitor Storwize V7000 Unified systems, including discovery, data collection (probes), performance monitoring, and trending. Tivoli Storage Productivity Center reports and topology views also provide information about the tier, tier capacity, and tier free capacity for Storwize V7000 Unified.

Similar to other systems, Storwize V7000 Unified appears as an entity within data sources, reports, data collection schedules, the topology viewer, and so on. Tivoli Storage Productivity Center also supports launch in context and single sign-on for Storwize V7000 Unified.

General procedure

To view the reports and topology for the Storwize V7000 Unified environment, complete the following steps:

1. Log on to the Storwize V7000 Unified console and create a user ID and password for use with Tivoli Storage Productivity Center.
2. Open the Tivoli Storage Productivity Center user interface and add the Storwize V7000 Unified to Tivoli Storage Productivity Center. Click **Administrative Services > Data Sources > Storage Subsystems**. Click **Add**. This opens the Configure Devices wizard. For information about adding the storage system, see “Configure storage subsystem connections page” on page 244.
3. Tivoli Storage Productivity Center supports Secure Shell (SSH) keys that are in OpenSSH format or PuTTY (.ppk) format that is not password protected.

If you have a password protected PuTTY .ppk key, use the PuTTYgen utility on a Windows operating system to convert that key to OpenSSH format by following these steps:

- a. Start the PuTTYgen utility.
- b. Load the .ppk file. Click **File > Load private key**.
- c. From the menu bar, click **Conversion > Export OpenSSH key conversion**.

This action saves the key in OpenSSH format.

Important: If you upload a new key for the existing user ID using Tivoli Storage Productivity Center, that new key overrides the existing key. If you have another application using that Storwize V7000 Unified user ID (for example Tivoli Storage Productivity Center for Replication), the application cannot log on to Storwize V7000 Unified because the key associated with the user ID is overwritten.

If you have this situation, create a user ID specifically for Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication on the Storwize V7000 Unified cluster. Add the Storwize V7000 Unified to Tivoli Storage Productivity Center for Replication using the new user ID and password (or change the user ID and password while in the disconnected state).

Note:

- Tivoli Storage Productivity Center ships default public and private keys in OpenSSH format. The public and private keys can be found in `tpc_installation_directory/device/conf`. The file names are `tpc_svc.pub` (public key) and `tpc_svc.pem` (private key).
4. After you add the Storwize V7000 Unified, an automatic discovery and probe are performed.
 5. Before you can view performance reports for the Storwize V7000 Unified, you must define and run a performance monitor.
 6. View the reports and topology for the Storwize V7000 Unified in Tivoli Storage Productivity Center.

To scan a Storwize V7000 Unified file system, complete the following steps:

1. Export the Storwize V7000 Unified file system to be scanned as a CIFS or NFS export.
2. Mount the export on a host which has a Storage Resource agent deployed on it. Exports are accessible to the TCP/IP data network through the Storwize V7000 Unified interface node **public IPs**. Exports must be mounted as a CIFS export on Windows agents or an NFS export on UNIX agents. On Windows operating systems, the export must be mounted as the local system user or from a service that is running under the local system account so that it is visible to all logon sessions. One way of mapping a drive as the local system user is to map the drive from a **psexec** command window. The **psexec** command window is available as part of the **Sysinternals** command-line tools for Windows operating systems.
3. Probe the Storwize V7000 Unified and the Storage Resource agent. The default behavior of the probe is to discard export paths that are not at the root of the file system. This is done so that a situation cannot occur where the product counts information twice on a file system. The reason for this behavior is because the probe is treating two exports to the same file system as separate file systems. This behavior can be overridden by setting the **saveNonRoot** parameter. The probe picks up all exports (shares) of a specific NAS device as

file systems. To set the **saveNonRoot** parameter, insert the line **saveNonRoot=1** in the server section of the `TPCD.config` file and restart the Data server to take effect. The `TPCD.config` file is located in the `TPC_installation_directory/data/config` directory where `TPC_installation_directory` is where Tivoli Storage Productivity Center is installed.

4. Assign a scan agent to Storwize V7000 Unified file system. Open the Tivoli Storage Productivity Center GUI and click **Administrative Services > Configuration > Scan/Probe Agent Administration**. Save the page.
5. A scan job for the Storwize V7000 Unified can be created and run from **Data Manager > Monitoring > Scans**.

Identifying Storwize V7000 Unified devices in the user interface

The way you configure a Storwize V7000 Unified device determines how the Storwize V7000 Unified device is represented in the Tivoli Storage Productivity Center user interface, as follows:

- Storwize V7000 File Module generally appears as a "computer" in the user interface. In **Data Manager > Alerting > Other NAS Alerts > NAS**, it appears as a "NAS cluster."
- Storwize V7000 storage system appears as a "subsystem" in the user interface.

Planning for the XIV system

Use this information to plan for using the XIV Storage System with Tivoli Storage Productivity Center.

IBM XIV Storage System is a disk storage architecture designed to eliminate the complexity of administration and management of storage. The XIV system parallelized architecture of the system, optimal exploitation of all system components (including disks, CPUs, and switches), and unique caching architecture all translate into excellent performance.

The unique balancing of all data across system components prevents the occurrence of hot spots. With all components working under the same load, performance and reliability are exceptional.

The XIV system uses large capacity Serial Advanced Technology Attachment (SATA) disk drives which optimizes the use of disk capacity, resulting in outstanding power consumption without compromising performance.

The XIV system is designed to be scalable in storage, interfaces, cache, CPU power, and internal bandwidth. The architecture supports each aspect to grow independently, resulting in a scalable system in both capacity and performance.

The XIV system provides data protection and availability. All disk drives, modules, switches, and uninterruptible power supply units are fully redundant, ensuring high reliability and excellent performance.

The built-in thin provisioning of the XIV system helps reduce direct and indirect costs by allowing users to install capacity only for data written. You can grow your data capacity over time with minimal management effort.

For more information about XIV system, see the XIV system Information Center.

Tivoli Storage Productivity Center support

Tivoli Storage Productivity Center uses the native interface for the XIV system to run discovery and probe jobs. Because of the dynamic XIV architecture, where volume data is spread over all disks, no data that describes the space allocation from volumes to a specific physical disk is available.

The XIV system includes a graphical user interface (GUI) for configuration and administration (XIV Storage Manager). XIV Storage Manager is the element manager for the XIV system. XIV Storage Manager version 4.2 or later must be installed on the same computer as the Tivoli Storage Productivity Center GUI. If XIV Storage Manager is not installed on the same computer as the Tivoli Storage Productivity Center GUI or is configured incorrectly, an error message is displayed when you attempt to start XIV Storage Manager from Tivoli Storage Productivity Center.

The XIV systems have the capability of defining up to three administrative nodes, each with their own IP address. In Tivoli Storage Productivity Center 4.2.1 or later, if an XIV system is configured with multiple administrative nodes, Tivoli Storage Productivity Center detects the IP addresses for these nodes. If Tivoli Storage Productivity Center fails to connect to one of the IP addresses, then an attempt is made to connect to the XIV system using one of the other IP addresses.

Note: The XIV system does not support IPv6.

Considerations for data collected for XIV system

The following considerations apply to values shown in the XIV system and Tivoli Storage Productivity Center GUIs.

- The XIV Storage System GUI and the XIV Storage System command-line interface (CLI) report sizes in GB of disks, volumes, pools, and so on, based on 10^9 . Tivoli Storage Productivity Center reports sizes in GB based on 2^{30} . For example, a volume showing a size of 17 GB in the XIV Storage System GUI is reported as 16 GB (XIV GB size * 1000 * 1000 * 1000 / 1024 / 1024 / 1024 rounded to the nearest GB) in Tivoli Storage Productivity Center.
- The XIV system differentiates thin provisioned pools from non-thin provisioned (regular) pools and volumes in a thin provisioned pool and a non-thin provisioned pool. When hard size is equal to soft size, the pool is represented as a regular pool, otherwise the pool is represented as thin provisioned. However, the XIV system CLI and CIM do not distinguish between regular and thin provisioned pools and all pool volume names in Tivoli Storage Productivity Center are prefixed with an asterisk (*) to indicate that the volume is thin provisioned.

General procedure

To view the reports and topology for the XIV system environment, complete the following steps:

1. Open the Tivoli Storage Productivity Center user interface and add the XIV system to Tivoli Storage Productivity Center. In the navigation tree, expand **IBM Tivoli Storage Productivity Center > Configure Devices**. This opens the Configure Devices wizard. For information about adding the storage system, see “Configure storage subsystem connections page” on page 244.
2. After you add the XIV system, an automatic discovery and probe are performed.

3. View the reports and topology for the XIV system in Tivoli Storage Productivity Center.

Planning for the Fabric Manager

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

The main features of Fabric Manager are shown in the following table.

Table 21. Main features of the Fabric Manager

Feature	Advantages	Benefit
Automatic device discovery	Enables you to see the data path from the servers to the switches and storage systems	Easily manage your SANs from a single console
Real-time monitoring and alerts	Monitors SAN events and alerts administrators of problems	Helps maintain the availability of your SAN
Zone control	Gives you the ability to add and remove devices from zones	Comprehensive fabric management from a single console
SAN reporting	Provides the ability to generate reports on the SAN devices	Quickly provides reports and inventories of your SAN
Enterprise scalability	Scale from SAN islands to enterprise SANs	Accommodates new and ever changing business needs
Switch performance management (1)	Provides the ability to monitor switches for performance metrics which allows for the triggering of switch threshold alerts and performance reporting. The performance reporting is at the switch and port levels and displays in the defined Switch Performance reports and in the topology viewer.	Helps you maintain high SAN availability

Note:

1. Tivoli Storage Productivity Center provides predictive fault analytics for the SAN infrastructure through reporting, monitoring, and alerting of link failure rates and error frame rates. These transient errors are indicative of potential pending SAN link failures.
2. Tivoli Storage Productivity Center does not support heterogeneous fabrics in versions 5.1 and later.

Planning for Brocade management agents

IBM Tivoli Storage Productivity Center uses a Storage Management Initiative (SMI) agent to manage Brocade switches.

Supported Brocade models

If you are using Brocade Data Center Fabric Manager (DCFM) 10.4 or later, the Storage Management Initiative (SMI) agent used by Tivoli Storage Productivity Center is embedded in the DCFM. If you are using Brocade Network Advisor (BNA) 11.1 or later, the SMI agent is embedded in the BNA. If you are not using Brocade DCFM 10.4 or later, and not using BNA 11.1 or later, you must use external SMI agents for Brocade switches.

For the Brocade models, firmware versions, and SMI agents that are used by Tivoli Storage Productivity Center, see the support matrix at <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

Discovering fabrics and switches that use an SMI agent

When you configure Tivoli Storage Productivity Center to connect to the SMI agent (embedded or external), you can run a CIM object manager (CIMOM) discovery job. The discovery job discovers the fabric and switches and the status of these fabrics and switches.

When you run a probe job using the SMI agent, you can see the topology of the fabric, the properties of the hardware that make up the fabric, zoning information, and the endpoint devices.

When you run a fabric performance monitor job, you can see the performance data for the fabric.

Using other agents for discovery: out-of-band SNMP agents or Storage Resource agents

SMI agents provide the following benefits for fabric management:

- Probes and zone control do not occur in the data path.
- Zone alias control is supported by Tivoli Storage Productivity Center only through the fabric SMI agents.
- Propagation of alerts for real-time SAN events is not degraded by using fabric SMI agents.

Using multiple agent types allows the redundant collection of information if one type of agent fails. Also, some types of agents provide certain features the other agents do not. However, using multiple agent types can increase network traffic and the management load on the switches.

The following table outlines the fabric management functions that are supported by agent type. Where multiple agents can be used, the preferred agent type is noted.

Table 22. Supported and preferred interfaces for fabric management

Function	Brocade
Performance monitoring	Fabric SMI agent

Table 22. Supported and preferred interfaces for fabric management (continued)

Function	Brocade
Topology information collected	Preferred: Fabric SMI agent Still supported: Storage Resource agent Out-of-band SNMP agent
Zone information collected	Fabric SMI agent
Zone control	Fabric SMI agent
Zone control with zone aliases	Fabric SMI agent
Switch and switch port information collected (for running performance monitor)	Preferred: Fabric SMI agent
Fabric and switch events (see table note)	Preferred: Fabric SMI agent Switch SNMP trap configuration
Launch of switch and fabric element managers	Preferred: Fabric SMI agent Still supported: Storage Resource agent Out-of-band SNMP agent
Hosts, endpoint devices, device-centric and host-centric information collected	Storage Resource agent
Switch sensors and events	Out-of-band SNMP agent
<p>Note:</p> <ul style="list-style-type: none"> The Storage Resource agent runs the following jobs: <ul style="list-style-type: none"> Discovery When a discovery job is run, Tivoli Storage Productivity Center gathers information about the fabrics, fabric to switch relationships, and other key attributes about the switches in the fabrics. Probe When a probe job is run, Tivoli Storage Productivity Center gathers additional details for topology and zoning information. <p>If you deploy a Storage Resource agent, and a fabric probe job has not been run, the connections from the Storage Resource agents to the switches are not displayed, and all the fibre channel (FC) ports are reported as unused.</p> Collected zone information indicates that the active zone set and the inactive zoning library information is collected during a fabric probe. Zone control refers to the ability to actively configure zoning. Tivoli Storage Productivity Center does not get a transaction lock on the switch while making zone changes on the fabric. For fabric and switch events, the Storage Resource agent relies on both SNMP traps and CIM indications for event notification and automatic fabric discovery and probe, if needed. <p>The preferred method includes both SMI agent and switch SNMP trap configuration. An out-of-band SNMP agent is not necessary for getting SNMP traps. You can configure the switch to send SNMP traps to the Device server. However, you can configure the out-of-band SNMP agent for other reasons.</p>	

Planning to use the embedded SMI agent

A software-embedded Storage Management Initiative (SMI) agent is provided in Brocade Data Center Fabric Manager (DCFM) 10.4 or later, or Brocade Network Advisor 11.1 or later. DCFM and Brocade Network Advisor manage multiple fabrics within and across data centers. When you configure DCFM or Brocade Network Advisor, you set up one switch to be the *seed* or master switch that interconnects to all the other switches in the fabric.

Configuring the Brocade switch

Before Tivoli Storage Productivity Center can discover the switches in a fabric through the software-embedded SMI agent, you must configure the Brocade DCFM or Brocade Network Advisor to manage the fabric. Use a seed switch to configure the Brocade DCFM or Brocade network advisor. This configuration is done through the seed switch. For heterogeneous Fabric OS (FOS) fabrics, the seed switch must be one of the FOS switches. There are special requirements for the seed switch:

- For pure FOS fabrics, the seed switch must have a firmware of 5.3 or later.
- For heterogeneous FOS fabrics, the seed switch must have a firmware of 6.0 or later.

Migrating from a legacy agent to an SMI agent

You can configure Tivoli Storage Productivity Center to use the software-embedded SMI agent instead of the out-of-band Fabric agent for the FOS fabric.

Limitations

These are some limitations using the software-embedded SMI agent:

- The software-embedded SMI agent has asset collection period logic for handling fabric events. If anything in the fabric topology changes and is restored within the asset collection period, the SMI agent does not send a CIM indication. In this situation, you must run a fabric probe to collect the most updated fabric topology information. For more information about the asset collection period, including the value for the collection period, refer to the DCFM or Brocade Network Advisor documentation.
- The software-embedded SMI agent is not supported on IBM AIX.
- Zone control sessions through the SMI agent do not establish a session with the seed switch until the session from Tivoli Storage Productivity Center is committed.
- DCFM Professional Edition does not have the software-embedded SMI agent.
- You can use either the HTTPS or HTTP protocol at one time with DCF or Brocade Network Advisor. You cannot use both protocols at the same time.
- Empty zones, zone sets, and zone aliases are not supported through FOS fabrics through the SMI agent.
- When you set up the seed switch, if the switch link is broken anywhere along the connection line, Tivoli Storage Productivity Center will not be able to report on the switches past the broken link.
- When a Brocade i10000 switch has a faulty control processor card, the SMI agent managing the switch can hang indefinitely. This issue causes time-outs and probe failures from the Tivoli Storage Productivity Center server.

Planning for external SMI agent support

Brocade fabrics that are not managed by Brocade Data Center Fabric Manager (DCFM) 10.4.0 or later and are not managed by Brocade Network Advisor 11.1 or later require a separate external Storage Management Initiative (SMI) agent for the Brocade hardware. The external SMI agents must be installed on separate hosts and can manage one or more switches.

Installing and Configuring external SMI agents

You need to install and configure the external SMI agents before adding the SMI agents to Tivoli Storage Productivity Center.

For more information about configuring the external SMI agents for Brocade switches, refer to the documentation for the switch.

Considerations

Consider these items when using external SMI agents:

- If Brocade fabrics are not laid out as a *mesh*, and one inter-switch link (ISL) break can segment the fabric, manage the fabric with both a fabric SMI agent and an out-of-band Fabric agent (SNMP). The reason for this is that Brocade SMI agents do not automatically manage newly created fabrics in the event of a fabric segmentation.
- Brocade SMI agents do not correctly report whether switches and fabrics are running in Interop Mode or Native Mode. Because of this limitation, Tivoli Storage Productivity Center is unable to determine whether a fabric or switch is in Interop Mode or Native Mode. For homogeneous Brocade fabrics that are configured in Interop Mode, the following actions from Tivoli Storage Productivity Center will fail:
 - Zone control with domain/port members and node WWN members
 - Creating zone aliases

Planning for QLogic or Cisco fabric management agents

Depending on the function that you want to perform (topology discovery, zone discovery, performance monitoring, and so forth), IBM Tivoli Storage Productivity Center uses a Storage Management Initiative (SMI) agent, out-of-band Simple Network Management Protocol (SNMP) agent, or Storage Resource agent to manage QLogic and Cisco switches.

Supported QLogic and Cisco models

For the QLogic and Cisco models, firmware versions, and agents that are supported by Tivoli Storage Productivity Center, see the support matrix at <http://www.ibm.com/support/docview.wss?uid=swg27019378>.

The following table outlines the fabric management functions that are supported by agent type.

Table 23. Supported and preferred interfaces for fabric management

Function	QLogic	Cisco
Performance monitoring	Fabric SMI agent	Fabric SMI agent

Table 23. Supported and preferred interfaces for fabric management (continued)

Function	QLogic	Cisco
Topology information collected	Storage Resource agent Out-of-band SNMP agent	Storage Resource agent Out-of-band SNMP agent SNMP is needed for physical infrastructure discovery.
Zone information collected	Storage Resource agent	Storage Resource agent required in each VSAN
Zone control	Storage Resource agent	Storage Resource agent required in each VSAN
Zone control with zone aliases	Not supported	Not supported
Switch and switch port information collected (for running performance monitor)	Preferred: Fabric SMI agent	Preferred: Fabric SMI agent
Fabric and switch events (see table note)	Out-of-band SNMP agent	Out-of-band SNMP agent
Launch of switch and fabric element managers	Storage Resource agent Out-of-band SNMP agent	Storage Resource agent Out-of-band SNMP agent
Hosts, endpoint devices, device-centric, and host-centric information collected	Storage Resource agent	Storage Resource agent
Switch sensors and events	Out-of-band SNMP agent	Out-of-band SNMP agent

Table 23. Supported and preferred interfaces for fabric management (continued)

Function	QLogic	Cisco
<p>Note:</p> <ol style="list-style-type: none"> The Storage Resource agent has 2 distinct functions: <ul style="list-style-type: none"> Discovery When a discovery job is run, Tivoli Storage Productivity Center gathers information about the fabrics, fabric to switch relationships, and other key attributes about the switches in the fabrics. Probe When a probe job is run, Tivoli Storage Productivity Center gathers additional details for topology and zoning information. <p>If you deploy a Storage Resource agent and a fabric probe job is not run, the connections from the Storage Resource agents to the switches are not displayed, and all the Fibre Channel (FC) ports are reported as unused.</p> Zone information collected indicates that the active zone set and the inactive zoning library information is collected during a fabric probe. Zone control is the ability to actively configure zoning. Tivoli Storage Productivity Center does not get a transaction lock on the switch while making zone changes on the fabric. The following notes are specific to Cisco: <ul style="list-style-type: none"> Each virtual storage area network (VSAN) is viewed as an individual storage area network (SAN). The out-of-band SNMP agent and SMI agent gets the physical fabric information and can correlate the VSAN information to a physical infrastructure. The out-of-band SNMP agent also collects some VSAN information. The GS-4 function is not supported by the Storage Resource agent for Cisco. For many Cisco proprietary zone members, the values are displayed as zeros. If a zoning configuration has Cisco proprietary zones, then you cannot modify the zoning for that configuration through Tivoli Storage Productivity Center. You must use the element manager to remove proprietary zone members or to perform zoning. No aliases for Cisco are reported. Orphan Cisco zones and aliases are not reported. For fabric and switch events, the Storage Resource agent relies on both SNMP traps and CIM indications for event notification and automatic fabric discovery and probe, if needed. The preferred method includes both SMI agent and switch SNMP trap configuration. An out-of-band SNMP agent is not necessary for getting SNMP traps. You can configure the switch to send SNMP traps to the Device server. However, you can configure the out-of-band SNMP agent for other reasons. 		

Discovering fabrics and switches

The type of discovery that is performed for fabrics and switches depends on the agent that you are using. For information about discovery by agent type, see the topic about discovering storage resources in the IBM Tivoli Storage Productivity Center User's Guide .

Collecting data with the Fabric Manager

Tivoli Storage Productivity Center uses different types of agents to gather data about switches. Different combinations of these agents are required to effectively

enable the functions of Fabric Manager. In addition, the topology viewer is greatly affected by the discovery of all the managed entities in the management scope of Tivoli Storage Productivity Center.

The following table shows the actions that trigger data collection and the data that is collected depending on the type of agent that is used.

Table 24. Information collected by agent

Action that triggers data collection	Storage Resource agent	Out-of-band fabric agent (SNMP)	SMI agent (Brocade)	SMI agent (Cisco and QLogic)
Administrative Services > Discovery > Out of Band Fabric	–	<ul style="list-style-type: none"> • Discovers switches and connections for switches with IP addresses in a specified range (if you selected this option) • Discovers switches and connections for switches that were previously added as out-of-band agents (either manually or switches with IP addresses in a specified range for which an out-of-band discovery was previously run) • Retrieves Cisco VSAN and physical infrastructure and correlates VSAN to physical infrastructure (see note 2 on page 94) 	–	
Administrative Services > Discovery > CIMOM	–	–	<ul style="list-style-type: none"> • Discovers Storage Management Initiative (SMI) agents through Service Location Protocol (SLP) • Retrieves switch information • Retrieves fabric information 	<ul style="list-style-type: none"> • Discovers SMI agents through SLP • Retrieves switch information • Retrieves fabric information
IBM Tivoli Storage Productivity Center > Probe (including fabric) (see notes 3 on page 94 and 9 on page 94)	<ul style="list-style-type: none"> • Retrieves zoning information, excluding zone alias information, for QLogic and Cisco • Retrieves switch and topology information • Retrieves attribute information for hosts and devices connected to the fabric 	<ul style="list-style-type: none"> • Retrieves switch and topology information • Retrieves Cisco VSAN and physical infrastructure and correlates VSAN to physical infrastructure (see note 2 on page 94) 	<ul style="list-style-type: none"> • 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 	<ul style="list-style-type: none"> • Retrieves detailed fabric, switch, switch port, and blade information

Table 24. Information collected by agent (continued)

Action that triggers data collection	Storage Resource agent	Out-of-band fabric agent (SNMP)	SMI agent (Brocade)	SMI agent (Cisco and QLogic)
<p>Action of Storage Resource agent (see note 1):</p> <ul style="list-style-type: none"> • Turns on • Connectivity changes • Change event detected 	<ul style="list-style-type: none"> • Retrieves zoning information, excluding zone alias information • Retrieves switch and topology information • Retrieves attribute information for hosts and devices connected to the fabric 	<ul style="list-style-type: none"> • Retrieves switch and topology information • Retrieves Cisco VSAN and physical infrastructure and correlates VSAN to physical infrastructure (see note 2 on page 94) 		
<p>SNMP trap received from switch (see note 6 on page 94)</p>	<ul style="list-style-type: none"> • Retrieves zoning information, excluding zone alias information • Retrieves switch and topology information • Retrieves attribute information for hosts and devices connected to the fabric 	<ul style="list-style-type: none"> • Retrieves switch and topology information • Retrieves Cisco VSAN and physical infrastructure and correlates VSAN to physical infrastructure (see note 2 on page 94) 		
<p>CIM indication received from fabric SMI agent (see note 7 on page 94)</p>			<p>Performs a mini-probe to collect information that is relevant to the indication received, or sets entities as missing. Might do one of the following actions:</p> <ul style="list-style-type: none"> • Retrieves switches and fabrics, or • Retrieves zoning data, or • Retrieves switch port status and connection to node, or • Retrieves switch blade status and all port connections for the blade, or • Sets fabric, switch, blades, connections, or nodes as missing 	

Table notes:

1. The Storage Resource agent has two distinct functions:
 - **Discovery**

When a discovery job is run, Tivoli Storage Productivity Center gathers information about the fabrics, fabric to switch relationships, and other key attributes about the switches in the fabrics.

- **Probe**

When a probe job is run, Tivoli Storage Productivity Center gathers additional details for topology and zoning information.

If you deploy a Storage Resource agent and a fabric probe job is not run, the connections from the Storage Resource agents to the switches are not displayed, and all the Fibre Channel (FC) ports are reported as unused

2. SNMP traps received result in a fabric probe of fabrics for the switches (for which traps were received). Agent assignment determines the agents that need to be used for that probe. Event notifications from the event agent could result in any of the agent types to be used (Storage Resource agent, SNMP agent, and SMI agent) for the probe.
3. A fabric probe always results in a fabric discovery using Storage Resource agents and SMIS agents that can report on fabrics been probed.
4. With agent assignment, all Storage Resource agents capable of reporting on a specific fabric might not be run if Tivoli Storage Productivity Center does not deem it necessary to do so.
5. For Storage Resource agents:
 - Starting a Storage Resource agent results in a fabric discovery using that Storage Resource agent. If the fabrics discovered are part of some fabric probe definition, then a fabric probe is also done. Agent assignment is used to determine the actual agents needed for that fabric probe.
 - Stopping a storage resource agent does not cause any data collection
6. SNMP traps result in a fabric probe of fabrics for which the traps were received. Agent assignment determines the agents that need to be used for that probe. So, SNMP traps might result in any of the agent types being used (Storage Resource agent, SNMP agent, and SMIS agent).
7. A CIM indication results in a Storage Resource agent fabric discovery for Storage Resource agents that can report on fabrics for which CIM indications were received. Also, if a mini-probe fails and a Storage Resource agent can report on the fabric for which the indication is received, the Storage Resource agent is run to collect fabric information.
8. Starting or restarting a device or the Data server or both can result in a fabric probe of fabrics which are part of some probe definition defined by the user. Agent assignment determines the agents that must be used for that probe. This action might result in any of the agent types being used (Storage Resource agent, SNMP agent, and SMIS agent).

Deleting or uninstalling a Storage Resource agent can result in agent assignment attempting to use an alternative agent to probe fabrics that the uninstalled Storage Resource agent might report on if such fabrics were part of a probe definition defined by the user.
9. If a fabric is in the process of rebuilding itself, a fabric probe job might fail. The fabric rebuilds itself if you add or remove a switch, or if you change the principal switch. If you experience a fabric probe job failure, suspect recent changes to the fabric which might have caused the fabric to rebuild. After the fabric rebuilds itself, the fabric probe job will likely succeed. Consider implementing a change-control process for the storage environment so that the previously mentioned changes are planned, approved, and recorded.
10. Your fabric probe or zone control job can fail on the AIX operating system with error message: AGT0430I.

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

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

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

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

For information about the supported agent types by switch vendor, see “Planning for Brocade management agents” on page 85 or “Planning for QLogic or Cisco fabric management agents” on page 89.

Storage Resource agent

The Storage Resource agent collects information about the SAN and sends that information to the Tivoli Storage Productivity Center Device server. The Storage Resource agent can gather topology information for the entire fabric. To gather host-level and detailed Host Bus Adapter (HBA) information, the agent must be installed on each host where that information is wanted.

The Storage Resource agent performs the following functions:

- Gathers information about the SAN by querying switches and devices for topology and identification information.
- Gathers host-level information for the local system.
- Gathers information about the zoning and allows zone control (except for Brocade switches) of the fabric.
- Gathers information about the HBAs installed on the local system, including make, model, and driver versions.
- Gathers event information detected by the HBAs and forwards it to the Device server.

Limitations to note are:

- Can gather only host-level and HBA information about the host containing the Storage Resource agent.
- Can provide only detailed identification for devices that are in the same zone as the Storage Resource agent. The Storage Resource agent identifies the device by worldwide name (WWN) if the Storage Resource agent is not in the same zone as the device rather than being able to uniquely identify the device by device type (for example, host, and subsystem).
- Identification of endpoint devices is reliant on the endpoint device responding to RNID or other queries.
- Can gather only a subset of the switch attributes. There are some switch attributes that can be collected only by using out-of-band Fabric agents.
- Tivoli Storage Productivity Center does not collect information about aliases through the Storage Resource agents.

Out-of-band Fabric agent

IBM Tivoli Storage Productivity Center uses SNMP queries to discover information about the SAN. Management Information Base (MIB) information is collected from the switches and directors by the out-of-band Fabric agent. Switches and directors are added as out-of-band Fabric agents and contacted from the Tivoli Storage Productivity Center Device server by SNMP.

The out-of-band Fabric agent performs the following functions:

- Gathers information about the fabric by querying the switch or director for topology information.
- Gathers virtual SAN information for Cisco switches.

Limitations to note are:

- Topology information is only gathered for the switch added as an out-of-band Fabric agent. The agent cannot gather the topology information for any connected switches unless they are also defined as out-of-band Fabric agents. If you are performing out-of-band discovery only on a fabric with several switches, you must install an out-of-band Fabric agent for each switch in the fabric to discover the whole fabric.
- Device information is limited; most devices are unknown with a type equal to "Other" and identified by their WWN.
- There must be a TCP/IP connection between the switch and the Tivoli Storage Productivity Center Device server.
- To enable events from the switch to the Device server, the switch must be configured to send SNMP traps to the Device server. The SNMP traps are synonymous to the events that are received by the in-band Fabric agent.

Configuration guidelines for agent placement

This section provides configuration guidelines for various agent types required for fabric functions.

The placement of out-of-band agents, Storage Resource agents, and CIM agents determines the information displayed in the topology. Switches must be configured to send SNMP traps to the server for fabric event notification. Follow these guidelines to determine the best agent placement strategy for these agent types.

- For a topology map, if no CIM agents are present, you must define Storage Resource agents and out-of-band agents to discover all of your topology. If CIM agents are present for Brocade fabrics, Storage Resource agents are still required to collect detailed information about endpoints and HBA information. If no CIM agents are present, switch zoning and LUN masking might restrict access to some devices, and some zone control features are unavailable.
- For a complete topology map, including correct device icons, you must define Storage Resource agents and out-of-band agents, or Storage Resource agents and CIM agents (for Brocade), on all servers and switches.
- For information about Brocade zoning, you must configure CIM agents for the Brocade fabric.
- For complete device centric and host centric views, you must place Storage Resource agents on all servers you want displayed.
- Cisco requires an out-of-band discovery for VSAN information.

Before implementing out-of-band agents, Storage Resource agents, and CIM agents, formulate a clear plan of your environment and the information you want to collect.

For more information about what the agents collect and for the supported agent types for switch performance management and fabric zone configuration, see “Collecting data with the Fabric Manager” on page 91.

Planning for Fabric Manager configuration

You must configure the main SAN devices properly before using IBM Tivoli Storage Productivity Center.

There are two main components to setting up Tivoli Storage Productivity Center:

- SAN environment setup
- Tivoli Storage Productivity Center setup

The essential elements of the SAN environment needed for Tivoli Storage Productivity Center are:

- Switches and directors
- Host Bus Adapters (HBAs)
- CIM agents for switches

The Storage Resource agent collects the full zone database information from the same switch in the fabric during a fabric probe. This condition assumes that there are no topology changes.

When there are topology changes and the agent state changes, Tivoli Storage Productivity Center might switch to reporting on the full zone database from a different switch in the fabric. When zone control is done through Tivoli Storage Productivity Center, a check is made to see if the zoning on the switch is the same as zoning information in the Tivoli Storage Productivity Center database. If the information is not the same, a warning message is issued before proceeding with zone changes.

This function is similar to selecting a data source for zoning for CIM agents. For CIM agents, the agent managing the principal switch in the fabric is the default data source for zoning. You have the option of specifying an alternative data source for zoning using the Tivoli Storage Productivity Center user interface.

For any fabric, the active configuration information is the same regardless of which switch in the fabric is queried by Tivoli Storage Productivity Center. However, zoning entities in the inactive configuration (also referred to as full zone database) can be different depending on the switch in the fabric that was queried.

Switches have a configuration parameter that controls propagation of inactive zoning entities to all switches in the fabric. If the parameter is set to propagate inactive zoning entities to every switch in the fabric, then any time an inactive zone entity changes, this change is propagated to all switches in the fabric so all switches have the same full zone database.

For Storage Resource agents, the selected data source is automatically determined by Tivoli Storage Productivity Center. If the Storage Resource agent is connected to multiple switches in the fabric, then the agent collects zoning information using only one switch (so that redundant fabric information is not collected from

multiple switches). Not all Storage Resource agents connected to a fabric are used to collect zoning information (only one switch is used).

You might see different inactive zone sets displayed because of the following conditions for the Storage Resource agent:

- A different Storage Resource agent is used to collect information
- There was a topology change
- The Storage Resource agent was stopped
- The Storage Resource agent was uninstalled
- A new Storage Resource agent was deployed

You might see different inactive zone sets displayed because of the following conditions for the switch:

- A different CIM agent is used
- The switch is not configured to propagate inactive zoning entities

If the switch configuration parameter is not set to propagate the information, then it is possible for switches in the fabric that have full zone database to have different zoning entities (active configuration is always the same for all switches in the fabric).

To configure the fabric, complete the following steps:

1. Install and configure the SMI-S agent (CIM agent) for the switch. Contact your switch vendor for instructions on how to install and configure the SMI-S agent.
2. For every fabric you want to monitor, configure one or more switches in the fabric to send SNMP traps to the Device server. This action causes SNMP traps to be sent to Tivoli Storage Productivity Center whenever there are changes in the fabric. Tivoli Storage Productivity Center can then perform automatic fabric discoveries and probes in response to these traps.
3. For each fabric you want to monitor, install a Storage Resource agent on each host that is connected to the fabric of interest.
4. Run a CIMOM discovery job for the CIM agents. Wait for the job to complete successfully for the CIM agents.
5. Run a fabric probe job. Wait for the job to complete successfully for the fabric.
6. You can now view information for the fabric and create a switch performance monitor job.

These are limitations for switches:

SNMP V1 protocol

Switches that are being managed out-of-band such as SNMP agents must be configured to use the SNMP V1 protocol. Some switches use SNMP V2 or V3 by default and must be reconfigured before they can work correctly with Tivoli Storage Productivity Center. Refer to the switch vendor documentation for specific configuration information.

HBA API for HP-UX

The HBA API for the HP-UX platform does not support the API function that is required for the in-band Fabric agent to support the collection of SAN events through the HBA. If there are no other Fabric agents (out-of-band or in-band agent) that is covering a particular SAN, or no switch in the fabric is configured to send SNMP traps to the Device server, an automatic discovery that would have normally been triggered due to reporting of such events will not occur.

Users might still be able to have a Tivoli Storage Productivity Center agent report SAN events for a particular SAN and have automatic triggering of discovery if one of the following conditions is true:

- Installing a Storage Resource agent on another supported host (non-HP-UX platform) that is connected to the same SAN that the HP-UX Fabric agent or Storage Resource agent is connected to.
- Defining out-of-band agents that are configured to get SNMP traps from switches in that SAN. You must consider this option carefully because there can be performance implications for this option.

Performance statistics

Some switch vendors do not return all the performance statistics required by Tivoli Storage Productivity Center. In these cases, some performance statistics might report as zero.

Note: A CIMOM discovery and fabric probe are required for all switch vendors supporting SMIS 1.1 and higher.

In the Storage Management Initiative Specification (SMI-S), peak data rates for a port are reported in the CIM_FcPortRateStatistics class. For SMI-S Version 1.1, this class is an optional class, meaning that it might not be supported by a given vendor. This can be the case if the Tivoli Storage Productivity Center user interface shows zero values for the peak rates, even when there is non-zero traffic flowing through the port. When a vendor does support this class, the peak data rate represents the peak value of the associated metric counter (as reported by the CIMOM used to access the device) over a recent, vendor-defined measurement window (for example, 20 milliseconds) of the Tivoli Storage Productivity Center measurement window (for example, 15 minutes).

Index number of a fiber channel port on a blade

Currently there is no SMI-S standard property that holds the index value number of a fiber channel port on a blade. For this reason, the Tivoli Storage Productivity Center SMIS Fabric Topology Table Mappers must have vendor-specific code for calculating the index. At this time, only Cisco and Brocade implement the Blades Subprofile, so index calculations are done only for Cisco and Brocade.

Switch blade port index values

The switch blade port index values might be incorrect for Brocade switches where one of the following scenarios apply:

- When a port swap operation is performed
- When the Extended Edge PID mode has been enabled on the director. For more information about the Extended Edge PID mode, see the Brocade document *Fabric OS Administrator's Guide*.

Planning for Fabric Manager switches and directors

For the Fabric Manager to gather and display information about switches and directors, the switches and directors must be configured correctly. The configuration varies between vendors. Incorrect switch configuration can lead to missing information and misconceptions about the Tivoli Storage Productivity Center product not working properly with certain switches.

For the switch to successfully receive and respond to queries, there are some basic requirements:

- The FC Management MIB and FE MIB must be enabled on the switch.

- The switch must be configured to receive SNMPv1 queries and respond in SNMPv1. Some switches are configured to use SNMPv2 or SNMPv3 by default.
- The community string configured in Tivoli Storage Productivity Center must match one of the community strings configured on the switch with read access. Cisco switches must additionally have a community string match for write access. The default community strings in Tivoli Storage Productivity Center are "public" for read access and "private" for write access. Additional community strings can be defined on the switches, but are not used.
- SNMP access control lists need to include the Tivoli Storage Productivity Center Device Server. Some automatically include all hosts while others exclude all hosts by default.

Another aspect of the SNMP configuration includes trap notification. SNMP traps are generated by the switch and directed to the Tivoli Storage Productivity Center Device Server as an indication that something in the fabric has changed and that a discovery must occur to identify changes. The default configuration for handling switch traps is to send them from the switch to port 162 on the Tivoli Storage Productivity Center Device Server. For the successful generation and reception of traps, there are some configuration requirements:

- The trap destination must be set. The trap destination is typically the Tivoli Storage Productivity Center Device server, but it can also be an intermediary SNMP manager that receives the trap and sends it to Tivoli Storage Productivity Center. The trap destination is set on the switch.
- The destination port must be set on the switch. Tivoli Storage Productivity Center listens on port 162 by default.
- The traps must be sent as SNMPv1. The SNMPv1 designation is set on the switch.
- The trap severity level must be set to generate traps for change conditions. The trap severity level indicates that error level traps and anything more severe are sent. The trap severity level is set on the switch.

Note:

- Storage Resource agents do not collect HBA events.
- If you want to monitor the fabric, you should have one or more switches configured to send SNMP traps to the server.

Configuring these settings differs between switch vendors and models. For information about configuring supported switches, see <http://www.ibm.com/developerworks>. Use the search field provided to search for "configure switches." Open the article "Configure switches successfully for Tivoli Storage Productivity Center."

Planning for host bus adapters (HBAs)

Proper installation of the HBA driver is an essential element for enabling data collection and is one of the most common problems when an agent appears to be healthy, but is not discovering any data other than the host information. Some HBA vendors require a separate installation of the HBA API package while others include it with the driver.

For the switches to successfully receive and respond to queries, there are some basic requirements:

- The switch must support the FC-GS-3 standard interface for discovery:
 - Name server

- Configuration server
- Unzoned name server
- For zone control functions, the fabric zone server must be supported, except in the case of Brocade switches.

There are some basic requirements on the HBA for proper agent operation:

- The HBA driver must be installed.
- The HBA API package must be installed. Some vendors include the HBA API package as part of their driver.
- HBA firmware is at a level that IBM Tivoli Storage Productivity Center supports.
- The current state of the HBA as reported by the operating system is good.
- The HBA has a good connection to the fabric.

Fabric events are automatically sent from the agent to the Tivoli Storage Productivity Center Device server. There is no need for configuration.

Note: This item applies only to the in-band Fabric agents. Storage Resource agents do not send fabric events to the server.

Planning for private switch networks

Some switch vendors recommend a private IP network for the fiber channel switches. It is important to understand that the IBM Tivoli Storage Productivity Center Device server cannot communicate with the switches if they are on a private IP network. It is important to note that the out-of-band Fabric agents require a TCP/IP connection from the the Device server to the switch. Also, SNMP traps from the switches cannot travel directly from the switches to the the Device server.

If you are using a private switch IP network you must rely on Storage Resource agents to gather your SAN information and to forward SAN events to the Device server.

Another option that is sometimes used with a private switch network is to allow the Tivoli Storage Productivity Center Device server to communicate on the private switch network using a second network interface card (NIC).

Planning for performance management

IBM Tivoli Storage Productivity Center can collect performance data for devices that use the native interfaces (DS8000, XIV system, SAN Volume Controller, Storwize V7000 Unified, or Storwize V7000 storage systems), or for storage systems and fiber channel switches that use CIM agents. The devices that use CIM agents must be SMI-S 1.1 compliant.

There are three main functions for Tivoli Storage Productivity Center performance management:

- Performance data collection (performance monitoring)
- Performance thresholds and alerts
- Performance reports

You can define how long performance data is stored in the database. Tivoli Storage Productivity Center can set thresholds for important performance metrics. When any boundary condition is crossed, Tivoli Storage Productivity Center can notify

the user through email, SNMP, or other alerting mechanisms. Tivoli Storage Productivity Center can generate reports, historic trend charts, and help analyze the bottleneck of a performance congestion by drilling down to threshold violated components and the affected hosts. A combination of these functions can be used to:

- Monitor a complicated storage network environment
- Predict warning signs of system failure
- Do capacity planning as overall workload grows

For more information about performance planning, see <http://www.ibm.com/support/docview.wss?uid=swg27017573>.

Performance metrics

IBM Tivoli Storage Productivity Center can report on many different performance metrics, which indicate the particular performance characteristics of the monitored devices.

Two important metrics for storage systems are throughput in I/O per second, and the response time in milliseconds. Throughput is measured and reported in several different ways:

- Throughput of an entire box (subsystem)
- Each cluster (TotalStorage Enterprise Storage Server)
- Each controller (DS6000 or DS8000)
- Each I/O Group (SAN Volume Controller or Storwize V7000)

Throughputs are measured for:

- Each volume (or LUN)
- At the Fibre Channel interfaces (ports) on some of the storage boxes
- On fiber channel switches
- At the RAID array after cache hits have been filtered out

For storage systems, the performance statistics are separated into frontend I/O metrics and backend I/O metrics. Frontend I/O metrics are a measure of the traffic between the servers and storage systems. Backend I/O metrics are a measure of all traffic between the storage system cache and the disks in the RAID arrays in the backend of the subsystem. Most storage systems give metrics for both kinds of I/O operations: frontend and backend operations. It is important to know whether the throughput and response times are at the frontend (close to the system level response time as measured from a server) or backend (between the cache and disk).

Note: A CIMOM discovery and fabric probe are required for all switch vendors supporting SMIS 1.1 and higher.

The main frontend throughput metrics are:

- Total IO rate (overall)
- Read IO rate (overall)
- Write IO rate (overall)

The corresponding frontend response time metrics are:

- Overall response time

- Read response time
- Write response time

The main backend throughput metrics are:

- Total backend IO rate (overall)
- Backend read IO rate (overall)
- Backend write IO rate (overall)

The corresponding backend response time metrics are:

- Overall backend response time
- Backend read response time
- Backend write response time

It is desirable to track any growth or change in the rates and response times. It frequently happens that I/O rate grows over time, and that response time increases as the I/O rates increase. This relationship is what "capacity planning" is all about. As I/O rates increase, and as response times increase, you can use these trends to project when additional storage performance (as well as capacity) is required.

Depending on the particular storage environment, it might be that throughput or response time times change drastically from hour to hour or day to day. There might be periods when the values fall outside the expected range of values. In that case, other performance metrics can be used to understand what is happening. Here are some additional metrics that can be used to make sense of throughput and response times:

- Total cache hit percentage
- Read cache hit percentage
- Write-cache delay percentage (previously known as NVS full percentage)
- Read transfer size (KB/operation)
- Write transfer size (KB/operation)

Low cache hit percentages can drive up response times, because a cache miss requires access to backend storage. Low hit percentages also tend to increase the utilization percentage of the backend storage, which might adversely affect the backend throughput and response times. High write-cache delay percentage (previously known as NVS full percentage) can drive up the write response times. High transfer sizes typically indicate more of a batch workload, in which case the overall data rates are more important than the I/O rates and the response times.

All these metrics can be monitored through reports or graphs in the Tivoli Storage Productivity Center GUI. Some examples of supported thresholds are:

- Total I/O rate and total data rate thresholds
- Total backend I/O rate and total backend data rate thresholds
- Read backend response time and write backend response time thresholds
- Total port I/O rate (packet rate) and data rate thresholds
- Overall port response time threshold
- Port send utilization percentage and port receive utilization percentage thresholds
- Port send bandwidth percentage and port receive bandwidth percentage thresholds

For fiber channel switches, the important metrics are total port packet rate and total port data rate, which provide the traffic pattern over a particular switch port. Port bandwidth percentage metrics are also important to provide an indicator of bandwidth usage based on port speeds. When there are lost frames from the host to the switch port, or from the switch port to a storage device, the dumped frame rate on the port can be monitored.

The important things are:

- Monitor the throughput and response time patterns over time for your environment
- Develop an understanding of expected behaviors
- Investigate the deviations from normal patterns of behavior to get warning signs of abnormal behavior
- Generate the trend of workload changes

Performance management configuration

Performance data can be collected from devices that are managed through the native interfaces or with devices that are managed by a CIM agent. For those devices that use the CIM agent, the devices must be SMI-S compliant.

To get performance data for devices, complete the following steps:

1. Prepare for performance monitoring.

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

- The version of CIM agent and firmware for the device is supported. For information about the CIM agent and firmware supported, see <http://www.ibm.com/support/docview.wss?uid=swg21386446>. Select the appropriate version for **Storage**.
- A CIM agent is installed on a different server from IBM Tivoli Storage Productivity Center. For some switch vendors, the switch comes with an embedded CIM agent, and therefore, no separate CIM agent installation is required.
- For subsystems or switches on a private network, be sure to have the CIM agent installed on a gateway machine so that the Tivoli Storage Productivity Center server on a different machine can communicate with the CIM agent.
- The CIM agent is configured to manage the intended device.

For information about how to install and configure the CIM agent, contact your CIM agent provider. Also see *IBM Tivoli Storage Productivity Center Hints and Tips* for information about installing and configuring CIM agents. See <http://www.ibm.com/support/docview.wss?uid=swg27008254>.

2. Select the steps to follow for your device:

- For the devices that use the native interface (DS8000, the XIV system, SAN Volume Controller, and Storwize V7000), follow these steps:
 - a. Add the device to Tivoli Storage Productivity Center. Use the Configure Devices wizard to add the device. For information about the Configure Devices wizard, see “Configure Devices wizard” on page 240. On the Select Device Type panel, select **Storage Subsystem**.
 - b. After the device is added to Tivoli Storage Productivity Center, automatic discovery and probe operations are done for the storage systems.
- For devices that use the CIM agent (ESS, DS3000, DS4000, DS5000, DS6000 all non-IBM devices), complete these steps:

- a. Add the CIM agent to Tivoli Storage Productivity Center. For information about adding a CIM agent, use the Configure Devices wizard. For information about the Configure Devices wizard, see “Configure Devices wizard” on page 240.
 - b. Run a CIMOM discovery job for the storage system or switch.
 - c. Run a probe job for the storage system or switch. The probe job gets all component level information into the database. The probe job also includes getting all the instances of controllers, arrays, extents, volumes, LUN to host mappings, FC ports, and so on. Only those storage systems for which configuration information is successfully collected can have performance monitor jobs run on them.
 - d. For fiber channel switches, run a fabric probe job for the fabric.
3. Create threshold alerts. A performance threshold is a mechanism by which you can specify one or more boundary values for various performance metrics, and can specify to be notified if any collected performance data violates these boundaries.

Tivoli Storage Productivity Center has default thresholds defined (but no default alerts). If you explicitly define an alert for one of the default thresholds, that means that if the threshold is violated, the constraint violation report displays the violation but the alert log will not.

When thresholds are applied during the processing of data collection, a performance monitor for a device must be actively running for a threshold violation to be recognized. Chronologically, you would need to define the threshold alert before the performance monitor job starts.

There is a special report that can provide statistics about how often and when an alert threshold has been reached, even without setting up alert destinations like email receivers or SNMP trap destinations. This special report is called a constraint violation report. To get a constraint violation report, expand **Disk Manager > Reporting > Storage Subsystem Performance > Constraint Violations**.

Note: Alerts can be created, deleted, or modified at any time, even if the performance monitor is already running. Changes are applied dynamically to the running job and take effect the next time data is collected for the job.

4. Create a performance monitor job. Performance data collection is performed through a performance monitor job. This job can be run immediately, can be scheduled to run one time only, or can be scheduled to run repeatedly as wanted. Only after the device has been probed successfully is that device eligible for a performance monitor. If the probe is not successful, that device will not show up in the performance monitor creation panel.

Performance management for storage systems is done through the Disk Manager. Data collection for the storage systems can be scheduled through **Disk Manager > Monitoring > Subsystem Performance Monitors**. Right-click on **Subsystem Performance Monitors**. Click **Create Subsystem Performance Monitor**. In the content pane, enter the information schedule and run a performance monitor job. To view the performance reports, expand **Disk Manager > Reporting > Storage Subsystem Performance**.

Performance management for fiber channel switches is done through the Fabric Manager. Data collection for the switches can be scheduled through **Fabric Manager > Monitoring > Switch Performance Monitors**. Right-click on **Switch Performance Monitors** and click **Create Switch Performance Monitors**. To view the performance reports, expand **Fabric Manager > Monitoring > Switch Performance Monitors**.

5. Check the performance monitor status. When the performance monitor job starts to run, you begin to collect performance data for the device. Check the status of the performance monitor job to make sure that it runs and continues running. In the navigation tree, expand **IBM Tivoli Storage Productivity Center > Job Management**. In the content pane, click the **Refresh** button because refreshes are not automatic on the Job Management panel. See Jobs for Selected Schedule for a list of performance monitor jobs and the status of the jobs. For more information about the Job Management panel, see "Managing schedules, jobs, and runs" on page 257.
6. Specify the retention period for the performance data. After the monitor job is created, you can configure how long to keep performance data. You can specify a retention period for collected sample data, for aggregated hourly data, and for daily data. The retention period applies to all performance monitors. For more information about the retention period, see "Resource History Retention" on page 300.

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

For storage systems, the most numerous components are usually volumes, therefore, the largest amount of performance data collected is for volumes.
7. View the performance reports in Tivoli Storage Productivity Center.

Microsoft Cluster Server

Data Manager can monitor and report on Microsoft Cluster Server (MSCS) clustered nodes and cluster resource groups.

Microsoft Cluster Server (MSCS) is a built-in feature of the Windows operating system. It is software that supports the connection of up to eight servers into a "cluster" for higher availability and easier manageability of data and applications. MSCS can automatically detect and recover from server or application failures. It can be used to move server workload to balance utilization and to provide for planned maintenance without downtime.

As with standard Windows clusters, you perform most of the configuration tasks, and also the management tasks, associated with Exchange clusters using the Cluster Administrator. Cluster Administrator is installed by default on servers that have Cluster Service installed.

You can also use Cluster Administrator to remotely administer a server cluster. Computers that are used to administer a server cluster remotely must be secure and restricted to trusted personnel.

When a cluster node is probed, there are normal errors in the logs for the disks belonging to the cluster resource groups that are not currently hosted on this node.

When looking at an MSCS cluster problem, the current state of the clustering environment is often needed. Microsoft provides a tool to dump the MSCS cluster configuration into a set of files.

Note: The term "virtual servers" is replaced with "cluster resource groups" in Tivoli Storage Productivity Center. Microsoft still refers to "virtual servers."

Microsoft Cluster Server environment

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

Report information can include local resources for a clustered node and the resources in a cluster resource group. The Storage Resource agent is not cluster aware so the Storage Resource agent cannot run in the cluster resource group and fail over from one node to another. The Storage Resource agent supports failover of a clustered file system so that Tivoli Storage Productivity Center scan jobs can continue to work when the clustered file system is moved from one node to another. The Storage Resource agent running on the node hosting the cluster resource group discovers the resources on that cluster resource group. The Storage Resource agents must be installed on every node in the cluster.

The following information lists the requirements for using the Data server in an MSCS cluster environment.

Server requirement

Not applicable. The server component can be installed on a node in an MSCS cluster, but the Data Manager can monitor and report only on nodes that have an agent installed.

agent requirements

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

client requirements

Not applicable. The client component can be installed on a node in an MSCS cluster, but the Data Manager can monitor and report on only the nodes with an agent installed.

Note: The Data server does not currently support the monitoring of clustered database application.

If the Storage Resource agent port is behind a firewall, that port needs to be opened for all virtual server and cluster resource group addresses.

Microsoft Cluster Server support

Use this information to plan for running IBM Tivoli Storage Productivity Center with MSCS.

Cluster resource groups

A *cluster resource group* represents an MSCS entity.

You can configure MSCS into cluster resource groups. You can define resource group policies and attributes that dictate how MSCS manages resources to keep them highly available at different stages of cluster operation (startup, failover, and fallback). The types of resources you can put into clustered resource groups are:

Logical volumes

A set of logical partitions that MSCS makes available as a single storage unit that is, the logical volume is the “logical view” of a physical disk.

Filesystem

A file system is written to a single logical volume. Typically, you organize a set of files as a file system for convenience and speed in managing data.

Shared filesystems

A file system that resides entirely in a shared logical volume.

Physical disk

A physical disk.

Network name

A network name.

IP address

An IP address.

Reports

The asset reporting navigation tree includes a By Cluster report that shows the nodes and cluster resource groups in each cluster. The disks and file systems that are shared as clustered resources are associated with the cluster resource group and not with the node where they are hosted.

When generating reports, you can use the By Cluster report subtype wherever the By Computer subtype is available. The By Cluster report summarizes the results of the nodes and the cluster resource groups in the cluster.

Monitoring groups

When you create a monitoring group or create a data collection schedule, you can select individual nodes and cluster resource groups. You can create a file system monitoring group that includes clustered and non-clustered resources.

Computer alerts

You can register for cluster-specific alerts that notifies you that your cluster has gone through a configuration change such as a cluster resource being added or removed. You can also register for an alert that is triggered when a cluster resource group is moved. The same alert is used for a cluster resource group failover, cluster resource group fallback, and when the cluster resource group is manually moved to a new node.

Cluster alerts are not triggered until the next time the cluster is probed.

Upgrading agents in an MSCS cluster

Take the following steps when upgrading the Tivoli Storage Productivity Center agents in an MSCS cluster:

- Make sure that all agents in the cluster are configured to use the same port address.
- Use either the Tivoli Storage Productivity Center installation program or the graphical user interface to perform the agent upgrades.
- Do not probe or scan the agents in the cluster while the agent upgrade is in progress.

- Once all agents in the cluster have been upgraded successfully, create and run a new probe job on the cluster.

For more information about upgrading Tivoli Storage Productivity Center agents, see “Upgrading Storage Resource agents” on page 363.

Installing MSCS support

To install MSCS support, complete the following steps:

1. Make sure that your MSCS environment is configured for Tivoli Storage Productivity Center support.
2. Install a Storage Resource agent on each node of the cluster. All agents in the cluster must be configured to use the same listening port.
3. Run a discovery job for the agent.
4. Add a Network Name resource with an IP address that can be contacted from the Tivoli Storage Productivity Center server.
5. Run a scan job for the agent.
6. Run a probe job for the agent.
7. View reports for the MSCS entities.
8. View the MSCS environment in the topology viewer.

Planning for PowerHA

Tivoli Storage Productivity Center supports Storage Resource agents installed on PowerHA nodes (previously known as HACMP). Use this information to configure the PowerHA environment before using it with Tivoli Storage Productivity Center.

Cluster resource groups

A *cluster resource group* represents a PowerHA entity.

You can configure PowerHA into cluster resource groups so that they can be highly available. You can define resource group policies and attributes that dictate how PowerHA manages resources to keep them highly available at different stages of cluster operation (startup, failover, and fallback). The types of resources you can put into clustered resource groups are:

Volume groups

A set of physical volumes that AIX treats as a contiguous, addressable disk region.

Logical volumes

A set of logical partitions that AIX makes available as a single storage unit that is, the logical volume is the “logical view” of a physical disk.

Filesystem

A file system is written to a single logical volume. Typically, you organize a set of files as a file system for convenience and speed in managing data.

Shared filesystems

A journaled file system that resides entirely in a shared logical volume.

Applications

Critical services that are accessible to users.

Service IP labels or addresses

A way to establish communication between client nodes and the server

node. Services, such as a database application, are provided using the connection mode over the service IP label.

Tape resources

You can configure a SCSI or a Fiber Channel tape drive as a cluster resource in a non-concurrent resource group, making it highly available to two nodes in a cluster.

Communication links

You can define communication links as resources in a PowerHA resource group.

Planning for PowerHA support

With PowerHA, you can use the following types of configurations:

- PowerHA with cluster resources groups that are nonconcurrent. Concurrent cluster resource groups are not supported.
- The following volume groups in a PowerHA environment:
 - Standard volume groups
 - Enhanced concurrent-mode volume groups
 - Scalable volume groups

When you create a Tivoli Storage Productivity Center configuration, include the Storage Resource agent. A Storage Resource agent must be installed on each node of the cluster. All agents in a cluster must be configured to use the same listening port. You must also have a cluster resource group (CRG) with at least one IP address that is accessible from the Tivoli Storage Productivity Center server if you plan to perform scans on the cluster resource group. Each Storage Resource agent collects information about the local node and the clustered resources that are currently hosted on the node.

The following operations are not supported:

- Probing and scanning of clustered databases
- Reporting on cluster information when the cluster is configured to have concurrent cluster resource groups

Note:

- Cluster resource group scanning is only supported by the Tivoli Storage Productivity Center Advanced license.
- If the Storage Resource agent port is behind a firewall, that port needs to be opened for all virtual server and cluster resource group addresses.

Reports

The asset reporting navigation tree includes a By Cluster report that shows the nodes and cluster resource groups in each cluster. The disks, volume groups, and file systems that are shared as clustered resources are associated with the cluster resource group and not with the node where they are hosted.

When generating reports, you can use the By Cluster report subtype wherever the By Computer subtype is available. The By Cluster report summarizes the results of the nodes and the cluster resource groups in the cluster.

Monitoring groups

When you create a monitoring group or create a data collection schedule, you can select individual nodes and cluster resource groups. You can create a file system monitoring group that includes clustered and non-clustered resources.

Computer alerts

You can register for cluster-specific alerts that notify you that your cluster has gone through a configuration change such as a cluster resource being added or removed. You can also register for an alert that is triggered when a cluster resource group is moved. The same alert is used for a cluster resource group failover, cluster resource group fallback, and when the cluster resource group is manually moved to a new node.

Cluster alerts will not be triggered until the next time the cluster is probed.

Installing PowerHA support

To install PowerHA support, complete the following steps:

1. Ensure that your PowerHA environment is configured for Tivoli Storage Productivity Center support.
2. Create a service IP label for the CRG with an IP address that can be contacted from the Tivoli Storage Productivity Center server.
3. Install a Storage Resource agent on each node of the cluster. All agents in the cluster must be configured to use the same listening port.
4. Run a discovery job for the agent.
5. Run a probe job on the agents.
6. Run a scan job on the agents.
7. Run a scan job on the cluster resource groups.
8. View reports for the PowerHA entities.
9. View the PowerHA environment in the topology viewer.

PowerHA environment

This topic provides information about Tivoli Storage Productivity Center support in a PowerHA environment.

Tivoli Storage Productivity Center supports Storage Resource agents installed on PowerHA nodes. The Storage Resource agent must be installed on every node in the cluster. You cannot configure the Storage Resource agent as a clustered application.

The Storage Resource agent collects information from the cluster when the node is probed. The information that is collected is available in the Data Manager and Disk Manager reports and the topology viewer. Resources that are not clustered are reported under the node. The resources that are associated with a cluster resource group are reported under a computer entity that represents the cluster resource group. For example, if a single node cluster that has one cluster resource group is probed it produces two computer entities in the Data Manager and Disk Manager reports and the topology viewer.

The following entities are reported under the associated clustered resource group and not the node:

- Volume groups
- NFS shares
- Service IP labels

Physical volumes, logical volumes, and file systems for clustered volume groups are also reported under the cluster resource group. Scan and ping jobs can be created for a cluster resource group and these jobs work regardless of which node is hosting the cluster resource group. Scan and ping requests for a cluster resource group are sent to the service IP label defined for the cluster resource group. Therefore, in order for these jobs to succeed the service IP label must be accessible from the Data server and the CRG must be online.

Restriction: You can use PowerHA only with cluster resource groups that are nonconcurrent. You cannot use PowerHA with concurrent cluster resource groups.

The following information lists the support requirements for a PowerHA cluster environment.

Agent requirement

Each node in a PowerHA cluster must have a Storage Resource agent installed. The Data Manager can only monitor and report on nodes that have a daemon Storage Resource agent installed. If you are using daemon Storage Resource agents, all agents installed in a cluster must listen on the same port number. Agents in other clusters can be configured with a different port address. However, all agents in the cluster must use the port address used by the other agents in that cluster.

You can use either Storage Resource agents or Data agents on each node. You can have a mixed configuration when migrating Data agents to Storage Resource agents, but scan jobs of clustered file systems will fail until all the nodes in the cluster have the same type of agent.

Probe requirements

Probes are not automatically executed in response to cluster events. Schedule probes to run as appropriate for the needs of the environment.

Scan requirements

To perform scans on a cluster resource group, the cluster resource group must be configured with an IP address that the Tivoli Storage Productivity Center server can communicate with. This IP address is displayed in reports associated with PowerHA clusters. If the server cannot contact this address once it has been configured, the server tries the next known address for the cluster resource group until it successfully contacts an IP address.

Note: Cluster resource groups are now listed as computers in scan job results, and now have log files specific to the cluster resource group.

NAS support

This section provides information for the planning of NAS support.

Network Attached Storage system requirements

This section provides information about the requirements for installing and running the Data Manager within a Network Attached Storage (NAS) environment.

IBM Tivoli Storage Productivity Center supports the monitoring of the NetApp Data ONTAP.

To be supported, NAS devices other than Network Appliance Filers must meet the following criteria:

- If scanned from a UNIX agent, a NAS device must support Network File System (NFS) queries.
- If scanned from a Windows agent, a NAS device must support Common Internet File System (CIFS) queries.
- A NAS device must support SNMP queries (sysName and sysOID).
- A NAS device must supply a unique sysName.
- If a NAS device hosts both Windows and UNIX file systems, you need at least two proxy Storage Resource agents to monitor the device. At least one Windows proxy Storage Resource agent is required to probe and scan the Windows CIFS shares. At least one UNIX proxy Storage Resource agent is required to probe and scan UNIX NFS shares.

For detailed information about configuring NAS and Tivoli Storage Productivity Center, see <http://www.redbooks.ibm.com/>. Search for **sg247490**.

Here are the requirements for installing and running the Data Manager within a NAS environment.

Server requirement

Not applicable. You cannot install the server on a NAS filer.

Agent requirements

You must install the agent on a machine that has access to the NAS filers within your environment that you want to monitor.

Windows

The agent that is logging in to and scanning the NAS filer is not required to be in the same domain as the user or the NAS filer. If you install the agent on a different domain from the NAS filer, the agent scans the NAS filer *if* the domain of the agent computer is a "trusted domain" by the domain of the NAS filer.

UNIX and Linux

The agent computer must import the NAS filer export files as NFS mounts (or automounts on Oracle Solaris).

Note:

- You do not install agents to the NAS filers themselves.
- For Tivoli Storage Productivity Center 5196 Network Attached Storage 300 G machines only, install the agent directly on those machines. If the 300 G is clustered, you need to install an agent on each local node.

Client requirements

Not applicable. You cannot install the client component to a NAS filer.

NAS requirements

The NAS filers within your environment must be visible to the machines where you install the agent or agents. If you want to monitor NAS filers from Windows, you must configure those NAS filers to be members of a Windows domain.

Planning for NetApp device support

NetApp devices provide unified storage for organizations with Network Attached Storage or storage area network environments that have file and block-level services. IBM Tivoli Storage Productivity Center provides the same file and block-level support for NetApp devices as is provided for other filers and storage subsystem devices.

Overview

Tivoli Storage Productivity Center provides the following support for NetApp devices:

- supports NetApp Data ONTAP SMI-S Agent 3.0
- supports the SMI-S 1.2 Array profile implemented by the NetApp SMI-S agent
- supports all filer models running NetApp Data ONTAP version 7.2 and 7.3

NAS Gateway and V Series are not supported with the Array profile.

The following licenses are required:

- For performance monitoring, an IBM Tivoli Storage Productivity Center license is required
- For SNMP discovery, an IBM Tivoli Storage Productivity Center Advanced license is required

Configuration considerations

The way you use Tivoli Storage Productivity Center to interact with NetApp devices, and the data you can collect for NetApp devices, depends on how you configure your NetApp devices. For example:

- If you configure a Storage Resource agent to be a Scan/Proxy agent, you use Data Manager to work with NetApp devices. The Data Manager user interface displays a NetApp filer as a "computer". For NetApp devices configured this way, Tivoli Storage Productivity Center collects file storage information.
- If you configure a NetApp SMI-S agent (CIMOM), you use Disk Manager to work with NetApp devices. The Disk Manager user interface displays a NetApp device as a "subsystem". For NetApp devices configured this way, Tivoli Storage Productivity Center collects block storage information.
- If you configure a NetApp device as both a filer and a subsystem, Tivoli Storage Productivity Center collects both file and block storage information. You can use Disk Manager and Data Manager to work with the NetApp device. Keep in mind the following considerations for the user interface:
 - In Data Manager, a NetApp filer is shown as a "computer".
 - In Disk Manager, a NetApp device is shown as a "subsystem" and a Fibre Channel (FC) LUN as a "volume".
- If multiple Storage Resource agents are set up to probe or scan the same NetApp filer, the Storage Resource agent that was added to Tivoli Storage Productivity Center first is used for the probe or scan. Therefore, only data that is gathered by the first Storage Resource agent is shown. For example, if you add a Storage Resource Agent for the Windows operating system and then you add a Storage Resource Agent for the UNIX operating system, Windows operating system related data such as shares are shown, but not exports.

If you initially configure a NetApp device using a Storage Resource agent, you do not lose any function if you later decide to also configure the NetApp device using a CIM agent.

General procedure to install and configure NetApp devices

To install and configure a NetApp device to enable the Data Manager feature, complete the following steps:

1. Select a machine to act as the Storage Resource agent that meets the auto-discovery requirements for Tivoli Storage Productivity Center to auto-discover the NetApp filers.
2. Install a Storage Resource agent on that machine.
3. Set the license for the NetApp filers that are discovered.
4. Set the Scan/Proxy agent for the discovered filers.
5. Create and run a probe job for each of the filers to collect Data Manager information.

To install and configure a NetApp device to enable Disk Manager functionality, complete the following steps:

1. Install Tivoli Storage Productivity Center or upgrade to version 4.2 or later.
2. Install the NetApp SMI-S agent and add filers to its configuration using the NetApp SMI-S agent utility
3. Add the NetApp CIMOM by using **Administrative Services > Data Sources > CIMOM Agents**. In the content pane, click **Add CIMOM**. To automatically discover the NetApp SMI-S agent using SLP, skip this step.
4. Run a CIMOM discovery.
5. If you want to use SLP to discover the NetApp SMI-S agent, make sure that your SLP directory agents are configured or the **Scan local subnet** box is checked on the **Administrative Services > Discovery > CIMOM**. In the content pane, go to the Options page to enter the IP addresses or host names for the SLP directory agents to be used before you run the CIMOM discovery.
6. Create and run a probe job for the configured filers to collect Disk Manager information.
7. If wanted, use the command-line interface (CLI) **tpctool** command to view information for NetApp storage systems.

Refer to the *IBM Tivoli Storage Productivity Center User's Guide* for more information about how to collect and view information gathered by NetApp devices. Refer to the *IBM Tivoli Storage Productivity Center Command-Line Interface Reference* for information about using the **tpctool** command.

Removing a NetApp device

If you later decide to remove a NetApp device, the steps to remove it depend on how you configured it. For example, if the NetApp device is configured only as a filer, you can remove it using the **Administrative Services > Configuration > Scan/Probe Agent Administration** page. If the NetApp device is configured only as a subsystem, you can remove it using the **Disk Manager > Storage Subsystems** page. If the NetApp device is configured as both a filer and a subsystem, then you must perform both of the preceding steps.

Limitations of NetApp support

Keep in mind the following limitations for NetApp support:

- The NetApp Data ONTAP SMI-S Agent 3.0 or later implements the Block Server Performance subprofile. It provides volume performance data but, for now, does not provide performance data at the storage subsystem level.
- Tivoli Storage Productivity Center supports only the SMI-S Array profile of the NetApp Data ONTAP SMI-S Agent. Other SMI-S profiles, including the self-contained NAS profile and the NAS Head profile, are not supported.
- For the Data ONTAP SMI-S 3.0 agent or later, if a volume is offline, the performance monitor might fail with the message: PM HWNPM2132W Performance data could not be collected for device *device*.
- To work with Network Appliance quotas (using **Data Manager > Policy Management > Network Appliance Quotas**), the NetApp device must be configured as a filer. This function is not available if the NetApp device is configured only as a CIMOM.

Identifying NetApp devices in the user interface

The way you configure a NetApp device determines how the NetApp device is represented in the Tivoli Storage Productivity Center user interface, as follows:

- NetApp device configured as a filer: the NetApp device appears as a "computer" in the user interface.
- NetApp device configured as a storage subsystem using a CIMOM: the NetApp device appears as a "subsystem" in the user interface.
- NetApp device configured as both a filer and storage subsystem: the NetApp device appears as either a "computer" or a "subsystem" in different places in the user interface. For example, some reports under My Reports refer to the NetApp filer as a "computer" while others refer to the NetApp device as a "subsystem".

The NetApp device is represented either as a "computer" or a "subsystem", or both, depending on how you have configured the NetApp device, are:

- Dashboard
- My Reports
- Topology viewer
- Alerting
- Data Manager
- Disk Manager

For example:

- **Data Manager > Reporting > Asset > By Storage Subsystem** lists a NetApp device as a storage subsystem if the device was probed by the NetApp Data ONTAP SMI-S agent.
- **System Reports > Data > Disk Defects** displays a NetApp device name as either a "computer" or a "subsystem", depending on whether the device information was collected by the Storage Resource agent or the SMI-S agent.

In the topology viewer, the NetApp device appears as follows:

- If a Storage Resource agent is configured to probe the NetApp device, the NetApp device is displayed as a "computer"
- If a CIM agent is configured to probe the NetApp device, the NetApp device is displayed as a "storage subsystem"

Planning for System Storage N Series Gateway servers

Tivoli Storage Productivity Center supports the System Storage N Series Gateway servers as the **Other NAS** node.

Upgrade information

If you are upgrading Tivoli Storage Productivity Center from TotalStorage Productivity Center 3.3.x, there are specific steps to follow. You must delete and unlicense existing N Series Gateway servers first before adding the N Series Gateway servers. For information about upgrading Tivoli Storage Productivity Center with N Series Gateway servers, see “Upgrading System Storage N Series Gateway servers” on page 370.

New installation

If you are installing Tivoli Storage Productivity Center as a new installation, the N Series Gateway servers can be automatically discovered after the Storage Resource agent is installed. If the N Series Gateway servers are automatically discovered, they are listed under **Administrative Services > Configuration > License Keys**. In the content pane, click the magnifying glass icon for the product to display the **Licensing** page. The **Licensing** page displays the OS Type as **NetApp Data ONTAP** with the Licensed box cleared.

To manually add the N Series Gateway servers as **Other NAS**, the Licensed box must be cleared. The OS Type and Licensed fields are updated after the N Series Gateway Server is manually added as **Other NAS**. For more information, see “Manually adding an N Series Gateway server ” on page 371.

Planning for VMware

Tivoli Storage Productivity Center supports the VMware Virtual Infrastructure which consists of the ESX Server and VMware VirtualCenter.

Overview

The ESX Server is a true hypervisor product which can host multiple virtual machines that run independently of each other while sharing hardware resources. The VirtualCenter is the management application that is the central entry point for the management and monitoring of the ESX Servers for a data center. Install a Storage Resource agent on each virtual machine that you want to monitor.

For more information about the ESX Server or VMware VirtualCenter, see <http://www.vmware.com>.

For information about permissions required by Tivoli Storage Productivity Center to access information from VMware, see the information center. Search for “VMware data source”.

The storage systems and platforms supported on VMware are shown in Table 25 on page 118.

Table 25. Supported platforms and storage systems for VMware

Supported platform	Supported storage systems	Supported function
ESX Server 3.0.x	<ul style="list-style-type: none"> • 3PAR • EMC CLARiiON • EMC Symmetrix • Hitachi Data Systems 9000 series • Hewlett Packard Enterprise Virtual Arrays (EVA) • IBM DS4000 • IBM DS6000 • IBM DS8000 • IBM ESS • IBM SAN Volume Controller • IBM Storwize V7000 • IBM XIV Storage System 	Monitoring LUN correlation
ESX Server 3.5.x		
ESX Server 3.5i		
ESX Server 4.0.x		
ESX Server 4.1.x		
VirtualCenter 2.0.1		
VirtualCenter 2.5		

For full function, both the Storage Resource agent and Virtual Infrastructure must be up and running. If one of the items is not present in a given environment, only a limited picture is presented to the user. Some virtual machines might not be recognized.

The hierarchical mapping of storage allocated to the virtual machine is available for the virtual machines on the ESX Server.

Note: Tivoli Storage Productivity Center now supports the mapping of storage from the ESX Server to the disk drives for the ESX Server 3.5.

Note the following limitations:

- No HBA virtualization is available for the VMware virtual machines.
- No events directly generated by the Virtual Infrastructure is supported.
- Data path explorer is not supported for VMware ESX Server and virtual machines.

Tivoli Storage Productivity Center for Replication in a VMware environment

Use Tivoli Storage Productivity Center for Replication to maintain a consistent copy of customer data. You can install Tivoli Storage Productivity Center for Replication under a VMware image. However, there are some aspects of running in a virtualized environment that can interfere with the ability of the software to function correctly if a disaster occurs. Be aware of the following limitations:

- Virtual machines in a VMware environment can be moved from one physical server to another using VMotion or other tools. If you move your virtual machine with Tivoli Storage Productivity Center for Replication on it, unpredictable results can occur during the processing of certain critical operations. This action might prevent Tivoli Storage Productivity Center for Replication from maintaining data consistency or prevent you from being able to quickly recover your data at the remote site.

- If the VMware image resides on the same storage that is managed by Tivoli Storage Productivity Center for Replication, Tivoli Storage Productivity Center for Replication might attempt to "freeze" I/O on the storage system. This freezing of I/O on the storage system can cause Tivoli Storage Productivity Center for Replication to stop responding. If this freeze occurs, the I/O is not released until the storage system freeze timeout occurs. Any applications running against that storage system might be affected by a delay, which can prevent Tivoli Storage Productivity Center for Replication from maintaining data consistency.
- Virtual images are controlled and managed by the ESX server. You must ensure that the internal VM network is set up correctly so users can access Tivoli Storage Productivity Center for Replication. Correctly setting up the network ensures that:
 - Tivoli Storage Productivity Center for Replication can access the storage systems that it is managing
 - Helps users quickly recover data on the remote storage systems if a disaster occurs

Also ensure that the virtual environment is secure and that it is monitored to maintain data consistency.
- Because virtual environments can vary, ensure that you thoroughly test your environment with Tivoli Storage Productivity Center for Replication, and ensure that you test disaster recovery-related steps.

Software requirements

For information about the software requirements for your ESX Server or VMware VirtualCenter environment, see "Software requirements for operating systems" on page 4.

Installing Tivoli Storage Productivity Center server on a virtual machine

You can install the Tivoli Storage Productivity Center server on a virtual machine on VMware ESX server 3.5.x. The hardware and operating system requirements are the same requirements as for a physical machine.

In addition, the following requirements must be met:

- CPU** For the ESX server, do not have more virtual processors than there are physical cores in the system. Plan your system so that no processor scheduling is required by the VM kernel for the virtual machine.
- RAM** Ensure that you have enough RAM in the ESX server to service all the virtual machines with a maximum RAM usage. Plan your system so that the ESX server does not need to swap RAM for the virtual machine.
- Disk** Use the SAN-attached RDM with SCSI pass-through for DB2 data and log storage.

For more information about VMware, see the following publications:

- "Using VMware ESX Server with IBM WebSphere Application Server" at http://www.vmware.com/partners/vmware/ESX_WAS_WP.pdf.
- "Scaling IBM DB2 9 in a VMware Infrastructure 3 Environment" at http://www.vmware.com/pdf/db2_scalability_wp_vi3.pdf.

VMware configuration

Before you can display reports or see the topology for VMware Virtual Infrastructure, you must configure the VMware environment.

To configure the VMware environment, complete the following steps:

1. If the VMware Virtual Infrastructure uses SSL certificates for communication, you have to use **keytool** to manually import the SSL certificates into a truststore. Each Virtual Infrastructure data source provides an individual certificate. There is a default truststore registered in the Device server's system properties file. **Keytool** is a tool shipped with the Java run-time environment. For information about how to import the SSL certificates, go to the Tivoli Storage Productivity Center information center. Search for **Importing SSL certificates for VMware**.
2. Add the VMware VI data source. The data source can be a hypervisor (ESX Server or VirtualCenter). Adding the data source is the first step in getting information from VMware Virtual Infrastructure. Adding a VMware data source is like adding a CIM agent or Storage Resource agent.
3. Test the connection to the VMware VI data source. This test ensures that you can access information from the VMware data source.
4. Run a discovery job for the VMware environment. The discovery is needed to retrieve every ESX Server instance that is part of the Virtual Infrastructure that has been added. The discovery mechanism is like a discovery for storage subsystems. Discovery jobs can be scheduled and are performed on the complete list of known VMware data sources.
5. Run a probe job for the ESX Server, hypervisor, and virtual machines. This step gets the detailed information from the hypervisors and virtual machines for IBM Tivoli Storage Productivity Center.
6. Configure alerts for VMware. You can create alerts for the following alert conditions:
 - Hypervisor discovered
 - Hypervisor missing
 - Virtual Machine added
 - Virtual Machine deleted
7. Install the Storage Resource agent on each of the virtual machines you want to monitor. For full functionality, you need two data sources.
8. You can now view VMware reports and VMware topology.

VMware capacity reports

You must probe both the ESX Server and the Storage Resource agent on the virtual machines before you can generate accurate reports for disk and file system capacity.

The TOTAL row in the capacity report shows the capacity for the file system or the disk capacity. The total includes virtualized disks, virtual machines, non-virtualized disks, and non-virtualized machines.

For example, you have an ESX Server that has a capacity of 100 GB and 60 GB is allocated to the virtual machine. The virtual machine uses 5 GB of space. Both the ESX Server (H1) and the virtual machine (VC1) are probed. You also have a physical machine (PC1) that is probed. The capacity is shown in the report, as follows:

Column heading	Capacity	Used space (calculated as Capacity minus Free Space)	Free space
TOTAL	130 GB	25 GB	105 GB
H1	100 GB	60 GB	40 GB
VC1	60 GB	5 GB	55 GB
PC1	30 GB	20 GB	10 GB

If you probed the physical machine (PC1) and the virtual machine (VC1) but did not probe the ESX Server (hypervisor), the capacity is shown as follows:

Column heading	Capacity	Used Space (calculated as Capacity minus Free Space; negative values are shown as 0)	Free Space
TOTAL	30 GB	0 GB	65 GB
PC1	30 GB	20 GB	10 GB
VC1	60 GB	5 GB	55 GB

If you probed the hypervisor (H1) and the physical machine (PC1) but did not probe the virtual machine (VC1), the capacity is shown as follows:

Column heading	Capacity	Used Space (calculated as Capacity minus Free Space)	Free Space
TOTAL	130 GB	80 GB	50 GB
H1	100 GB	60 GB	40 GB
PC1	30 GB	20 GB	10 GB

If you have an ESX Server that has devices that are directly attached to the virtual machine, the file system capacity created on the virtual machine is not added in the TOTAL row.

For example, the file system capacity shown in the table is based on the following factors:

- The file system (FS1) created on the hypervisor has a capacity of 100 GB and uses 60 GB of space.
- The file system (FS1) created from internal storage on the hypervisor (H1) has a capacity of 60 GB and uses 5 GB of space.
- The file system (FS2) on the device that is directly attached to the virtual machine (VC1) has a capacity of 5 GB and uses 4 GB of space.
- The file system created on the physical machine (PC1) has a capacity of 30 GB and uses 20 GB of space.

Column heading	Capacity	Used Space (calculated as Capacity minus Free Space)	Free Space
TOTAL	130 GB	24 GB	106 GB
H1 FS1	100 GB	60 GB	40 GB
VC1 FS1	60 GB	5 GB	55 GB
VC1 FS2	5 GB	4 GB	1 GB
PC1 FS1	30 GB	20 GB	10 GB

Planning for the Virtual I/O Server

You can use the Storage Resource agent to gather information about Virtual I/O Servers. Before you can monitor Virtual I/O Servers, you must plan on how to install the agents in your environment.

Overview

The Virtual I/O Server is part of the IBM PowerVM[®] hardware feature. The Virtual I/O Server allows the sharing of physical resources between LPARs including Virtual SCSI and virtual networking. Sharing of physical resources allows more efficient utilization of physical resources through sharing between LPARs and facilitates server consolidation.

Support for Virtual SCSI environment

The Virtual SCSI environment consists of a Virtual SCSI Server Adapter or Adapters that are created on the Virtual I/O Server and mapped to a Virtual I/O Client. If you have a Virtual SCSI environment and the virtual target device has a one to one mapping to a storage volume, Tivoli Storage Productivity Center can extract the storage system serial number and correlate the Virtual SCSI disk on the LPAR to a corresponding storage volume. This function means that the Storage Resource agent installed on the Virtual I/O Client (LPAR), can collect this information. The Linux system must be running on the IBM Power[®] System.

Note:

- Tivoli Storage Productivity Center performs the correlation only if the virtual target device is a one to one mapping to a single storage system volume.

The storage systems supported are:

- DS8000
- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- DS3000,
- DS4000
- DS5000

After the Storage Resource agent is installed on the Virtual I/O Server, you can see the following information:

- Storage system serial number

- Virtual SCSI Client Adapter
- Virtual SCSI disk
- Volume groups, logical volumes, and file system

Note:

- Currently, the correlation of Virtual SCSI disks that originate from the XIV system is not performed.
- Tivoli Storage Productivity Center can perform the correlation of Virtual SCSI disks only if AIX Multiple Path I/O (MPIO) or SDDPCM multipath driver is used on the Virtual I/O Server.
- If there is one or more Virtual SCSI disks present on an LPAR, Tivoli Storage Productivity Center considers the LPAR as a virtual machine. In this case, Tivoli Storage Productivity Center does not include disk space for this LPAR in the TOTAL disk space reports.

Disks that are assigned directly through the physical Fibre Channel adapter to the LPAR in a Virtual SCSI environment are listed as virtual disks in Tivoli Storage Productivity Center.

Login requirements

Use the **padmin** user ID when logging in to a Virtual I/O Server, **padmin** is the main administrator ID for Virtual I/O Servers. Upon login, a password change is required. There is no default password to remember.

The Virtual I/O Server does not support users logging in as root to install or configure Tivoli Storage Productivity Center agents. However, you can debug errors generated by agents as a root user. For example, you can run the `service.sh` script to gather debugging information.

Upon logging in to the Virtual I/O Server, you are placed into a restricted Korn shell. The restricted Korn shell works the same way as a regular Korn shell, except you cannot:

- Change the current working directory.
- Set the value of the SHELL, ENV, or PATH variable.
- Specify the path name of the command that contains a redirect output of a command with a `>`, `>|`, `<>`, or `>`.

As a result of these restrictions, you cannot run commands that are in locations not defined in your PATH variable. These restrictions prevent you from directly sending the output of the command to a file, instead requiring you to pipe the output to the **tee** command.

Installation requirements

You must use the **padmin** user ID to perform the following actions:

- Install the Storage Resource agent. See “Installing the Storage Resource agent on the Virtual I/O Server” on page 172.
- Run the commands to configure, start, and stop the agent. Go to the information center. Search for **managing the daemon**.
- For information about disk space requirements for the Storage Resource agent, see “Hardware requirements” on page 1 (Hardware requirements for the Tivoli Storage Productivity Center agents).

The installation requirements are:

- Tivoli Storage Productivity Center supports Virtual I/O Server version 1.5.2 or later.
- The Virtual I/O Server for POWER5 processor-based systems, must have a minimum of 512 MB of memory.
- The Virtual I/O Server for POWER6 processor-based systems, must have a minimum of 768 MB of memory.

General procedure to install and configure agents

To install and configure the agents on the Virtual I/O Server, complete the following steps:

1. Install Tivoli Storage Productivity Center.
2. Install and configure the agents on the Virtual I/O Server.
3. When you install the agents, they are automatically started
4. Run a discovery job for the agents on the Virtual I/O Server.
5. Collect information about a Virtual I/O Server by running probe, scan, and ping jobs.
6. View the storage information gathered by the monitoring jobs through the topology viewer and reports that you can generate through the Fabric Manager and Data Manager.

See the *IBM Tivoli Storage Productivity Center User's Guide* for more information about how to use the Tivoli Storage Productivity Center user interface to collect and view information gathered by agents on the Virtual I/O Servers.

Planning for SQL access to Tivoli Storage Productivity Center's views

Use this topic for planning to retrieve storage information that is stored in the Tivoli Storage Productivity Center database repository.

Tivoli Storage Productivity Center stores the information collected by data collection jobs in its DB2 database repository. The information in the database repository is organized into a set of tables. A view is a way of describing data that exists in one or more of these tables.

You can use the Structured Query Language (SQL) to retrieve the information from these views and then use that information in reports and applications outside of Tivoli Storage Productivity Center, such as Business Intelligence and Reporting Tools (BIRT), Microsoft Excel, and Crystal Reports. Other applications can also use these views to gather and import information that is collected by Tivoli Storage Productivity Center.

The views are created within a DB2 schema when you install the Tivoli Storage Productivity Center server. A schema defines the views and the fields within each view. The name of the schema containing the views for Tivoli Storage Productivity Center is TPCREPORT.

To view this schema using the DB2 Control Center, complete the following steps:

1. Start the DB2 Control Center.
2. Expand **All Databases > TPCDB** in the left-hand pane.
3. Click **Schemas**.
4. Click **TPCREPORT** in the content pane.

Before you can work with the data stored in views, keep in mind the following planning considerations:

- A view is a logical table that consists of data that is generated by a query. A view is based on an underlying set of base tables in the database repository, and the data in a view is determined by a SELECT statement that is run on the base tables.
- The views are available in Tivoli Storage Productivity Center version 4.1 or higher. If you want to retrieve data from the database repositories of previous versions of the product, you must first upgrade to version 4.1 or higher. See Chapter 4, "Upgrading and migrating Tivoli Storage Productivity Center," on page 341 for information about how to upgrade Tivoli Storage Productivity Center.
- You must run data collection jobs to collect data about the storage resources in your environment before you can retrieve data from views. The data collected by probes, scans, pings, switch performance monitors, and subsystem performance monitors populates those views.
- You must have proper authority to access and retrieve data from the views. The views can be accessed by the following users:
 - DB2 administrator
 - The DB2 user that Tivoli Storage Productivity Center uses to store and retrieve information in the database repository
 - A view-only user that uses the operating system credentials associated to DB2 (for example, db2admin/password).
- For information about performing the following tasks, go to the information center. Search for *Accessing views in the database repository*.
 - Use the Business Intelligence and Reporting Tools (BIRT) report designer to create a report based on the information available in Tivoli Storage Productivity Center's exposed views. See the following Web site for more information about BIRT: <http://www.eclipse.org/birt/phoenix/>.
 - View the report you created in BIRT through the Tivoli Common Reporting v1.2 feature of Tivoli Enterprise Portal. See the following web site for more information about Tivoli Common Reporting: <http://www.ibm.com/developerworks/spaces/tcr>.
- See the following page on Tivoli Storage Productivity Center's support web site for detailed information about the views provided with Tivoli Storage Productivity Center: <http://www.ibm.com/systems/support/storage/software/tpc>. Search for **Views Documentation**.

Setting up a view-only user on Windows

This topic provides information about how to set up a user with view-only privileges for the exposed views in the TPCREPORT schema. This procedure is for Windows only.

To set up a user with view-only privileges on Windows, complete these steps:

1. **Create a new user on Windows.** For example, create a user named TPCRPT.
 - a. Open **Administrative Tools > Computer Management** in the Control Panel.
 - b. Click **New User** on the **Action** menu.
 - c. Enter information about the new user and click **Create**.
 - d. Click **Close**.

2. **Add the new user to the DB2USERS group.** When you add the new user to the DB2USERS group it automatically provides that user with SQL connect authority only.
 - a. Open **User Accounts** in the Control Panel.
 - b. On the **Users** tab under **Users for this computer**, select the user account name and click **Properties**.
 - c. On the **Group Membership** tab, select the **DB2USERS** group and then **OK**.
3. **Use the DB2 Control Center to assign SELECT only authority to the new user.**
 - a. Start the DB2 Control Panel.
 - b. In the Control Center navigation tree, expand **TPCDB > User and Group Objects**.
 - c. Select the **DB Users** folder. A list of users appears in the content pane.
 - d. Select the new user that you created in step 1.
 - e. Select **Change...** from the **Selected** menu while the new user is highlighted.
 - f. Click the **View** tab. A list of views appears in the content pane. A green check mark in the columns next to each view indicates that the user can perform the actions represented by those columns. A red circle with a line through it indicates that the user cannot perform the actions represented by those columns.
 - g. Do not change the privileges for the views in the TPCREPORT schema. The default settings for the views in the TPCREPORT schema indicate that the new user cannot perform the following SQL commands against those views: **INSERT**, **UPDATE**, **DELETE**, **CONTROL**. Accept the default settings that restrict the new user from performing actions against the views.
 - h. Click the **Schema** tab. A list of schemas and the users level of authority for each appears in the content pane.
 - i. Make sure that the user does not have **CREATEIN**, **DROPIN**, or **ALTERIN** privileges. This check ensures that the user has read-only access to the exposed Tivoli Storage Productivity Center views.

Monitoring your DB2 environment

It is a good practice to monitor your DB2 environment to better understand what is happening inside your DB2 data server. DB2 9.7 includes a new monitoring infrastructure that you can access through new event monitors.

This infrastructure is a superior alternative to the existing system monitor, event monitors, snapshot commands, and snapshot SQL interfaces.

For more information about monitoring your DB2 environment, see <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.wn.doc/doc/c0055019.html>.

Planning for the Discovery Library Adapter

Use this topic to plan for using the Discovery Library Adapter.

IBM Tivoli Storage Productivity Center provides a Discovery Library Adapter which can be used in a multiple IBM Tivoli products environment to exchange data with other Tivoli products. The data gathered by IBM Tivoli Storage Productivity Center is put into a common data model format and written to a file using the Identification Markup Language. This file can be loaded into the change

and configuration management database (CCMDB) so that other Tivoli products such as Tivoli Application Dependency Discovery Manager or Tivoli Business Service Manager can use that data.

The Discovery Library Adapter also enables Tivoli Application Dependency Discovery Manager to be able to perform a launch in context operation to the IBM Tivoli Storage Productivity Center server. The following IBM Tivoli Storage Productivity Center panels are accessible:

- Topology view for computers
- Topology view for fabrics
- Topology view for switches
- Topology view for storage systems

Tivoli Storage Productivity Center Monitoring Agent

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

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

For performance reasons, install the Tivoli Storage Productivity Center server on a separate system from the IBM Tivoli Monitoring server. You will be installing the Tivoli Storage Productivity Center Monitoring Agent in the Tivoli Storage Productivity Center environment. You will also be installing the Tivoli Storage Productivity Center Monitoring Agent support files in the IBM Tivoli Monitoring server environment.

Before using the Tivoli Storage Productivity Center Monitoring Agent, you must have IBM Tivoli Enterprise Monitoring Server installed. For information about the IBM Tivoli Enterprise Monitoring Server, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.2/welcome.htm.

For information about installing IBM Tivoli Monitoring, see the Quick Start Guide at http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.2fp2/itm622fp2_qsg_en.htm.

This agent is on the installation media disk 1, part 4.

Hardware and software requirements

You must install Tivoli Agent Management Services version 6.2 or later on Windows, AIX or Linux. The following Tivoli Monitoring components are required:

Tivoli Enterprise Monitoring Agent Framework - TEMA

- Tivoli Enterprise Monitoring Agent Framework
- Warehouse Proxy
- Summarization and Pruning Agent
- Tivoli Enterprise Services User Interface Extensions

Tivoli Enterprise Monitoring Server - TEMS

- Tivoli Enterprise Monitoring Server
- Warehouse Proxy
- Summarization and Pruning Agent
- Monitoring Agent support for the corresponding operating system on which the agent is installed

Tivoli Enterprise Portal Server - TEPS

- Tivoli Enterprise Portal Server
- Warehouse proxy
- Monitoring Agent support for the operating system on which the agent is installed
- Summarization and Pruning Agent

Tivoli Enterprise Desktop Client - TEPD

- Tivoli Enterprise Desktop Client
- Warehouse Proxy
- Monitoring Agent support for the operating system on which the agent is installed
- Summarization and Pruning Agent

Tivoli Enterprise Portal Browser Client

IBM Eclipse Help Server

Note: IBM Tivoli Monitor does not support the Tivoli Enterprise Portal Desktop Client (TEPDC) on AIX.

The following disk space is required:

Memory

Approximately 100 MB of RAM per agent instance.

Disk space

Approximately 300 MB of disk space for the IBM Tivoli Storage Productivity Center Monitoring Agent. For information about disk space required for historical data, see the *IBM Tivoli Monitoring Administrator's Guide* at http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.2/itm_admin.htm.

To install the IBM Tivoli Storage Productivity Center Monitoring Agent, you must have Administrator authority on Windows and root authority on AIX and Linux.

For information about installing the IBM Tivoli Enterprise Monitoring Server, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.2/itm_install86.htm.

Tivoli Agent Management Services

The Tivoli Agent Management Services extension of IBM Tivoli Monitoring monitors the health and availability of Storage Resource agents that are run as daemon processes.

The purpose of Agent Management Services is to provide customers with a central view of various Tivoli agents, and to perform the following functions:

- Monitors the health of Storage Resource agents
- Starts, stops, and manages agents

Through these functions, Tivoli Agent Management Services provides the following benefits for agent management:

- Adds a dimension of fault tolerance to agents by keeping a deployed Storage Resource agent up and running.
- Reduces the mean time to recovery of failed agents to seconds, limiting the number of situations that need to be configured to those reported as unrecoverable faults.
- Reduces the number of events flowing to your event managers to those reported by Tivoli Agent Management Services.

For information about Tivoli Monitoring, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.2fp2/itm622fp2_qsg_en.htm .

A Tivoli Monitoring Agent for OS must be installed on every computer that you want to monitor with a Storage Resource agent. The Tivoli Monitoring Agent for OS monitors all agents that have CAP files located in the CAP file directories. The Tivoli Agent Management Services can stop, start, and check the status of agents listed in the CAP file. The CAP file is created when a daemon Storage Resource agent is installed on a system.

For more information about the Tivoli Agent Management Services, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.2fp2/agentmgmtsvcs_intro_tema.htm .

For information about Tivoli Monitoring supported platforms, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.2fp2/ch2.4.htm#ch2.4.

On Windows, the status, start, and stop functions are handled through the Windows services.

The CAP file is installed in the following directories:

Windows

`%ALLUSERSPROFILE%\Application Data\IBM\CAP`

For example:

Windows 2008: `C:\Program Data\IBM\CAP`

Windows 2003: `C:\Documents and Settings\All Users\Application Data\IBM\CAP`

Non-Windows

`/opt/IBM/CAP`

CAP files are not configurable. You cannot change the location of these directories.

If the creation of the CAP file fails, you receive an error message in the installation log file and trace log file. If the CAP is not created, the lack of a CAP file does not affect the function of the Storage Resource agent. You can reinstall the agent with the **-force** option to create the CAP file.

For more information about the CAP file or Tivoli Agent Management Services, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.2fp2/agentmgmtsvcs_intro_tema.htm .

Planning to install Tivoli Storage Productivity Center

Install Tivoli Storage Productivity Center on the system to manage storage resources, storage system management, fabric management, replication management, alerting, and analytical functions. Install Storage Resource agents on each computer to collect information about the hosts, applications, storage systems, and SAN fabrics. The Storage Resource agents transfer that data to the Tivoli Storage Productivity Center server.

Hardware and software requirements

Tivoli Storage Productivity Center can require a large amount of memory, disk space, network bandwidth, and processor resources.

Tip: Use a dedicated server that is not running other applications when you install Tivoli Storage Productivity Center. A dedicated server improves performance.

For information about hardware and software requirements, see “Planning for installation” on page 1. See the following website for any changes to product and platform support: <https://www.ibm.com/support/docview.wss?uid=swg21386446>.

Read the *Readme for IBM Tivoli Storage Productivity Center Version 5.1* on http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center. The readme file contains last-minute information that could not be included in the documentation.

Choosing the installation method

You can install Tivoli Storage Productivity Center by using the installation wizard or the command-line interface in silent mode. When you use silent mode, you run a command and provide the installation values in a response file.

To use the installation wizard on operating systems such as UNIX and Linux, you must have X Window System software. For more information about the installation wizard, see “Installing Tivoli Storage Productivity Center in a single-server environment by using the wizard” on page 141.

The installation wizard helps you install Tivoli Storage Productivity Center with minimal user interaction. The silent-mode installation is useful if your system is running from a terminal that cannot display graphics. For more information about using the installation wizard, see “Installing Tivoli Storage Productivity Center in a single-server environment by using the wizard” on page 141. For more information about installing the product in silent mode, see “Installing Tivoli Storage Productivity Center in a single-server environment by using silent mode” on page 142 or “Installing Tivoli Storage Productivity Center with a remote database by using silent mode” on page 154.

The installation procedures use the following directory to indicate where Tivoli Storage Productivity Center is installed:

TPC_installation_directory

where the default directory for Windows operating systems is C:\Program Files\IBM\TPC, and the default directory for AIX or Linux operating systems is /opt/IBM/TPC.

In Tivoli Storage Productivity Center, servers can include any of the following types:

- Data server
- Device server
- Replication server

Choosing an interface

Tivoli Storage Productivity Center has the following graphical user interfaces (GUIs):

- web-based GUI, which provides simplified monitoring and troubleshooting capabilities and helps you to complete the following tasks:
 - Find information about the capacity and health of the resources managed by Tivoli Storage Productivity Center.
 - Display information about the web server, such as the internal and related resources for the server.In the web-based GUI, you can only acknowledge and delete alerts.
- stand-alone GUI, which helps you add devices, manage devices, provision, schedule jobs, and create alerts.

For more information about the GUIs, go to the Tivoli Storage Productivity Center Information Center and search for *Interfaces to Tivoli Storage Productivity Center*.

Languages

When you install Tivoli Storage Productivity Center, you must select one of the following languages:

- English
- Czech
- French
- German
- Hungarian
- Italian
- Japanese
- Korean
- Polish
- Brazilian Portuguese
- Russian
- Spanish
- Chinese (Simplified)
- Chinese (Traditional)

Messages, online help, and text are displayed in the language that you select. You are not required to install a language pack after installation. When you use Tivoli Storage Productivity Center, the language that is displayed is the language setting of the operating system.

If you are installing Tivoli Storage Productivity Center on a Windows operating system and want to authenticate users against an instance of Microsoft Active Directory, see “Configuring Tivoli Storage Productivity Center for a Windows domain” on page 310.

If you are installing Tivoli Storage Productivity Center and want to use an existing instance of Tivoli Integrated Portal 2.2, see “Tivoli Integrated Portal and Tivoli Common Reporting” on page 135.

Installation restrictions

Before you install Tivoli Storage Productivity Center, ensure that you are familiar with following issues:

Installation images on AIX

When you extract installation images, use the GNU tar program instead of the AIX tar program. The AIX tar program might truncate long file names, which can cause installation errors in Tivoli Storage Productivity Center and Tivoli Integrated Portal.

Go to <http://www.ibm.com/systems/power/software/aix/linux/toolbox/alpha.html> and install GNU 1.14 or later. You need to specify this program as the default tar program in the PATH environment variable.

Cygwin software

When you install Tivoli Storage Productivity Center on a Windows operating system and have Cygwin software on the same server, do not include Cygwin software as a path-related environment variable, for example, the PATH variable.

Tivoli Common Reporting requirements

When you install Tivoli Storage Productivity Center on the Linux operating system, there are other packages that you may need to install. See “Software requirements for operating systems” on page 4 for more information about the packages that you need to install.

DB2 on AIX or Linux

When you install DB2 on operating systems such as AIX or Linux, by default DB2 is not configured to start when the server restarts. To ensure that DB2 starts when the server starts, after you install DB2 on AIX or Linux, restart DB2, and run the following command:

```
DB2_installation_directory/bin/db2iauto -on instance_name
```

where *DB2_installation_directory* is the location of DB2 and *instance_name* is db2inst1.

Related tasks:

“Installing Tivoli Storage Productivity Center with a remote database by using the installation wizard” on page 150

You can install Tivoli Storage Productivity Center in a multiple-server environment by using the installation wizard.

“Installing Tivoli Integrated Portal and Tivoli Common Reporting by using the installation wizard” on page 152

You can install Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI in a multiple-server environment using the installation wizard.

Tivoli Storage Productivity Center components

You can install Tivoli Storage Productivity Center in single- or multiple-server environments. In a single-server environment, all components are installed on one server.

When you install Tivoli Storage Productivity Center, the following components are installed:

- Database repository
- Data server
- Device server
- Replication server
- Stand-alone graphical user interface (GUI)
- Web-based graphical user interface (GUI)
- Command-line interface (CLI)
- Storage Resource agent
- IBM Tivoli Integrated Portal
- IBM Tivoli Common Reporting

This table contains configuration scenarios, which are valid for both installation types. You can use the information in this table to help you install Tivoli Storage Productivity Center:

Table 26. Tivoli Storage Productivity Center configurations

If server A includes...	If server B includes...	Install in this order...
Data server, Device server, Replication server, database repository, web-based GUI, stand-alone GUI, CLI, Storage Resource agent, Tivoli Integrated Portal, Tivoli Common Reporting (single-server environment)	This scenario applies to a single-server environment. There is no server B.	This scenario applies to a single-server environment. There is no server B.

Table 26. Tivoli Storage Productivity Center configurations (continued)

If server A includes...	If server B includes...	Install in this order...
Data server, Device server, Replication server, web-based GUI, stand-alone GUI, CLI, Storage Resource agent, Tivoli Integrated Portal, Tivoli Common Reporting (multiple-server environment)	Database repository	<ol style="list-style-type: none"> 1. Install the DB2 server on server B. 2. Run the Tivoli Storage Productivity Center installation program on server B to install the database repository. 3. Install the DB2 client on server A. 4. Run the Tivoli Storage Productivity Center installation program on server A to install the remainder of the Tivoli Storage Productivity Center components.
Data server, Device server, Replication server, database repository, stand-alone GUI, CLI, and Storage Resource agent (multiple-server environment)	Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI	<ol style="list-style-type: none"> 1. Install the DB2 server on server A. 2. Run the Tivoli Storage Productivity Center installation program on server A to install all the components except for Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI. 3. Install the DB2 client on server B. 4. Run the Tivoli Storage Productivity Center installation program on server B to install Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI.

Tivoli Storage Productivity Center in an IPv4 or IPv6 environment

Tivoli Storage Productivity Center can use both IPv4 and IPv6 addresses for communication between its components.

If you have a system that is configured for dual-stack networking (with both IPv4 and IPv6 addresses), Tivoli Storage Productivity Center defaults to IPv4 addressing.

For information about using IPv4 and IPv6, see “Planning for Internet Protocol Version 6” on page 46.

Chapter 2. Installing Tivoli Storage Productivity Center

Before you install IBM Tivoli Storage Productivity Center, you must install IBM DB2 because DB2 is required for the Tivoli Storage Productivity Center database repository.

Installing Tivoli Storage Productivity Center

Tivoli Storage Productivity Center provides an installation program that guides you through the installation process. You can use the installation wizard or the command line in silent mode.

Tivoli Storage Productivity Center installation

You can install Tivoli Storage Productivity Center by using the installation wizard or the command line in silent mode.

Overview

When you install Tivoli Storage Productivity Center, Tivoli Integrated Portal and Tivoli Common Reporting are automatically installed. You can specify a location to install the new instance of Tivoli Integrated Portal. You can reuse an existing Tivoli Integrated Portal instance, if its authentication mechanism is using the operating system repository or the Federated repositories.

Fully qualified host names

Some systems might be configured to return a short host name such as `server22`, instead of a fully qualified host name such as `server22.myorg.mycompany.com`. Tivoli Storage Productivity Center requires fully qualified host names, so you must install the software on a computer that has a fully qualified host name.

Installation method

You can install Tivoli Storage Productivity Center by using the installation wizard or the command line in silent mode. In silent mode, a command is provided with the values in a response file.

Use the installation wizard to install Tivoli Storage Productivity Center with minimal user interaction. Use the silent-mode installation if your system is running from a terminal that cannot display graphics.

Important: On UNIX operating systems (for example, AIX or Linux) you must have X Windows support to display the installation wizard graphical user interface (GUI).

Tivoli Integrated Portal and Tivoli Common Reporting

You can install a new instance of Tivoli Integrated Portal 2.2 and Tivoli Common Reporting 2.1.1, or you can reuse the existing instances.

The installation program detects the existing instances of Tivoli Integrated Portal. The unsupported instances are disabled or grayed out. A Tivoli Integrated Portal instance is shown as unsupported if any of the following is true:

- Tivoli Integrated Portal version is lower than Version 2.2.
- Tivoli Common Reporting version is lower than Version 2.1.1.
- Tivoli Integrated Portal is a 64-bit instance and does not have 64-bit Tivoli Common Reporting and 64-bit Embedded Security Services components that are installed on it.

During the installation, if Tivoli Integrated Portal Version 2.2 is not detected, a new instance of Tivoli Integrated Portal Version 2.2 is installed. Tivoli Common Reporting 2.1.1, Embedded Security Services, and the web-based GUI components are also installed. During this process, you must enter a Tivoli Common Reporting port. Tivoli Storage Productivity Center always installs 32-bit versions of these components.

If an instance of Tivoli Common Reporting 2.1.1 and Tivoli Integrated Portal 2.2 are detected, those instances are used by Tivoli Storage Productivity Center. When you use an existing instance of Tivoli Integrated Portal, you must enter the same administrative user name and password that you used when you installed the Tivoli Integrated Portal instance.

When you uninstall Tivoli Integrated Portal 1.1 and an instance of Tivoli Integrated Portal 2.2 does not exist, the ports, administrative user name, and password that were used in Version 1.1 are reused for Version 2.2.

Important:

- Tivoli Storage Productivity Center can only detect the Tivoli Integrated Portal instances that were installed by a root user. If you are planning to share a Tivoli Integrated Portal instance with Tivoli Storage Productivity Center and other products, such as IBM® Tivoli® Business Service Manager 6.1, you must also install those products as a root user.
- Tivoli Storage Productivity Center always installs a 32-bit version of Tivoli Common Reporting 2.1.1 and Embedded Security Services. As a result, Tivoli Storage Productivity Center can only reuse an existing 64-bit Tivoli Integrated Portal instance that already has the 64-bit Tivoli Common Reporting and the 64-bit Embedded Security Services components that are installed on that instance.

Installation images

The following installation images are included with Tivoli Storage Productivity Center:

- Tivoli Storage Productivity Center for AIX
- Tivoli Storage Productivity Center for Linux
- Tivoli Storage Productivity Center for Windows
- Tivoli Storage Productivity Center Storage Resource agent

The Tivoli Storage Productivity Center installation image has the following disks:

- **Disk 1**, which comprises the following parts:
 - **Disk 1, Part 1**, the base Tivoli Storage Productivity Center installation program.

- **Disk 1, Part 2**, the Tivoli Storage Productivity Center Device and Replication Server installation programs.
- **Disk 1, Part 3**, the base Tivoli Integrated Portal installation program.
- **Disk 1, Part 4**, the Tivoli Common Reporting and Embedded Security Services installation programs.
- **Disk 1, Part 5**, the Tivoli Integrated Portal Temporary Fix (PTF) installation program.

Download each of the **Disk 1** compressed image files to the same directory (for example, c:\tpcdownload). After you download the files, extract each file into the same target directory (for example, c:\tpcinstall) in order (for example, first **Disk 1- part 1**, then **Disk1 - part 2**, and so on).

While you extract the files, if you are prompted to merge (or replace) folders or files with the same name, click **Merge/Replace**.

- **Disk 2**, which comprises the following:
 - **Disk 2, itm_unix1 and itm_unix2**, the Tivoli Storage Productivity Center Monitoring Agent installation program for IBM Tivoli Monitoring on the UNIX operating system.
 - **Disk 2, itm_windows**, the Tivoli Storage Productivity Center Monitoring Agent installation program for IBM Tivoli Monitoring on the Windows operating system.
 - **Disk 2, sde_unix**, which is the IBM System Discovery Engine on the UNIX operating system.

Tivoli Storage Productivity Center installation package

The following software is included with Tivoli Storage Productivity Center:

- DB2 9.7 for Windows (32-bit)
- DB2 9.7 for Windows (64-bit)
- DB2 9.7 for Linux (64-bit)
- DB2 9.7 for AIX (64-bit)

Tivoli Monitoring server and DB2 installation images are included with Tivoli Storage Productivity Center.

The Tivoli Monitoring server that is included with Tivoli Storage Productivity Center has a limited license and that can be used to manage or monitor your Tivoli Storage Productivity Center environment. The product license defines the specific entitlements.

The following software is included with the Tivoli Monitoring server installation images:

- IBM Tivoli Monitoring 6.2.2.2 Server for Windows
- IBM Tivoli Monitoring 6.2.2.2 Server for Linux
- IBM Tivoli Monitoring 6.2.2.2 Server for UNIX
- IBM Tivoli Monitoring 6.2.2.2 Agent (this is the Tivoli Monitoring Agent for the operating system)
- IBM Tivoli Monitoring 6.2.2.2 National Language Support

For information about installing these components, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.2fp2/installation.htm.

For information about the supported platforms for Tivoli Monitoring, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.2fp2/ch2.4.htm#ch2.4.

Starting the installation program

This topic describes how to start the Tivoli Storage Productivity Center installation programs.

Installation images

If you are using electronic images, you must select an installation image. For a list of installation images, see “Tivoli Storage Productivity Center installation” on page 135.

Storage Resource Agent

Contains the files to perform local Storage Resource agent installations. You must download the file for the platform on which you want the agent to reside.

- The **Storage Resource Agent** location and operating system file name is the same as the **disk 1** image.
- The Storage Resource agent image is in *DVD/data/sra/operating_system_name*.

The operating systems supported are listed in the following table.

Table 27. Operating system for Storage Resource agents

Operating system	Operating system name
AIX	aix_power
HP-UX	hp-ux_itanium
Linux x86	linux_ix86
Linux for Power Systems Servers	linux_power
Linux s390	linux_s390
Oracle Solaris	solaris_spark
Windows	windows

Note:

- If you are using Tivoli Storage Productivity Center electronic installation images that are downloaded to a Windows directory, and the directory has spaces or unusual special characters in the name, Tivoli Storage Productivity Center does not install correctly.

For example, if you have a directory name:

```
C:\tpc 51 edition\disk1
```

Tivoli Storage Productivity Center does not install correctly.

When you remove the spaces, for example:

```
C:\tpc51edition\disk1
```

Tivoli Storage Productivity Center installs correctly.

- If you are deploying the Storage Resource agent on the UNIX or Linux operating system, you must use root as the user name.

Starting the installation program on the Windows operating system

To start the Tivoli Storage Productivity Center installation program using an electronic image, complete the following steps:

1. Download the image into a directory.
2. Extract the file.
3. At the command prompt, enter `setup.bat` from the source installation directory.

To start the Tivoli Storage Productivity Center installation program using the disk 1 DVD, complete the following steps:

1. Open Windows Explorer and navigate to the Tivoli Storage Productivity Center DVD drive or directory.
2. At the command prompt, enter `cd DVD_drive\disk1` and navigate to the Tivoli Storage Productivity Center DVD drive or directory.
3. Double-click **setup.bat**.

Starting the installation program on the Linux operating system

To start the Tivoli Storage Productivity Center installation program from the electronic image (disk 1), complete the following steps:

1. Enter the following to create a directory:
`mkdir /tpcinst`
2. Enter `cd tpcinst`.
3. Enter `ftp filename.tar` to transfer the Tivoli Storage Productivity Center installation tar file into a directory.
4. Enter this command:
`tar -xvf filename.tar`
5. Set up your shell environment to point to the instance where the database repository is installed. Source the `db2profile` for the wanted instance. For example, if the DB2 instance is **db2inst1**, source the `db2profile` by entering:
`./home/db2inst1/sqllib/db2profile`

Remember: There is a space between `.` and `/home`.

6. Start the installation program by running the `./setup.bin` command.

To start the Tivoli Storage Productivity Center installation program from the disk 1 DVD, complete the following steps:

1. At the command prompt, enter `mkdir /cdrom` to create a mount point.
2. Insert the DVD into the DVD drive.
3. Mount the DVD file system at the appropriate mount point. For example, to mount a DVD into a DVD drive known as `/dev/cd0` at mount point `/cdrom`, enter the following command:
`mount -o ro /dev/cdrom /cdrom`
4. Enter `cd /cdrom` to navigate to the directory where the DVD is mounted.
5. Set up your shell environment to point to the instance where the database repository is installed.
6. Source the `db2profile` for the appropriate instance. For example, if the DB2 instance is **db2inst1**, source the file by entering:
`./home/db2inst1/sqllib/db2profile`

Remember: There is a space between `.` and `/home`.

7. Start the installation program from the *disk1* directory (for example, *TPC_installation_directory/disk1*).
8. Enter `/setup.bin`. If `setup.bin` exits prematurely because the media (DVD) is not ready, run the `/setup.bin` command again until the installation wizard is displayed.

Starting the installation program on the AIX operating system

To start the Tivoli Storage Productivity Center installation program from the electronic image (disk 1), complete the following steps:

1. Enter the following to create a directory:


```
mkdir /tpcinst
```
2. Enter `cd tpcinst`.
3. Enter `ftp filename.tar` to transfer the Tivoli Storage Productivity Center installation tar file into a directory.
4. Enter this command:


```
tar -xvf filename.tar
```
5. Set up your shell environment to point to the instance where the database repository is installed. Source the `db2profile` for the wanted instance. For example, if the DB2 instance is **db2inst1**, source the file by entering:


```
./home/db2inst1/sqllib/db2profile
```

Remember: There is a space between `.` and `/home`.

6. Start the installation program by running the `./setup.bin` command.

To start the Tivoli Storage Productivity Center installation program from the disk 1 DVD, complete the following steps:

1. Insert the DVD into the DVD drive and mount the DVD file system at the wanted mount point. On the AIX operating system, you can use the `crfs` command to add an entry to `/etc/filesystems` for the mount point. Enter the following commands:

```
/usr/sbin/crfs -v cdrfs -p ro -d'cd0' -m'/cdrom' -A'no'
mount /cdrom
```

The `crfs` command must be run once for a mount point. After that you can use **mount** and **umount** for each DVD you insert, for example, **mount /cdrom** and **umount /cdrom**.

2. Enter `cd /cdrom` to navigate to the directory where the DVD is mounted.
3. Set up your shell environment to point to the instance where the database repository is installed.
4. Source the `db2profile` for the appropriate instance.

For example, if the DB2 instance is **db2inst1**, source the file by entering:

```
./home/db2inst1/sqllib/db2profile
```

Remember: There is a space between `.` and `/home`.

5. Start the installation program from the *disk1* directory (for example, *TPC_installation_directory/disk1*).
6. Run the `./setup.bin` command. If the `setup.bin` command exits prematurely because the media (DVD) is not ready, run the `./setup.bin` command again until the installation wizard is displayed.

Installing Tivoli Storage Productivity Center in a single-server environment

Install Tivoli Storage Productivity Center by using the installation wizard or the command line in silent mode.

Installing Tivoli Storage Productivity Center in a single-server environment is a simple process and can be completed successfully by most Tivoli Storage Productivity Center customers.

Related tasks:

“Installing Tivoli Storage Productivity Center in a single-server environment by using the wizard”

You can use the installation wizard to install Tivoli Storage Productivity Center in a single-server environment.

“Installing Tivoli Storage Productivity Center in a single-server environment by using silent mode” on page 142

You can install Tivoli Storage Productivity Center using the command line in silent mode.

Installing Tivoli Storage Productivity Center in a single-server environment by using the wizard

You can use the installation wizard to install Tivoli Storage Productivity Center in a single-server environment.

Before you start the Tivoli Storage Productivity Center installation wizard on AIX or Linux operating system, you must source the user profile (db2profile) for the instance owner of the DB2 database. For example:

```
. /home/db2inst1/sqllib/db2profile
```

When you install Tivoli Storage Productivity Center, a database called TPCDB is created. In this repository, the TPC database schema is created, which describes the structure of the database repository. If you already have the TPCDB database on your system, Tivoli Storage Productivity Center assigns a new default database repository name with a number as a suffix for example, TPCDB1.

To install Tivoli Storage Productivity Center in a single-server environment by using the wizard, complete the following steps:

1. Log on to the server where you want to install Tivoli Storage Productivity Center. Ensure that you have the required administrator privileges to install applications on the server.
2. Install DB2 Enterprise Server Edition. For more information about installing DB2, see “Installing DB2” on page 180.

If you are installing DB2 on a Windows operating system, log out and then log on again before you install Tivoli Storage Productivity Center.

3. Start the Tivoli Storage Productivity Center installation wizard.
4. On the Choose Installation Location and Type page, select **Single server**.
5. Follow the instructions and complete the installation by using the wizard.

If an error occurred during the installation process, review the installation log files to find details about the error.

Related tasks:

“Installing DB2” on page 180

This topic describes how to install DB2 Version 9.7 for operating systems such as Linux, AIX, or Windows.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Starting the installation program” on page 138

This topic describes how to start the Tivoli Storage Productivity Center installation programs.

“Verifying the installation” on page 163

After you install Tivoli Storage Productivity Center, you can verify whether the installation was successful.

“Reviewing the log files to resolve installation issues” on page 165

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

Installing Tivoli Storage Productivity Center in a single-server environment by using silent mode

You can install Tivoli Storage Productivity Center using the command line in silent mode.

Before you start the Tivoli Storage Productivity Center installation program on AIX or Linux operating system, you must source the user profile (db2profile) for the instance owner of the DB2 database. For example:

```
. /home/db2inst1/sqllib/db2profile
```

To install Tivoli Storage Productivity Center in a single-server environment by using silent mode, complete the following steps:

1. Log on to the server where you want to install Tivoli Storage Productivity Center. Ensure that you have the required administrator privileges to install applications on the server.
2. Install DB2 Enterprise Server Edition. For more information about installing DB2, see “Installing DB2” on page 180.

If you are installing on a Windows operating system, log out and then log on again before you install Tivoli Storage Productivity Center.

3. Edit and save the appropriate response file. See “Editing the response file” on page 143.

Restriction: If you install Tivoli Storage Productivity Center by using silent mode, and are not using the response file, you must set `-DLICENSE_ACCEPTED=true` to accept the license agreement. For example:

```
setup.bat/bin -i silent -DLICENSE_ACCEPTED=true
```

4. Run the installation program.

- For Windows operating systems, run the following command:
`setup.bat -l language -i silent -f absolute_path_to_response_file`
where *language* can be one of the following values:
 - Czech - cs

- English - en
- French - fr
- German - de
- Hungarian - hu
- Italian - it
- Japanese - ja
- Korean - ko
- Polish - pl
- Brazilian Portuguese - pt_BR
- Russian - ru
- Spanish - es
- Chinese (Simplified) - zh_CN
- Chinese (Traditional) - zh_TW

`absolute_path_to_response_file` is the absolute path to the response file. For example, the following command specifies the language and the path:

```
setup.bat -l fr -i silent -f
c:\TPC51\silent_SingleServerTypical.properties
```

- For AIX or Linux operating systems, run the following command:
`./setup.bin -l language -i silent -f /absolute_path_to_response_file`

For example, the following command specifies the language and the path:
`./setup.bin -l de -i silent -f /TPC51/silent_SingleServerTypical.properties`

If you do not specify a language or if the language is not specified correctly, the installation program uses the English language as the default.

5. Optional: Monitor the progress of the installation.

- To view the progress of the silent mode installation on Windows operating systems, check the trace log in the following path:

```
TPC_installation_directory\logs\traceTPCInstall.log
```

- To view the progress of the silent mode installation on AIX or Linux operating systems, check the trace log in the following path:

```
TPC_installation_directory/logs/traceTPCInstall.log
```

Editing the response file

You must edit and save the appropriate response file when you install Tivoli Storage Productivity Center in silent mode. The silent mode install option for the response file is `-f absolute_path_to_response_file`. For example, in the Windows operating system, you enter `-f C:\installimage\silent_SingleServerTypical.properties`.

Use the following response files during a silent mode installation:

silent_SingleServerTypical.properties

This file specifies that all the Tivoli Storage Productivity Center components are installed on one server. There are no customization options.

silent_SingleServerCustom.properties

This file specifies that all the Tivoli Storage Productivity Center components are installed on one server with the following customization options:

- Change the user name and password to install the database repository.

- Change the name of the database repository.
- Change the paths for the database repository.
- Change the path for the database repository log.
- Change the ports for the Data server, Device server, Replication server, and Storage Resource agent.
- Change the user name and password to install Tivoli Integrated Portal
- Change the installation location for Tivoli Integrated Portal
- Change the ports for Tivoli Integrated Portal and Tivoli Common Reporting
- Install Tivoli Storage Productivity Center to use an existing instance of Tivoli Integrated Portal.
- Provide installation location and authentication information for that existing instance of Tivoli Integrated Portal.

silent_MultipleServer.properties

This file specifies the information that is required for a multiple-server environment. You can have these configurations:

- The database repository on a remote server and the other Tivoli Storage Productivity Center components on a local server.
- Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI components on a remote server and the other Tivoli Storage Productivity Center components on a local server.

Common response file parameters

These installation parameters are valid for all the response files.

CHOSEN_INSTALL_TYPE="option"

Specifies the installation type. Not all options are available in all the response files. Table 28 lists the options you can specify.

Table 28. Options for the CHOSEN_INSTALL_TYPE parameter

Properties file	Valid options
silent_SingleServerTypical	<ul style="list-style-type: none"> • Single Server Install • License Upgrade
silent_SingleServerCustom	<ul style="list-style-type: none"> • Single Server Install • License Upgrade
silent_MultipleServer	<ul style="list-style-type: none"> • Multiple Server Install • License Upgrade

For example:

```
CHOSEN_INSTALL_TYPE="Single Server Install"
```

LICENSE_ACCEPTED=false

Specifies whether the user accepts the Tivoli Storage Productivity Center license agreement. The options are:

true

The user accepts all the terms and conditions of the Tivoli Storage Productivity Center license agreement.

false or any other value

The user does not accept the Tivoli Storage Productivity Center license agreement. The installation program will exit.

USER_INSTALL_DIR=option

Specifies the location where you want to install Tivoli Storage Productivity Center.

On Windows operating systems, you must use double backslashes. For example:

```
USER_INSTALL_DIR=C:\\Program Files\\IBM\\TPC
```

The default value for Windows operating systems is C:\\Program Files\\IBM\\TPC.

The default value for AIX or Linux operating systems is /opt/IBM/TPC.

varSrvName=option

Specifies the fully qualified host name of the server.

If your system is configured for dual stack networking (with both IPv4 and IPv6 addresses), Tivoli Storage Productivity Center defaults to IPv4 addressing.

Ensure that DNS is configured correctly on your server and verify that these files exist on your target server:

- For Windows operating systems, C:\\Windows\\system32\\drivers\\etc\\hosts.
- For AIX or Linux operating systems, /etc/hosts.

These files must have an entry similar to the following example:

```
myserver.example.mycompany.com myservers
```

varTPCPortRangeSP=port

Specifies the first port in a range of ports. The default beginning port is 9549. These ports are used by the Tivoli Storage Productivity Center servers, Storage Resource agent, Tivoli Integrated Portal, and Tivoli Common Reporting.

varCommonUsrID=user_name

Specifies the user name to install Tivoli Storage Productivity Center. The default user name for Windows operating system is db2admin. The default user name for AIX or Linux operating system is db2inst1.

The user name that you select must have the following operating system privileges:

- For Windows operating systems, the user name must be in the administrators group.
- For Linux operating systems, the user name must be in the root group.
- For AIX operating systems, the user name must be in the system group.

If you are using the common user name to also install the database repository, the user name must have these DB2 privileges:

- For Windows operating systems, the user name must be in the DB2ADMNS group.
- For AIX or Linux operating systems, the user name must be in the db2iadm1 group.

varCommonUsrPw=password

Specifies the password for the common user name.

varFullRollback=*value*

Specifies whether a full rollback or partial rollback is done when an installation failure occurs. The value is:

0 Specifies that a partial rollback occurs. This is the default.

1 Specifies that a full rollback occurs.

varUseLicenseKeyOnImage=*value*

Specifies whether to use the license key file that is present on the installation image. The value can be:

0 Uses the license key that is present on the installation image (default value)

1 Provides the location of the license key file

Depending on the license package you purchased, one of these license files is included with Tivoli Storage Productivity Center in the root installation directory and is on disk 1, part 1 of the DVD or electronic image:

nodelock

Tivoli Storage Productivity Center license

nodelock.SEL

Tivoli Storage Productivity Center Select license

nodelock.AE

Tivoli Storage Productivity Center Advanced license

nodelock.EE

Tivoli Storage Productivity Center product included with IBM Systems Director

varLicenseKeyFile=*location_of_license_key_file*

Specifies the path for the license key file. This value must be set if the **varUseLicenseKeyOnImage=1** parameter is specified.

On the Windows operating system, you must use double backslashes. For example:

```
varLicenseKeyFile=C:\\license\\key\\nodelock.SEL
```

Response file parameter to customize Tivoli Storage Productivity Center ports

This parameter is used for the `silent_SingleServerCustom.properties` and `silent_MultipleServer.properties` response files to configure the ports.

varTPCServersPortRangeSP=*port*

Specifies the first port in a range of ports for the Tivoli Storage Productivity Center servers, and Storage Resource agent. The default beginning port is 9549.

This value overrides the value specified in the **varTPCPortRangeSP** parameter.

Response file parameters for advanced customization of the database repository

These parameters are used for the `silent_SingleServerCustom.properties` and `silent_MultipleServer.properties` response files to configure the database repository.

varDBAdmUsr=*user_name*

Specifies the user name that you must use to install the Tivoli Storage Productivity Center database. The user who logs in with this user name must be the owner of the DB2 instance where the database is created.

This user name must have these DB2 privileges:

- For Windows operating systems, the user name must be in the DB2ADMNS group. The default user name is db2admin.
- For AIX or Linux operating systems, the user name must be in the db2iadm1 group. The default user name is db2inst1.

varDBAdmPW=*password*

Specifies the password that is associated with the DB2 user name.

varDBName=*database_name*

Specifies the name of the Tivoli Storage Productivity Center database. The default database name is TPCDB.

Specify the database name using the following rules:

- The name can contain the following characters:
 - a - z
 - A - Z
 - 0 - 9
- The name must be 1 - 8 characters long.

Restriction: The database name cannot have the following conditions:

- Contain a space or a blank.
- Begin with SYS, DBM, or IBM.
- Be the name of an existing database.

varDBPath=*database_location, database_location, ...*

Specifies the absolute paths for the database.

The database path must follow these rules:

- The path cannot be longer than 242 bytes.
- The path must be an absolute path and not a relative path, for example, C:\

You can specify a maximum of 10 database paths and separate them with commas. For example, in the Windows operating system, you would specify C:,D:. In the AIX or Linux operating system, you would specify /home/db2inst1,/testPath.

- For example, on the Windows operating system, if the database is on the C drive, you can add a D drive, an E drive, and so on.

On Windows operating systems, you must use double backslashes. For example:

```
varDBPath=C:\\DB2\\TPCDB,D:\\TPCDB,E:\\TPCDB
```

- For example, on the AIX operating system:

```
/home/db2inst1,/testPath
```

varDBLogPath=*log_location*

Specifies the absolute path where the database log files are stored. The default value of the variable is set to <dftdbpath>/<instance name>/<varDBName>/SQLLOG.

For example:

- On Windows operating systems, c:\\DB2\\TPCDB\\SQLLOG

- On AIX operating systems, /home/db2inst1/db2inst1/TPCDB/SQLLOG

Response file parameters for advanced customization of Tivoli Integrated Portal and Tivoli Common Reporting

These parameters are used for the `silent_SingleServerCustom.properties` and `silent_MultipleServer.properties` response files to configure Tivoli Integrated Portal and Tivoli Common Reporting.

varTIPAdmUsr=*user_name*

Specifies the administrative user name to install Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI.

The default user name for the Windows operating system is `db2admin`. The default user name for the AIX and Linux operating system is `db2inst1`.

When you use an existing instance of Tivoli Integrated Portal, you must enter the Tivoli Integrated Portal administrative user name that was used when you installed that instance. You cannot enter a new user name.

varTIPAdmPW=*password*

Specifies the password for the user name.

varReuseExistingTIP=*value*

Specifies whether to install a new instance of Tivoli Integrated Portal or use an existing instance of Tivoli Integrated Portal that is installed on the server. The values are:

- 0** Install a new instance of Tivoli Integrated Portal 2.2. This option is the default.
- 1** Use an existing instance of Tivoli Integrated Portal 2.2 that is installed on the server.

varTIPPath=*location*

Specifies the location where Tivoli Integrated Portal will be installed. If you are using an existing instance of Tivoli Integrated Portal, enter the location where the program is installed.

For Windows operating systems, use two backslashes. An example of this parameter is:

```
varTIPPath=C:\\Program Files\\IBM\\TIP
```

varTIPPortRangeSP=*port*

Specifies the first port in a range of ports for Tivoli Integrated Portal and Tivoli Common Reporting servers. The default beginning port is 9568.

Tivoli Integrated Portal and Tivoli Common Reporting use these default port numbers:

Tivoli Integrated Portal
9568 - 9580.

Tip: Tivoli Integrated Portal port numbers can vary. To view information about Tivoli Integrated Portal port numbers, see “Tivoli Integrated Portal port numbers” on page 35.

Tivoli Common Reporting
9581

This value overrides the value specified in the **varTPCPortRangeSP** parameter.

varTCRPort=*port*

Specifies the Tivoli Common Reporting port. The default port is 9581. This value overrides the **varTPCServersPortRangeSP** and **varTIPPortRangeSP** parameters.

Response file parameter for the multiple-server environment

This parameter is valid for the `silent_MultipleServer.properties` file.

varMultipleServerComponentList=*components*

Specifies the Tivoli Storage Productivity Center components to install. Separate the components with commas:

database

Installs the Tivoli Storage Productivity Center database repository.

servers

Installs the Data server, Device server, Replication server, the Storage Resource agent, CLI, and the stand-alone GUI.

tiptcr

Installs Tivoli Integrated Portal, Tivoli Common Reporting, the web-based GUI, and Storage Resource agent. You can reuse an existing instance of Tivoli Integrated Portal and Tivoli Common Reporting.

Response file parameters for the remote database

The following parameters are used only in the `silent_MultipleServer.properties` response file to access the Tivoli Storage Productivity Center database on a remote server.

Remember: The DB2 database must be installed before you can use these parameters to access the database.

varRemoteDBSrvName=*host_name*

Specifies the fully qualified host name or IP address of the remote server where the Tivoli Storage Productivity Center database repository is installed.

Tivoli Storage Productivity Center accepts both IPv4 and IPv6 addresses. If you have a system that is configured for dual stack networking (with both IPv4 and IPv6 addresses), Tivoli Storage Productivity Center defaults to IPv4 addressing.

For example, an IPv6 address in long form can be 2001:DB8:0:0:0:0:0:0. An IPv6 address in short form can be 2001:DB8:: or 2001:DB8:0.0.0.0.0.0.

varRemoteDBPort=*port*

Specifies the DB2 port on the remote server. The default port is 50000.

varRemoteDBAdmUsr=*user_name*

Specifies the administrative user name to connect to the remote database. The default user name for the Windows operating system is db2admin. The default user name for the AIX or Linux operating system is db2inst1.

varRemoteDBAdmPW=*password*

Specifies the password for the user name.

varRemoteDBName=*database_name*

Specifies the Tivoli Storage Productivity Center database on the remote server. The default name is TPCDB.

Related reference:

“TCP/IP ports used by Tivoli Storage Productivity Center” on page 33

When you install Tivoli Storage Productivity Center, default ports must be opened through the firewall. You must disable the firewall program or open the ports to allow incoming requests to the Tivoli Storage Productivity Center ports. Review these ports before you install the Tivoli Storage Productivity Center.

Installing Tivoli Storage Productivity Center in a multiple-server environment

You can install Tivoli Storage Productivity Center by using the installation wizard or by using silent mode from the command line.

A multiple-server environment is ideal if you are monitoring large storage environments, where one server is not sufficient to manage the Tivoli Storage Productivity Center components. In this environment, Tivoli Integrated Portal and Tivoli Common Reporting can run on a separate server.

This environment facilitates the integration with other Tivoli products that are using Tivoli Integrated Portal and Tivoli Common Reporting components and allows the sharing of components that are installed on a separate server. If you have different administrators for DB2 and Tivoli Storage Productivity Center, this environment allows the database repository and DB2 to be installed and managed on a separate server.

Related tasks:

“Installing Tivoli Storage Productivity Center with a remote database by using the installation wizard”

You can install Tivoli Storage Productivity Center in a multiple-server environment by using the installation wizard.

“Installing Tivoli Integrated Portal and Tivoli Common Reporting by using the installation wizard” on page 152

You can install Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI in a multiple-server environment using the installation wizard.

“Installing Tivoli Storage Productivity Center with a remote database by using silent mode” on page 154

You can install Tivoli Storage Productivity Center in a multiple-server environment using silent mode. With this configuration, you will be installing the Tivoli Storage Productivity Center components on two servers.

“Installing Tivoli Integrated Portal and Tivoli Common Reporting by using silent mode” on page 157

You can install Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI in a multiple-server environment using silent mode. With this configuration, you will be installing the Tivoli Storage Productivity Center components on two servers.

Installing Tivoli Storage Productivity Center with a remote database by using the installation wizard

You can install Tivoli Storage Productivity Center in a multiple-server environment by using the installation wizard.

Before you start the Tivoli Storage Productivity Center installation wizard on AIX or Linux operating system, you must source the user profile (db2profile) for the instance owner of the DB2 database. For example:

```
. /home/db2inst1/sqllib/db2profile
```

When you install Tivoli Storage Productivity Center, a database called TPCDB is created. In this repository, the TPC database schema is created, which describes the structure of the database repository. If you already have the TPCDB on your system, Tivoli Storage Productivity Center assigns a new default database repository name with a number as a suffix for example, TPCDB1.

DB2 Enterprise Server Edition and the Tivoli Storage Productivity Center database repository are installed on one server, and the DB2 Client and the other Tivoli Storage Productivity Center components are installed on the second server.

For this procedure, the terms *Server A* and *Server B* denote the two servers. Server A has DB2 Enterprise Server Edition and the TPCDB database repository installed. Server B has the DB2 client and the other Tivoli Storage Productivity Center components installed.

To install Tivoli Storage Productivity Center with a remote database by using the installation wizard, complete the following steps:

1. Complete the following steps on Server A:
 - a. Log on to Server A.
 - b. Install DB2 Enterprise Server Edition. For more information about installing DB2, see “Installing DB2” on page 180.
 - c. If you are installing on a Windows operating system, log out and then log on again before installing Tivoli Storage Productivity Center.
 - d. Start the Tivoli Storage Productivity Center installation wizard.
 - e. Select the language for your installation and click **OK**.
 - f. On the Choose Installation Location and Type page, select **Multiple servers**.
 - g. Follow the instructions in the installation wizard and install the Tivoli Storage Productivity Center database repository on this server.
 - h. Select **Database repository**.
 - i. Follow the installation wizard to complete the installation.
 - j. Review the message on the summary page. The message indicates that you have successfully installed the database repository.
2. Complete the following steps on Server B:
 - a. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
 - b. Install the DB2 client. For more information about installing DB2, see “Installing DB2” on page 180.
 - c. If you are installing the DB2 client on a Windows operating system, log out and then log on again before installing Tivoli Storage Productivity Center.
 - d. On the Launch Pad page, select **Install a Product**.
 - e. Install the IBM Data Server Client. Only the DB2 client ID, db2caep1, is required on the client host.
 - f. Follow the steps for creating the instance owner ID. The administrative ID, dasusr1, and the fenced ID, db2fenc1, are not required.
 - g. Start the Tivoli Storage Productivity Center installation wizard. Before you start the Tivoli Storage Productivity Center installation wizard on AIX or Linux operating system, you must source the user profile, db2profile, for the instance owner of the DB2 database. For example:

```
. /home/db2inst1/sqllib/db2profile
```
 - h. Select the language for your installation and click **OK**.
 - i. On the Choose Installation Location and Type page, select **Multiple servers**.

- j. Follow the instructions in the installation wizard to install the following components on this server:
 - Tivoli Storage Productivity Center servers
 - web-based GUI
 - stand-alone GUI
 - Tivoli Integrated Portal
 - Tivoli Common Reporting
 - CLI
 - Storage Resource agent
- k. Select **Tivoli Storage Productivity Center servers and Tivoli Integrated Portal**.
- l. After the installation is finished, review the message log to ensure that no errors occurred.

Related tasks:

“Installing DB2” on page 180

This topic describes how to install DB2 Version 9.7 for operating systems such as Linux, AIX, or Windows.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Starting the installation program” on page 138

This topic describes how to start the Tivoli Storage Productivity Center installation programs.

“Tivoli Storage Productivity Center components” on page 133

You can install Tivoli Storage Productivity Center in single- or multiple-server environments. In a single-server environment, all components are installed on one server.

“Verifying the installation” on page 163

After you install Tivoli Storage Productivity Center, you can verify whether the installation was successful.

“Reviewing the log files to resolve installation issues” on page 165

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

Installing Tivoli Integrated Portal and Tivoli Common Reporting by using the installation wizard

You can install Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI in a multiple-server environment using the installation wizard.

Before you start the Tivoli Storage Productivity Center installation wizard on the AIX or Linux operating systems, you must source the user profile, `db2profile`, for the instance owner of the DB2 database. For example:

```
. /home/db2inst1/sqllib/db2profile
```

The DB2 Enterprise Server Edition, the Tivoli Storage Productivity Center servers, database repository, the stand-alone GUI, CLI, and Storage Resource agent are installed on one server, and the DB2 client, Tivoli Integrated Portal, Tivoli Common Reporting, and web-based GUI are installed on the second server.

For this procedure, the terms *Server A* and *Server B* denote the two servers. Server A has DB2 Enterprise Server Edition, the Tivoli Storage Productivity Center servers, the database repository, CLI, Storage Resource agent, and stand-alone GUI installed. Server B has the DB2 client, Tivoli Integrated Portal, Tivoli Common Reporting, and web-based GUI installed.

To install Tivoli Integrated Portal and Tivoli Common Reporting by using the installation wizard:

1. Complete the following steps on Server A:
 - a. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
 - b. Install DB2 Enterprise Server Edition. For more information about installing DB2, see “Installing DB2” on page 180.
 - c. If you are installing on a Windows operating system, log out and then log on again before installing Tivoli Storage Productivity Center.
 - d. Start the Tivoli Storage Productivity Center installation wizard.
 - e. Select the language for your installation and click **OK**. The language that you select is used for product installation.
 - f. On the Choose Installation Location and Type page, select **Multiple servers**.
 - g. Follow the instructions in the installation wizard and install the Tivoli Storage Productivity Center components and database repository on this server. On the Multiple Servers Option page, select **Database repository and Tivoli Storage Productivity Center servers**. Follow the instructions in the installation wizard to install the Tivoli Storage Productivity Center components.
 - h. After the installation is finished, review the message log to ensure that no errors occurred.
2. Complete the following steps on Server B:
 - a. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
 - b. Install the DB2 Client. For more information about installing DB2, see “Installing DB2” on page 180.
 - c. If you are installing on a Windows operating system, log out and then log on again before installing Tivoli Storage Productivity Center.
 - d. On the Launch Pad page, select **Install a Product**.
 - e. Install the IBM Data Server Client. Only the DB2 client ID (db2caep1) is required on the client host.
 - f. Follow the steps for creating the instance owner ID. The administrative ID (dasusr1) and the fenced ID (db2fenc1) are not required.
 - g. Start the Tivoli Storage Productivity Center installation wizard. Before you start the Tivoli Storage Productivity Center installation wizard on AIX or Linux operating system, you must source the user profile (db2profile) for the instance owner of the DB2 database. For example:

```
. /home/db2inst1/sqllib/db2profile
```
 - h. Select the language for your installation and click **OK**.
 - i. On the Choose Installation Location and Type page, select **Multiple servers**.
 - j. Follow the instructions in the installation wizard and install Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI on this server.
 - k. On the Choose Multiple Server Option page, select **Tivoli Integrated Portal**.

- l. After the installation is finished, review the message log to ensure that no errors occurred.

If you install Tivoli Storage Productivity Center using LDAP authentication, you can log in to Tivoli Integrated Portal using either LDAP or the local administrative user. The local administrative user can log in to Tivoli Integrated Portal to make administrative changes even if the LDAP server is offline or configured incorrectly. To restrict log in access to LDAP authentication, remove the local administrative user from Tivoli Integrated Portal after you install Tivoli Storage Productivity Center.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Starting the installation program” on page 138

This topic describes how to start the Tivoli Storage Productivity Center installation programs.

“Verifying the installation” on page 163

After you install Tivoli Storage Productivity Center, you can verify whether the installation was successful.

“Reviewing the log files to resolve installation issues” on page 165

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

Installing Tivoli Storage Productivity Center with a remote database by using silent mode

You can install Tivoli Storage Productivity Center in a multiple-server environment using silent mode. With this configuration, you will be installing the Tivoli Storage Productivity Center components on two servers.

Before you start the Tivoli Storage Productivity Center installation program on the AIX or Linux operating system, you must source the user profile, `db2profile`, for the instance owner of the DB2 database. For example:

```
. /home/db2inst1/sqllib/db2profile
```

The DB2 Enterprise Server Edition and the Tivoli Storage Productivity Center database repository are installed on one server, and the DB2 client and the other Tivoli Storage Productivity Center components are installed on the second server.

When you install Tivoli Storage Productivity Center, a database called TPCDB is created. In this repository, the TPC database schema is created, which describes the structure of the database repository. If you already have the TPCDB on your system, Tivoli Storage Productivity Center assigns a new default database repository name with a number as a suffix for example, TPCDB1.

For this procedure, the terms *Server A* and *Server B* denote the two servers. Server A has DB2 Enterprise Server Edition and the TPCDB database repository installed. Server B has the DB2 client and the other Tivoli Storage Productivity Center components installed.

To install Tivoli Storage Productivity Center with a remote database by using silent mode, complete the following steps:

1. Complete the following steps on Server A:
 - a. Log on to Server A to install DB2.
 - b. Install DB2 Enterprise Server Edition. For more information about installing DB2, see "Installing DB2" on page 180
 - c. If you are installing on a Windows operating system, log out and then log on again before installing Tivoli Storage Productivity Center.
 - d. In `silent_MultipleServer.properties`, set the following parameters for the remote database repository:
 - `CHOSEN_INSTALL_TYPE="Multiple Server Install"`
 - `varMultipleServerComponentList=database`
 - `USER_INSTALL_DIR=option`
 - `varDBAdmUsr=user_name`
 - `varDBAdmPW=password`
 - `varDBName=database_name`
 - `varDBPath=database_location`
 - `varDBLogPath=log_location`

Restriction: If you install Tivoli Storage Productivity Center by using silent mode, and are not using the response file, you must set `-DLICENSE_ACCEPTED=true` to accept the license agreement. For example:
`setup.exe/bin -i silent -DLICENSE_ACCEPTED=true`

- e. Save the response file.
- f. Run the Tivoli Storage Productivity Center installation program on the DB2 server to install the database repository.
 - For Windows operating systems, run the following command:
`setup.bat -l language -i silent -f absolute_path_to_response_file`

where *language* is one of the following languages:

language

- Czech - cs
- English - en
- French - fr
- German - de
- Hungarian - hu
- Italian - it
- Japanese - ja
- Korean - ko
- Polish - pl
- Brazilian Portuguese - pt_BR
- Russian - ru
- Spanish - es
- Chinese (Simplified) - zh_CN
- Chinese (Traditional) - zh_TW

absolute_path_to_response_file

The absolute path to the response file. For example, the following command specifies the language and the path:

```
setup.bat -l fr -i silent -f c:\TPC51\  
silent_MultipleServer.properties.
```

- For AIX or Linux operating systems, run the following command:

```
./setup.bin -l language -i silent  
-f absolute_path_to_response_file
```

For example, the following command specifies the language and the path:

```
./setup.bin -l de -i silent -f /TPC51/silent_MultipleServer.properties
```

If you do not specify a language or specify the language incorrectly, the installation program uses the English language as the default.

- g. Optional: Monitor the progress of the installation.

To view the progress of the silent mode installation on Windows operating systems, check the trace log in the following path:

```
TPC_installation_directory\logs\traceTPCInstall.log
```

To view the progress of the silent mode installation on AIX or Linux operating systems, check the trace log in the following path:

```
TPC_installation_directory/logs/traceTPCInstall.log
```

2. Complete the following steps on Server B:

- a. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.

- b. Install the DB2 Client.

If you are installing on a Windows operating system, log out and then log on again before installing Tivoli Storage Productivity Center.

- c. Before you start the Tivoli Storage Productivity Center installation program on the AIX or Linux operating system, you must source the user profile (db2profile) for the instance owner of the DB2 database. For example:

```
. /home/db2inst1/sqllib/db2profile
```

- d. Edit and save the `silent_MultipleServer.properties` file. Set the following parameters for Tivoli Storage Productivity Center.

- CHOSEN_INSTALL_TYPE="Multiple Server Install"
- varMultipleServerComponentList=*servers, tiptcr*
- USER_INSTALL_DIR=*option*
- varSrvName=*option*
- varCommonUsrID=*user_name*
- varCommonUsrPw=*password*
- varRemoteDBSrvName=*host_name*
- varRemoteDBPort=*port*
- varRemoteDBAdmUsr=*user_name*
- varRemoteDBAdmPW=*password*
- varRemoteDBName=*database_name*
- varTPCServersPortRangeSP=*port*
- varTIPAdmUsr=*user_name*
- varTIPAdmPW=*password*
- varTIPPath=*location*
- varTIPPortRangeSP=*port*
- varTIPPortRangeSP=*port*

- e. Run the Tivoli Storage Productivity Center installation program to install the Tivoli Storage Productivity Center components (excluding the database repository).

- For Windows operating systems, run the following command:

```
setup.bat -l language -i silent -f  
absolute_path_to_response_file
```

- For AIX or Linux operating systems, run the following command:

```
./setup.bin -l language -i silent  
-f /absolute_path_to_response_file
```

For example, the following command specifies that the language to install is German:

```
./setup.bin -l de -i silent -f /TPC51/  
silent_MultipleServer.properties
```

If you do not specify a language or specify the language incorrectly, the installation program uses the English language as the default.

- f. Optional: Monitor the progress of the installation.

To view the progress of the silent mode installation on Windows operating systems, check the trace log in the following path:

```
TPC_installation_directory\logs\traceTPCInstall1.log
```

To view the progress of the silent mode installation on AIX or Linux operating systems, check the trace log in the following path:

```
TPC_installation_directory/logs/traceTPCInstall1.log
```

Related reference:

“Editing the response file” on page 143

You must edit and save the appropriate response file when you install Tivoli Storage Productivity Center in silent mode. The silent mode install option for the response file is `-f absolute_path_to_response_file`. For example, in the Windows operating system, you enter `-f C:\installimage\
silent_SingleServerTypical.properties`.

Installing Tivoli Integrated Portal and Tivoli Common Reporting by using silent mode

You can install Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI in a multiple-server environment using silent mode. With this configuration, you will be installing the Tivoli Storage Productivity Center components on two servers.

Before you start the Tivoli Storage Productivity Center installation program on AIX or Linux operating system, you must source the user profile, `db2profile`, for the instance owner of the DB2 database. For example:

```
. /home/db2inst1/sqllib/db2profile
```

The DB2 Enterprise Server Edition, the Tivoli Storage Productivity Center servers, database repository, the stand-alone GUI, CLI, and Storage Resource agent are installed on one server, and the DB2 client, Tivoli Integrated Portal, Tivoli Common Reporting, and web-based GUI are installed on the second server.

For this procedure, the terms *Server A* and *Server B* denote the two servers. Server A has DB2 Enterprise Server Edition, the Tivoli Storage Productivity Center servers, the database repository, CLI, Storage Resource agent, and stand-alone GUI installed. Server B has the DB2 client, Tivoli Integrated Portal, Tivoli Common Reporting, and web-based GUI installed.

To install Tivoli Integrated Portal and Tivoli Common Reporting by using silent mode, complete the following steps:

1. Complete the following steps on Server A:
 - a. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
 - b. Install DB2 Enterprise Server Edition. For more information about installing DB2, see “Installing DB2” on page 180
 - c. If you are installing on a Windows operating system, log out and then log on again before installing Tivoli Storage Productivity Center.
 - d. In `silent_MultipleServer.properties` file, set all the parameters in the file except for the Tivoli Integrated Portal and Tivoli Common Reporting customization parameters.

Restriction: If you install Tivoli Storage Productivity Center by using silent mode, and are not using the response file, you must set

`-DLICENSE_ACCEPTED=true` to accept the license agreement. For example:
`setup.exe/bin -i silent -DLICENSE_ACCEPTED=true`

- e. Save `silent_MultipleServer.properties`.
- f. Run the Tivoli Storage Productivity Center installation program on the DB2 server to install the database repository and Tivoli Storage Productivity Center components.

- For Windows operating systems, run the following command:
`setup.bat -l language -i silent -f absolute_path_to_response_file`

where *language* can be one of the following languages:

language

- Czech - cs
- English - en
- French - fr
- German - de
- Hungarian - hu
- Italian - it
- Japanese - ja
- Korean - ko
- Polish - pl
- Brazilian Portuguese - pt_BR
- Russian - ru
- Spanish - es
- Chinese (Simplified) - zh_CN
- Chinese (Traditional) - zh_TW

absolute_path_to_response_file

The absolute path to the response file. For example, the following command specifies the language and the path:

```
setup.bat -l fr -i silent -f c:\TPC51\  
silent_MultipleServer.properties.
```

- For AIX or Linux operating systems, run the following command:
`./setup.bin -l language -i silent -f /absolute_path_to_response_file`

For example, the following command specifies the language and the path:

```
./setup.bin -l de -i silent -f /TPC51/silent_MultipleServer.properties
```

If you do not specify a language or specify the language incorrectly, the installation program uses the English language as the default.

- g. Optional: Monitor the progress of the installation.

To view the progress of the silent mode installation on Windows operating systems, check the trace log in the following path:

```
TPC_installation_directory\logs\traceTPCInstall.log
```

To view the progress of the silent mode installation on AIX or Linux operating systems, check the trace log in the following path:

```
TPC_installation_directory/logs/traceTPCInstall.log
```

2. Complete the following steps on Server B:

- a. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
- b. Install the DB2 Client.
- c. If you are installing on a Windows operating system, log out and then log on again before installing Tivoli Storage Productivity Center.
- d. Before you start the Tivoli Storage Productivity Center installation program on AIX or Linux operating system, you must source the user profile, `db2profile`, for the instance owner of the DB2 database. For example:

```
./home/db2inst1/sqllib/db2profile
```

- e. In the `silent_MultipleServer.properties` file, set the following parameters for Tivoli Integrated Portal and Tivoli Common Reporting:

- `CHOSEN_INSTALL_TYPE="Multiple Server Install"`
- `varMultipleServerComponentList=tiptcr`
- `USER_INSTALL_DIR=option`
- `varRemoteDBSrvName=host_name`
- `varRemoteDBPort=port`
- `varRemoteDBAdmUsr=user_name`
- `varRemoteDBAdmPW=password`
- `varRemoteDBName=database_name`
- `varTIPAdmUsr=option`
- `varTIPAdmPW=password`
- `varTIPPath=TIP_location`
- `varTIPPortRangeSP=log_location`
- `varTCRPort=port`

- f. Run the Tivoli Storage Productivity Center installation program on Server B to install Tivoli Integrated Portal and Tivoli Common Reporting.

- For Windows operating systems, run the following command:

```
setup.bat -l language -i silent -f  
absolute_path_to_response_file
```

where *language* can be one of the following languages:

language

- Czech - cs
- English - en
- French - fr
- German - de
- Hungarian - hu

- Italian - it
- Japanese - ja
- Korean - ko
- Polish - pl
- Brazilian Portuguese - pt_BR
- Russian - ru
- Spanish - es
- Chinese (Simplified) - zh_CN
- Chinese (Traditional) - zh_TW

absolute_path_to_response_file

The absolute path to the response file. For example, the following command specifies the language and the path:

```
setup.bat -l fr -i silent -f c:\TPC51\
silent_MultipleServer.properties.
```

- For AIX or Linux operating systems, run the following command:

```
./setup.bin -l language -i silent -f /
absolute_path_to_response_file
```

For example, the following command specifies the language and the path:

```
./setup.bin -l de -i silent -f /TPC51/silent_MultipleServer.properties
```

If you do not specify a language or specify the language incorrectly, the installation program uses the English language as the default.

- g. Optional: Monitor the progress of the installation.

To view the progress of the silent mode installation on Windows operating systems, check the trace log in the following path:

```
TPC_installation_directory\logs\traceTPCInstall.log
```

To view the progress of the silent mode installation on AIX or Linux operating systems, check the trace log in the following path:

```
TPC_installation_directory/logs/traceTPCInstall.log
```

If you experience an error, check the installation log files.

If you install Tivoli Storage Productivity Center using LDAP authentication, you can log in to Tivoli Integrated Portal using either LDAP or the local administrative user. The local administrative user can log in to Tivoli Integrated Portal to make administrative changes even if the LDAP server is offline or configured incorrectly. To restrict log in access to LDAP authentication, remove the local administrative user from Tivoli Integrated Portal after you install Tivoli Storage Productivity Center.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Starting the installation program” on page 138

This topic describes how to start the Tivoli Storage Productivity Center installation programs.

“Verifying the installation” on page 163

After you install Tivoli Storage Productivity Center, you can verify whether the installation was successful.

“Reviewing the log files to resolve installation issues” on page 165

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

Installing Tivoli Storage Productivity Center client components on a separate computer

You can install IBM Tivoli Storage Productivity Center client components on a computer other than the Tivoli Storage Productivity Center Server.

To install Tivoli Storage Productivity Center client components on a separate computer, complete the following steps:

1. Install the Tivoli Storage Productivity Center server.
2. Navigate to the `client_images` subdirectory in the Tivoli Storage Productivity Center installation directory. In the Windows operating system, for example, the path is `C:\Program Files\IBM\TPC\client_images`.
3. Locate the appropriate compressed file for your operating system of the computer where you want to install the stand-alone GUI and command-line interface (CLI). Examples of compressed bundles include `TPC_CLIENT_AIX.tar`, `TPC_CLIENT_LINUX.zip`, or `TPC_CLIENT_WIN.zip`.

Restriction: You must use only the client images from Tivoli Storage Productivity Center installation directory. The client images in the Tivoli Storage Productivity Center image or download directory must not be used, because they are updated by the installation program.

4. Extract the appropriate compressed file on another computer. For example, in the AIX operating system, you must extract the `TPC_CLIENT_AIX.tar` into `/opt/IBM/TPCClient` folder and run `/opt/IBM/TPCClient/gui/TPCD.sh` to start stand-alone GUI.

Tivoli Storage Productivity Center GUI and CLI components are now installed on a separate computer. You can issue Tivoli Storage Productivity Center commands and run the stand-alone GUI from that computer.

Installing Tivoli Common Reporting on a remote server

Use this procedure to install and configure Tivoli Common Reporting on a remote server.

To install Tivoli Common Reporting on a remote server, complete these steps:

1. Install the DB2 client on the remote server. For information about installing the DB2 client, see the DB2 information center. For more information about setting up multiple connections to the TPCDB data source, go to the Tivoli Storage Productivity information center and search for *Setting up multiple connections to the TPCDB data source in Tivoli Common Reporting*.
2. Install Tivoli Common Reporting on the remote server.
3. Create the data source in DB2 for the remote server by completing the following steps:

- a. Catalog the DB2 TCP/IP node:

Run the following command:

```
CATALOG TCPIP NODE db2_node_name
REMOTE host_name
SERVER port_number
REMOTE_INSTANCE db2_instance_name
SYSTEM system_name
OSTYPE operating_system_type
```

In the preceding command, specify the following values:

CATALOG TCPIP NODE *db2_node_name*

This value specifies the node name of the DB2 server.

REMOTE *host_name*

This value specifies the host name or IP address of the DB2 server.

SERVER *port_number*

This value specifies the port number of the DB2 server. The default port is 50000.

REMOTE_INSTANCE *db2_instance_name*

This value specifies the DB2 instance name on the remote server.

SYSTEM *system_name*

This value specifies the DB2 system name that is used to identify the server system.

OSTYPE *operating_system_type*

This value specifies the operating system type of the server system.
Valid values for Tivoli Common Reporting are: AIX, WIN, and LINUX.

An example for this command is:

```
CATALOG TCPIP NODE db2node
REMOTE system53.ibm.com
SERVER 50000
REMOTE_INSTANCE db2 system53
SYSTEM system53
OSTYPE win
```

- b. Catalog the Tivoli Storage Productivity Center database.

Run the following command:

```
CATALOG DATABASE tpc_database_name AS tpc_database_alias_name
AT NODE db2_node_name AUTHENTICATION SERVER
```

In the preceding command, specify the following values:

CATALOG DATABASE *tpc_database_name*

This value specifies the name of the remote database to catalog.

AS *tpc_database_alias_name*

This value specifies the alias for the remote database that is being cataloged. If you do not specify an alias, the database manager uses the name of the remote database as the alias.

AT NODE *db2_node_name*

This value specifies the name of the node that you previously cataloged.

AUTHENTICATION SERVER

This value specifies that authentication takes place on the DB2 data source node.

An example for this command is:

```
CATALOG DATABASE tpc_database_name AS tpc_database_alias_name
AT NODE db2node AUTHENTICATION SERVER
```

4. Type the following command on the client to connect to the remote database:
`CONNECT TO tpc_database_alias_name USER userid`
5. When prompted, enter your password.
6. If the connection is successful, move to step 7. If not, verify that step 3 was done correctly.
7. Configure Tivoli Common Reporting. For information about configuring Tivoli Common Reporting, go to the Tivoli Common Reporting information center and search for *Configuring IBM Tivoli Common Reporting*.

Related tasks:

“Installing Tivoli Integrated Portal and Tivoli Common Reporting by using the installation wizard” on page 152

You can install Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI in a multiple-server environment using the installation wizard.

“Installing Tivoli Integrated Portal and Tivoli Common Reporting by using silent mode” on page 157

You can install Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI in a multiple-server environment using silent mode. With this configuration, you will be installing the Tivoli Storage Productivity Center components on two servers.

Verifying the installation

After you install Tivoli Storage Productivity Center, you can verify whether the installation was successful.

Use the following questions to guide you through the verification process.

Restriction: You must do the following tasks:

- Use a user name with DB2 administrative privileges to run the DB2 commands.
- Run the commands on the DB2 command line.

Was the TPCDB database created?

To verify that the TPCDB database exists, run the following command:

```
db2 list db directory
```

This command returns information about the TPCDB database.

Are the tables and views present in the TPCDB database?

To verify that the tables and views are present, run the following DB2 commands:

```
db2 connect to TPCDB
db2 list tables for schema TPC
db2 disconnect TPCDB
```

These commands indicate that you can connect to and disconnect from the Tivoli Storage Productivity Center database repository and list the tables and views in the TPCDB database.

Is Tivoli Integrated Portal running?

For Windows operating systems:

To verify that Tivoli Integrated Portal is running, run the following command:

```
ps -eaf | grep tip
```

This command returns a list of the Tivoli Integrated Portal processes in the process table.

For the AIX or Linux operating systems:

To verify that Tivoli Integrated Portal is running, go to the Services listing and find **Tivoli Integrated Portal - V2.2_TIPProfile_Port_***port*. *port* is the port used by Tivoli Integrated Portal, and the default port is 9568.

Is Tivoli Common Reporting running?

For Windows operating systems:

To verify that Tivoli Common Reporting is running, go to the Services listing and find **IBM Cognos Content Database**.

For the AIX or Linux operating systems:

To verify that Tivoli Common Reporting is running, run the following command:

```
ps -eaf | grep NetworkServerControl
```

This command returns a list of the IBM® Cognos® Business Intelligence server process in the processes table.

Is the Device server running?

For Windows operating systems:

To verify that the Device server is running, go to the Services window and find **IBM WebSphere Application Server V8.0 - TPCDeviceServer**.

For the AIX or Linux operating systems:

To verify that the Device server is running, run the following command:

```
ps -eaf | grep deviceServer
```

This command returns a list of the Device server processes in the process table.

Is the Data server running?

For the Windows operating systems:

To verify that the Data server is running, go to the Services listing and find **IBM Tivoli Storage Productivity Center - Data Server**.

For the AIX or Linux operating systems:

To verify that the Data server is running, run the following command:

```
ps -eaf | grep DataServer.jar
```

This command returns a list of the Data server processes in the process table.

Is the Replication server running?

For the Windows operating system

To verify that the Replication server is running, go to the Services window and find **IBM WebSphere Application Server V8.0 - TPCReplicationServer**.

For the AIX or Linux operating systems:

To verify that the Replication server is running, run the following command:

```
ps -eaf | grep replicationServer
```

This command returns a list of the Replication server processes in the process table.

Is the Storage Resource agent running?

For the Windows operating systems:

To verify that the Storage Resource agent is running, go to the Services window and find **IBM Tivoli Storage Resource Agent**.

For the AIX or Linux operating systems:

To run the following command:

```
ps -eaf | grep Agent
```

This command returns a list of the Storage Resource agent processes in the process table.

Can I log on to the web-based GUI ?

Open a web browser and enter the following web address:

```
http://host_name:port/srm
```

where *host_name* is the fully qualified host name where the web-based GUI is installed, and *port* is the port that is used by the web-based GUI. The default port is 9568.

Are you running in a multiple-server environment?

In a multiple-server environment, Tivoli Storage Productivity Center authenticates using the following rules:

- Device server authentication is done through the Tivoli Integrated Portal server. To run the **tpctool** command, you must authenticate from the Tivoli Integrated Portal server and not from the Tivoli Storage Productivity Center server.
- For the stand-alone GUI, you must have a user name from the Tivoli Integrated Portal server and not from the Tivoli Storage Productivity Center server. The user name must be in the system group for the AIX operating system, in the root group for the Linux operating system, and in the Administrators group for the Windows operating system. You can also use a user name that is in a group that is mapped to the Tivoli Storage Productivity Center superuser role.
- For the Replication server, authorization is done through the access list. You can configure the user access list in the administration console or the CLI. At installation time, the common user name and `tpcFileRegistryUser` are the only users configured. If you want other user names authorized, you must add those user names.

Reviewing the log files to resolve installation issues

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

To resolve errors before you install Tivoli Storage Productivity Center, complete the following steps:

1. Navigate to one of the following directories:
 - For the Windows operating system, go to %TEMP% directory
 - For the UNIX or Linux operating system, go to /tmp directory
2. Review the following log files:
 - lax-xxxxx-out.txt
 - lax-xxxxx-err.txt

Tip: The files are InstallAnywhere log files and are automatically generated. The *xxxxx* are numeric characters that are assigned by the InstallAnywhere software.

To resolve errors after the Tivoli Storage Productivity Center installation starts, complete the following steps:

1. Navigate to one of the following default directories:
 - For the Windows operating system, go to C:\Program Files\IBM\TPC\logs
 - For the AIX or Linux operating system, go to /opt/IBM/TPC/logs
2. Review these log files:
 - msgTPCInstall.log
 - traceTPCInstall.log
 - sra\agent.trace
 - sra\agent_number.log
3. Review the Tivoli Integrated Portal log file:
 - On the Windows operating system, navigate to the C:\Users\Administrator\TIPInstaller-00.log file.
You are installing Tivoli Integrated Portal as an administrator.
 - On the Linux operating system, navigate to the /root/TIPInstaller-00.log file.
 - On the AIX operating system, navigate to the /home/root/TIPInstaller-00.log file.
4. For agent log files, navigate to one of the following default directories:
 - For the Windows operating system, go to *SRA_location*\install
 - For the UNIX or Linux operating system, go to /*SRA_location*/install where *SRA_location* is where the Storage Resource agent is installed.
5. Navigate to one of the following directories:
 - For the Windows operating system, go to C:\Users\administrator\
 - For the Linux operating system, go to /root
 - For the AIX operating system, go to /home/root

If an installation error occurred and you decide to do a partial rollback, the msgTPCInstall.log and traceTPCInstall.log files are updated with additional information when you resume the installation.

Starting the stand-alone GUI by using a web browser

You can start the stand-alone GUI by using Java Web Start.

For information about starting the stand-alone GUI in a web browser, see “Starting the Tivoli Storage Productivity Center stand-alone GUI by using Java Web Start” on page 196.

Adding an installation license

If you installed Tivoli Storage Productivity Center and want to add a Tivoli Storage Productivity Center license to use additional functions of the software, you can use the installation wizard or silent-mode installation.

To determine the type of license you have, find these files in the `TPC_installation_directory/properties/version` directory: :

Tivoli_Storage_Productivity_Center.5.1.0.swtag

Indicates that you have a Tivoli Storage Productivity Center license.

Tivoli_Storage_Productivity_Center_Advanced.5.1.0.swtag

Indicates that you have a Tivoli Storage Productivity Center Advanced license.

Tivoli_Storage_Productivity_Center_Select.5.1.0.swtag

Indicates that you have a Storage Productivity Center Select license.

Tivoli_Storage_Productivity_Center_Embedded.5.1.0.swtag

Indicates that you have a license for the Tivoli Storage Productivity Center product included with IBM Systems Director.

Depending on the license you purchased, one of the following license key files can be found in the `license\key` sub-directory, on disk1, part1 of the DVD, or electronic image:

node1ock

Tivoli Storage Productivity Center license

node1ock.SEL

Storage Productivity Center Select license

node1ock.AE

Tivoli Storage Productivity Center Advanced license

node1ock.EE

Tivoli Storage Productivity Center product included with IBM Systems Director

Adding an installation license using the installation wizard

If you installed Tivoli Storage Productivity Center and want to add a Tivoli Storage Productivity Center license for additional function, you can do so by using the installation wizard.

To add the license to your system using the installation wizard, complete the following steps:

1. If you are adding a license on the AIX or Linux operating system, source the `db2profile`.
2. Log on to your Tivoli Storage Productivity Center system with the appropriate user privileges.
3. Start the Tivoli Storage Productivity Center installation wizard from the DVD or from the license image you want to add.
4. On the installation wizard page, select a language and click **OK**. The language you select is used to install the license.
5. On the Choose Installation Location and Type page, select **License upgrade**.
6. On the successfully installed page, click **Finish**.
7. Stop and restart the Tivoli Storage Productivity Center GUI.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

Adding an installation license using silent mode

If you installed Tivoli Storage Productivity Center and want to add a Tivoli Storage Productivity Center license to access additional function, you can do so by using silent-mode installation.

To add the license to your system using silent mode, complete the following steps:

1. If you are adding a license on the AIX or Linux operating system, source the `db2profile`.
2. Log on to your Tivoli Storage Productivity Center system with the appropriate user privileges.
3. Start the Tivoli Storage Productivity Center installation wizard from the DVD or from the license image you want to add.
4. Modify the following parameters in the `silent_SingleServerTypical.properties`, `silent_SingleServerCustom.properties`, or `silent_MultipleServer.properties` file:
 - `CHOSEN_INSTALL_TYPE="License Upgrade"`
 - `varUseLicenseKeyOnImage=`
 - `varLicenseKeyFile=`
5. Save the appropriate response file.
6. Stop and restart the Tivoli Storage Productivity Center GUI.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Editing the response file” on page 143

You must edit and save the appropriate response file when you install Tivoli Storage Productivity Center in silent mode. The silent mode install option for the response file is `-f absolute_path_to_response_file`. For example, in the Windows operating system, you enter `-f C:\installimage\silent_SingleServerTypical.properties`.

Installing the agents

You can install the Storage Resource agent using the installation program, the Tivoli Storage Productivity Center user interface, or a command.

You can select one of the following agent installation scenarios:

- To deploy the Storage Resource agent using the GUI, see “Storage Resource agent Deployments” on page 281.
- To install the agent using a command, see “Installing the agent using a command” on page 169.

- To install the agent on a Virtual I/O Server, see “Installing the Storage Resource agent on the Virtual I/O Server” on page 172.
- To deploy the agent remotely on the Virtual I/O Server, see “Deploying the Storage Resource agent remotely” on page 172.

Restriction:

- Before you install or deploy the Storage Resource agent on a system, you must disable the firewall on that system.
- If you are using PowerHA software, a Storage Resource agent must be installed on each node of the cluster and all agents in a cluster must be configured to use the same listening port.
- To scan a cluster resource group, you must configure the cluster resource group to have at least one IP address that is accessible from the Tivoli Storage Productivity Center server.

To install the Storage Resource agents, you must log in as a user with the following authority:

On Windows operating systems

You must have Windows Administrator authority.

On UNIX or Linux operating systems

You must be logged in as the root user.

Preparing to install Storage Resource agents

This section provides information on preparing to install the Storage Resource agents.

General information

For information about deploying Storage Resource agents, see “Deployment guidelines and limitations for Storage Resource agents” on page 282.

Installing the agent using a command

You can install the Storage Resource agent using a command.

You would typically install the Storage Resource agents using the Tivoli Storage Productivity Center GUI. However, if you need to install the Storage Resource agents locally, you can do so with limited support.

During this installation method, a return code of zero means that a successful installation occurred and a nonzero return code means that an unsuccessful installation occurred. If you have an unsuccessful installation, you must review the log files to determine the problem.

The **-force** option can be used when you have a Storage Resource agent pointing to multiple servers. If one server has installed an agent, and another server wants to install an agent in the same location, the second server can use the **-force** option to install the agent.

If you use this installation method, when you enter a directory to install the Storage Resource agent, do not add an ending slash mark (\). For example, do not specify C:\agent1\. This action causes the installation to fail.

If you run the agent as a non-daemon service (On-Demand service), you must make sure that at least one protocol is valid for a successful connection from the server to the agent. See Table 29 for the required parameters for each protocol.

Table 29. Parameters required for each protocol

Protocol	Description
SSH	Requires the user ID and password or user ID, certificate, and passphrase.
Windows (SMB)	Requires the user ID and password.
REXEC	Requires user ID and password.
RSH	Requires the user ID.

The Storage Resource Agent image contains the installation images for the Storage Resource agents in the following directory:

DVD_installation_image_location/data/sra/operating_system

See Table 30 for the Storage Resource agent installation images.

Table 30. Storage Resource agent installation images

Operating system	Operating system name
AIX	aix_power
HP-UX	hp-ux_itanium
Linux x86	linux_ix86
Linux for Power Systems Servers	linux_power
Linux s390	linux_s390
Oracle Solaris	solaris_sparc
Windows	windows

To install the Storage Resource agents locally, complete the following steps:

1. Go to the installation image location:

```
cd DVD_installation_image_location
```

2. Run the following command:

```

▶ bin/Agent -install [-force] -serverPort server_secure_port
▶ -serverIP server_IP_address -installLoc Agent_install_location [-debug MAX]
▶ -agentPort agent_port -commtype daemon (1)
▶ -userID user_ID -password password -certFile certificate_file (2)
▶ -passphrase passphrase (3)

```

Notes:

- 1 Parameters when the agent is run as a daemon service.

- 2 Parameters when the agent is run as a non-daemon service. See Table 29 on page 170 to determine which parameter is required for each protocol.
- 3 Parameters when the agent is run as a non-daemon service. See Table 29 on page 170 to determine which parameter is required for each protocol.

The parameters are as follows:

-install

Installs the Storage Resource agent.

-force

This forces the agent to be installed. There are two different situations in which this parameter must be specified:

- If an earlier installation failed and there is residue on the system which causes further installations to fail. You must make sure that all the parameters provided are valid, such as the installation location, port, and so on.
- If the agent is already installed from one server and you now must install the agent pointing to another server.

-serverPort *server_secure_port*

The default port for the Data server is 9549.

-serverIP *server_IP_address*

The IP address of the server. If the server can be reached through multiple IP addresses, then multiple IP addresses can be specified with IP addresses separated with a comma.

-installLoc "*Agent_install_location*"

Location where the agent is installed. Enclose the directory name in quotation marks, for example, "C:\Program Files\IBM\TPC_SRA\".

-debug *MAX*

Optional parameter for debugging purposes.

-agentPort *agent_port*

If the agent is run as a daemon service, then the agent port must be specified.

-commtype *daemon*

If the agent is run as a daemon service, then this parameter must be specified.

-userID *user_ID*

For non-daemon service. The user ID defined on the agent system. This user ID is used by the server to connect to the agent system.

-password *password*

For non-daemon service. Password for the user ID.

-certFile "*certificate_file*"

For non-daemon service. The certificate used for SSH communication between the server and agent. This certificate must be stored on the server system. Enclose the directory name in quotation marks, for example, "c:\keys\id_sra".

For information about SSH protocol and SSL protocol certificates, see "Creating a certificate for SSH protocol" on page 286 and "Replacing a custom certificate for SSL protocol" on page 291.

-passphrase *passphrase*

For non-daemon service. The passphrase defined for the certificate used in SSH communication.

Installing the Storage Resource agent on the Virtual I/O Server

You can install the Storage Resource agent on the Virtual I/O Server.

Local installation

To install the Storage Resource agent on the Virtual I/O Server, complete the following steps:

1. Log in to the Virtual I/O Server using the **padmin** user ID.
2. Run the following command to set up the AIX environment:
`oem_setup_env`
3. Change to the directory where the Storage Resource agent installation image is located.

Note: You can either mount the Storage Resource agent installation DVD or FTP to the Storage Resource agent installation image to the padmin home directory.

4. Run the Storage Resource agent installation command. For information about the Storage Resource agent installation command, see “Installing the agent using a command” on page 169.

Deploying the Storage Resource agent remotely

You can deploy the Storage Resource agent remotely for the Virtual I/O Server from your local workstation.

Remote deployment

To deploy the Storage Resource agent remotely, complete the following steps:

1. Enable remote access for the **root** user ID. Follow these steps:
 - a. Connect to the Virtual I/O Server using ssh/telnet with the padmin user ID.
 - b. Run the following command to set up the AIX environment:
`oem_setup_env`
 - c. Run the following command:
`smitty login_user`
 - d. In the **User NAME** field, type **root** and press **Enter**.
 - e. In the **User can LOGIN REMOTELY(rsh,tn,rlogin)?** field, set the value to **true**.
2. Click **Save**.
3. Create an SSH certificate for the **root** user ID. For information about creating the SSH certificate for the **root** user, see “Creating an SSH certificate for the root user ID” on page 336.
4. Start the remote deployment.

Installing the Tivoli Storage Productivity Center Monitoring Agent

The Tivoli Storage Productivity Center Monitoring Agent can be used by the IBM Tivoli Enterprise Monitoring Server to monitor systems in your enterprise. This

agent is an optional program you can install and use in your enterprise. The Monitoring Agent uses an existing Tivoli Storage Productivity Center server in the environment to obtain information.

You will be installing the Tivoli Storage Productivity Center Monitoring Agent in the Tivoli Storage Productivity Center environment. You will also be installing the agent-specific support files needed by the Tivoli Enterprise Monitoring Server.

Installing and configuring the Tivoli Storage Productivity Center Monitoring Agent on Windows:

This topic describes how to install the Tivoli Storage Productivity Center Monitoring agent on Windows.

Installing the Monitoring Agent in the Tivoli Storage Productivity Center environment

This procedure describes how to install the Monitoring Agent in the Tivoli Storage Productivity Center environment. You can install the agent on the Tivoli Storage Productivity Center server or on another computer. This example assumes the following:

- The Tivoli Monitoring Server environment is installed and running.
- You are installing the agent on a computer that is different from the server on which Tivoli Storage Productivity Center is installed.

To install the Tivoli Storage Productivity Center Monitoring Agent on Windows:

1. Gather information required by the Tivoli Storage Productivity Center Monitoring agent installation wizard. You will need to know the host name and authentication password of the Tivoli Storage Productivity Center server.
2. Log on to the system with administrator authority on Windows.
If you are using the installation DVD, the files are located in the itm folder. If you are using the installation image, use the Disk 2 image.
3. In the installation directory and run **setup.exe**.
4. Proceed through the wizard.
5. On the User Data Encryption Key page, enter a data encryption key for encrypted communication.

The Tivoli Enterprise Monitoring server uses a data encryption key for Secure Socket (SSL) communication between the agent and server.

6. Write down the key so that you can configure the Tivoli Monitoring Enterprise server with the encryption key and click **Next** when you have specified the key.

Note: For information about the data encryption key, see **Step 4: Define the node's security environment** in http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/index.jsp?topic=%2Fcom.ibm.itcamfad.doc_7.1%2FABD003%2Fsection1.html.

7. On the Select Program Folder page, enter a Program Folder name or select a folder from the Existing Folders list, and click **Next**.
8. On the Start Copying Files page, review your current setting, and click **Next**. Wait for the installation to complete.
9. On the Tivoli Storage Productivity Center Agent Configuration page, in the navigation tree, you see these configuration steps:
 - For **Log Configuration**, enter the following information, and click **Next**:

Log File

Directory and file name where the Monitoring Agent log file is created. The default is \$ITM_INSTALL_DIR/tma\itm6/logs/TPCAgent.log..

Log Level

Specify the debugging level for the log file.

- For **TPC Configuration** , enter this information, and click **Next**:

Address

IP address or host name of the Device server.

TPC Device Server Port

Port for connecting to the Device server. The default port is 9550.

TPC Host Authentication Password

Password to use to connect to the Device server.

Confirm TPC Host Authentication Password

Enter the password again.

- For the optional **Server and Data Sources Node**, click **New**, enter a node name, and click **Next**. This node name is internal to Tivoli Storage Productivity Center and will not be displayed in the Tivoli Storage Productivity Center user interface.
- For the optional **Computers and Hypervisors Node**, click **New**, enter a node name, and click **Next**. This node name is internal to Tivoli Storage Productivity Center and will not be displayed in the Tivoli Storage Productivity Center user interface.
- For the optional **Storage Systems Node**, click **New**, enter a node name, and click **Next**. This node name is internal to Tivoli Storage Productivity Center and will not be displayed in the Tivoli Storage Productivity Center user interface.
- For the optional **Fabrics and Switches Node**, click **New** to add a new data node, and enter a node name, and click **OK**. This node name is internal to Tivoli Storage Productivity Center will not be displayed in the Tivoli Storage Productivity Center user interface.

Note: You can specify information for the optional nodes if you want Tivoli Storage Productivity Center to collect information about those nodes.

Note: If you configure more than one Monitoring Agent to collect information for these optional nodes, use unique node names for the different Monitoring Agents. Using identical names for the same node on more than one agent can result in Tivoli Monitoring Server displaying information for only the nodes on the most recently refreshed Monitoring Agent. For example, do not use the same node name for **Server and Data Sources Node** on more than one Monitoring Agent.

10. If the TEMS connection panel is not displayed, open the Manage Tivoli Enterprise Monitoring Services page.
11. Select **Monitoring Agent for TPC**, right-click **Monitoring Agent for TPC**, and click **Reconfigure**.
12. On the Configure Monitoring Agent for TPC, for the TEMS Connection page, enter this information for the Protocol 1 tab:

TEMS Hostname

Host name of the Tivoli Enterprise Monitoring Server.

Port Number

The default port is 1918.

Protocol

Enter the protocol type. For IP.PIPE settings, enter the port number.

For information about the TEMS connection parameters, see <http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.omegamon.mes.doc/zosconfigguide601263.htm>. Your configuration is now complete.

13. To enable the Eclipse Help Server for Tivoli Monitoring Services, configure the Eclipse Help Server.
14. Specify port 80 and restart TEMS and TEPS if necessary.

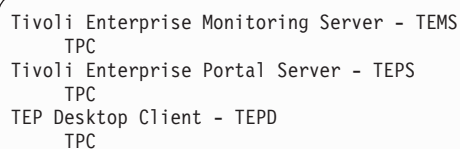
Installing the Monitoring Agent support files in the IBM Tivoli Monitoring Server environment

Tivoli Storage Productivity Center has support files that are installed on each of these Tivoli Monitoring components:

- Tivoli Enterprise Monitoring Server (TEMS)
- Tivoli Enterprise Portal Server (TEPS)
- Tivoli Enterprise Portal Desktop Client (TEPDC)

To install the Monitoring Agent in the IBM Tivoli Monitoring Server environment:

1. Log on to the IBM Tivoli Monitoring system with administrator authority on Windows.
 - If you are using the installation DVD, the files are located in the itm folder.
 - If you are using the installation image, use the Disk 1 part 4 image.
2. Go to the installation directory and run **setup.exe**.
3. Proceed through the wizard.
4. On the Select Features page, select the features you want to install:
 - a. Expand **Tivoli Enterprise Monitoring Agents - TEMA**, ensuring that you expand all nodes and select all checkboxes.



```
Tivoli Enterprise Monitoring Server - TEMS
  TPC
Tivoli Enterprise Portal Server - TEPS
  TPC
TEP Desktop Client - TEPD
  TPC
```

- b. Click **Next**
5. In the IBM Tivoli Productivity Center Monitoring - InstallShield Wizard, select all the components related to Tivoli Storage Productivity Center:

Note: This panel is displayed when the Monitoring Agent is installed on the Tivoli Monitoring Server.

```

Tivoli Enterprise Monitoring Agents - TEMA
  Tivoli Enterprise Monitoring Agent Framework
  TPC
  Tivoli Enterprise Services User Interface Extensions
Tivoli Enterprise Monitoring Server - TEMS
  TPC
Tivoli Enterprise Portal Server - TEPS
  TPC
TEP Desktop Client - TEPD
  TPC

```

6. Ensure that you expand all the nodes and select all the checkboxes and click **Next**.
7. Select the appropriate program folder in which the agent will be installed and click **Next**.
8. On the summary page, review the summary information and click **Next**.
9. After installation is complete, on the Configuration Defaults for Connecting to a TEMS page, you can enter connection information for the Tivoli Monitoring server and the Tivoli Storage Productivity Center Device server.

Note: This panel is displayed when the Monitoring Agent is installed on the Tivoli Monitoring server.

10. Enter this information for the Protocol 1 tab and click **OK**:

TEMS Hostname

Host name of the Tivoli Enterprise Monitoring Server.

Protocol

Enter the protocol type. For IP.PIPE settings, enter the port number.

For information about the TEMS connection parameters, go to <http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.omegamon.mes.doc/zosconfigguide601263.htm>

Your configuration is now complete. To see the reports that were generated, go to the Tivoli Storage Productivity Center Information Center and search for **monitoring agent reports**.

Installing and configuring the Tivoli Storage Productivity Center Monitoring Agent on AIX or Linux:

This topic describes how to install the Tivoli Storage Productivity Center Monitoring Agent on AIX or Linux.

Installing the Monitoring Agent and support files in the Tivoli Storage Productivity Center environment

This example assumes that the Tivoli Monitoring environment is installed and running and that you are installing the agent in the Tivoli Storage Productivity Center environment.

Tivoli Storage Productivity Center install support files on each of these Tivoli Monitoring components:

- Tivoli Enterprise Monitoring Server (TEMS)
- Tivoli Enterprise Portal Server (TEPS)
- Tivoli Enterprise Portal Desktop Client (TEPDC)

Note:

- IBM Tivoli Monitor does not support the Tivoli Enterprise Portal Desktop Client (TEPDC) on AIX.
- If you have the Tivoli Enterprise Portal Browser Client installed, you must also install the Monitoring Agent support files on the Browser Client.

To install the Tivoli Storage Productivity Center Monitoring agent on AIX or Linux:

1. Gather the following information required by the Tivoli Storage Productivity Center Monitoring agent installation wizard, such as the host name and IP address of the Tivoli Enterprise Monitoring server.
2. Log on to the system with the root user ID or a user ID that has root access rights.
 - If you are using the installation DVD, the files are located in the itm folder.
 - If you are using the installation image, use the Disk 1 part 4 image.
3. Go to the installation directory and run **./install.sh**.
4. Here is an example of the prompts and responses you will see. This example assumes that you do not have a directory and are creating a directory and are installing the agent on a local host.

```
UPGRADE

Enter the name of the IBM Tivoli Monitoring directory
[ default = /opt/IBM/ITM ]:

"/opt/IBM/ITM" does not exist
Try to create it [ 1=yes, 2=no; "1" is default ]? 1

Select one of the following:

1) Install products to the local host.
2) Install products to depot for remote deployment (requires TEMS).
3) Install TEMS support for remote seeding
4) Exit install.

Please enter a valid number: 1
```

5. This example is a continuation from the previous step and lets you read the license agreement.

```
International Program License Agreement

...
...
```

6. Read and agree with the license agreement, make a selection, and press **Enter**.
7. This example is a continuation from the previous step and lets you enter an encryption key for a secure connection between the monitoring agent and the Tivoli Enterprise Monitoring Server.

```
Enter a 32-character encryption key, or just press Enter to use the default
Default = IBM Tivoli Monitoring EncryptionKey

GSKit encryption key has been set.
Key File directory: /opt/IBM/ITM/keyfiles
```

Note: For information about the encryption key, go to http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/index.jsp?topic=%2Fcom.ibm.itcamfad.doc_7.1%2FABD003%2Finstall_unix.html and see **Step 5: Provide an encryption key.**

8. Enter the encryption key and press **Enter**.
9. Use the default key or enter a new key.
10. Write down and save the key so that you can configure the Tivoli Enterprise Monitoring Server later and press **Enter**.
11. This example is a continuation from the previous step and lets you select the product packages to install. Make a selection and press **Enter**. For an agent only installation, select 1.

Product packages are available for this operating system and component support categories:

- 1) IBM Tivoli Monitoring components for this operating system
- 2) Tivoli Enterprise Portal Browser Client support
- 3) Tivoli Enterprise Portal Desktop Client support
- 4) Tivoli Enterprise Portal Server support
- 5) Tivoli Enterprise Monitoring Server support
- 6) Other operating systems

12. This example is a continuation from the previous step and lets you select the product to install:

- a. Enter **3** (for Tivoli Enterprise Portal Desktop Client support) on the Product packages panel and press **Enter**.

The following products are available for installation:

- 1) Monitoring Agent for TPC V06.22.00.00
- 2) All of the above

Select **2** and press **Enter**. Wait for this step to complete.

- b. Enter **4** (for Tivoli Enterprise Portal Server support) on the Product packages panel and press **Enter**.

The following products are available for installation:

- 1) Monitoring Agent for TPC V06.22.00.00
- 2) All of the above

Select **2** and press **Enter**. Wait for this step to complete.

- c. Enter **5** (for Tivoli Enterprise Monitoring Server support) on the Product packages panel and press **Enter**.

The following products are available for installation:

- 1) Monitoring Agent for TPC V06.22.00.00
- 2) All of the above

Select **2** and press **Enter**. Wait for this step to complete.

13. You see this prompt. Enter **2** and press **Enter**.

Do you want to install additional products or product support packages?

You see the following message:

Installation step complete.

14. You see this prompt.

You may now configure any locally installed IBM Tivoli Monitoring product via the "/opt/IBM/ITM/bin/itmcmd config" command.

Go to the following directory:

```
cd /opt/IBM/ITM/bin
```

Run this command:

```
./itmcmd manage &
```

15. The Manage Tivoli Enterprise Monitoring Services window is displayed. Select **Monitoring Agent for TPC**. Right-click **Monitoring Agent for TPC** and click **Configure**.
16. The Tivoli Storage Productivity Center Agent Configuration panel is displayed. In the navigation tree, you see these configuration steps:
 - For Log Configuration, enter the following information, and click **Next**:

Log File

Directory and file name where the Monitoring Agent log file is created. The default is \$ITM_INSTALL/li6263/p1/bin/logs/TPCAgent.log.

Log Level

Specify the debugging level for the log file.

- TPC Configuration, enter this information, and click **Next**:

Address

IP address or host name of the Device server.

TPC Device Server Port

Port for connecting to the Device server.

TPC Host Authentication Password

Password to use to connect to the Device server.

Confirm TPC Host Authentication Password

Enter the password again.

- For the optional **Server and Data Sources Node**, click **New**, enter a node name, and click **Next**. This internal Tivoli Storage Productivity Center node name will not be displayed in the Tivoli Storage Productivity Center user interface.
- For the optional **Computers and Hypervisors Node**, click **New**, enter a node name, and click **Next**. This internal Tivoli Storage Productivity Center node name will not be displayed in the Tivoli Storage Productivity Center user interface.
- For the optional **Storage Systems Node**, click **New**, enter a node name, and click **Next**. This internal Tivoli Storage Productivity Center node name will not be displayed in the Tivoli Storage Productivity Center user interface.
- For the optional **Fabrics and Switches Node**, click **New**, enter a node name, and click **OK**. This internal Tivoli Storage Productivity Center node name will not be displayed in the Tivoli Storage Productivity Center user interface.

Note: You can specify information for the optional nodes if you want Tivoli Storage Productivity Center to collect information about those nodes.

Note: If you configure more than one Monitoring Agent to collect information for these optional nodes, use unique node names for the different Monitoring Agents. Using identical names for the same node on more than one agent can result in Tivoli Monitoring Server displaying information for only the nodes on the most recently refreshed Monitoring Agent. For example, do not use the same node name for **Server and Data Sources Node** on more than one Monitoring Agent.

17. Open the Manage Tivoli Enterprise Monitoring Services window and select **Monitoring Agent for TPC**.
18. Right-click **Monitoring Agent for TPC** and click **Configure**.
19. On the Configure Monitoring Agent for TPC page, for the TEMS Connection page, enter this information for the Protocol 1 tab:

TEMS Hostname

Host name of the Tivoli Enterprise Monitoring Server.

Protocol

Enter the protocol type. For IP.PIPE settings, enter the port number.

For information about the TEMS connection parameters, go to <http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.omegamon.mes.doc/zosconfigguide601263.htm>.

Your configuration is now complete.

Note: If you need to reconfigure the **Tivoli Enterprise Portal Server**, **Tivoli Enterprise Monitoring Server**, and **Tivoli Enterprise Portal Desktop Client**, do the following:

- a. Open the Manage Tivoli Enterprise Monitoring Services console by running the command:
`itmcmd manage`
 - b. Right-click **Tivoli Enterprise Portal Server** and click **Rebuild Configuration**.
 - c. Right-click **Tivoli Enterprise Monitoring Server** and click **Rebuild Configuration**.
 - d. Right-click **Tivoli Enterprise Portal Desktop Client** and click **Rebuild Configuration**.
20. To enable the Eclipse Help Server for Tivoli Monitoring Services, open Manage Tivoli Monitoring Services and configure the Eclipse Help Server.
 21. Specify port 80 be used and restart TEMS and TEPS if necessary.

Installing DB2

This topic describes how to install DB2 Version 9.7 for operating systems such as Linux, AIX, or Windows.

Note: If you are installing DB2 9.7, see “Software requirements for the database repository” on page 12 for the versions supported.

Preparing to install DB2

Before you install DB2, it is important to understand that default the user name and groups are created. The DB2 administrator user name and password are used to install IBM Tivoli Storage Productivity Center. You can use the default user name and password or provide your own.

Note: Windows platforms do not support systems that are configured only for IPv6. (Windows is enabled for IPv6 and IPv4.)

Preparing to install DB2 on Windows

If you have an administrator user ID that you want to use to install DB2 and IBM Tivoli Storage Productivity Center, this user ID must be a member of the DB2ADMNS and the Administrator groups.

When you install DB2, these groups are created:

- DB2ADMNS, which has all the required user rights assigned for administrative authority.
- DB2USERS, which has user rights assigned to DB2 users.

Assign users who need administrative authority to the DB2ADMNS group and DB2 users to the DB2USERS group. When you assign a user ID to either group, the user ID has all the user rights required for that group. You do not have to assign individual user rights to each user ID.

Important:

- The Windows user ID *Administrator* does not have the authority to create the Tivoli Storage Productivity Center database TPCDB. The DB2 administrator user ID that you need when you install IBM Tivoli Storage Productivity Center must have DB2 administrative rights. The user ID needed to install IBM Tivoli Storage Productivity Center must be a member of the DB2ADMNS and Administrators groups. The DB2 user ID *db2admin*, which was created when you installed DB2, has the authority for Tivoli Storage Productivity Center. For valid DB2 user IDs and passwords, see <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.admin.dboj.doc/doc/c0007248.html>.

Note: When you assign the Windows user ID *Administrator* to the DB2ADMNS group, this authority is active the next time the *Administrator* user ID logs in.

Preparing to install DB2 on UNIX or Linux

To operate DB2 on operating systems such as UNIX or Linux, the instance owner, the fenced user, and the DB2 administration server user are required. These users and groups are automatically created when DB2 is installed.

The user ID you need to install IBM Tivoli Storage Productivity Center must be the owner of the instance that you want to use. This user ID and password are created when you install DB2.

Table 31. DB2 Setup wizard default user and group names

Required user	Default user name	Default group name	Description
instance owner	db2inst1	db2iadm1	Created in the instance owner home directory. This user ID controls all DB2 processes and owns all file systems and devices used by the databases contained within the instance.

Table 31. DB2 Setup wizard default user and group names (continued)

Required user	Default user name	Default group name	Description
fenced user	db2fenc1	db2fadm1	Used to run user-defined functions (UDFs) and stored procedures outside of the address space used by the DB2 database.
DB2 administration server user	dasusr1	dasadm1	Used to run the DB2 administration server on your system. Used by the DB2 GUI tools to perform administration tasks. This user does not contain databases and there is only one administration server per machine.

Note: For valid DB2 user IDs and passwords, see <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.admin.dboobj.doc/doc/c0007248.html>.

Installing DB2 on a Windows operating system

Before you can install Tivoli Storage Productivity Center, you must install DB2 and register the DB2 license key.

To install DB2 on a Windows operating system, complete the following steps:

1. Log on with a user ID that has Administrator authority on Windows.
2. Insert the DB2 product DVD into the DVD drive or use the extracted DB2 installation image. Windows Autorun starts the Launchpad.

If you are using Windows Autorun, the installation program should start in 15-30 seconds. If the installation program does not start, complete one of the following steps:

- In a command prompt, to start DB2 set up page, navigate to the DVD-ROM\server directory or the *web image extraction*\server folder and run the **setup.exe** command.
 - In Windows Explorer, navigate to the DVD-ROM\server directory or the *web image extraction*\server folder, and double-click the setup.exe file.
3. On the DB2 Setup Launchpad page, click **Install a Product**.
 4. In the content pane, click **Install New for DB2 Enterprise Server Edition Version 9.7**.
 5. In the Welcome to the DB2 Setup wizard for DB2 Enterprise Server Edition, Version 9.7 panel, click **Next**.
 6. In the Software License Agreement page, read, accept the license agreement, and click **Next**.
 7. In the Select the installation type page, select **Typical** and click **Next**.
 8. In Select installation, response file creation, or both page, select **Install DB2 Enterprise Server Edition on this computer and save my settings in a response file**.

9. Enter a response file name or accept the default and click **Next**.
10. In the Select the installation folder page, enter a directory or use the default and click **Next**.
11. In the Set user information for the DB2 Administration Server page, enter the following information for User information:
 - a. Leave **Domain** blank.
 - b. In User name, type a user name of the DB2 user account that you want to use.
DB2 adds this user ID to the DB2ADMNS group.
 - c. Type your password and retype to confirm it.
12. Select the **Use the same user name and password for the remaining DB2 services** checkbox and click **Next**.
The user ID and password are used to install Tivoli Storage Productivity Center.
13. In the Configure DB2 instances page, click **Next**.

Note: This panel appears if you are installing DB2 on a computer that belongs to an accessible domain.

14. In the Prepare the DB2 tools catalog page, accept the defaults, click **Next**.
The default is to not prepare the DB2 tools catalog.
15. In the Set up notifications page, clear the **Set up your DB2 server to send notifications** checkbox, and click **Next**.
16. In the Enable operating system security for DB2 objects page, accept the defaults, and click **Next**.
The default is to enable operating system security.

Note: If you have installed DB2 before on this system, and the DB2ADMS group exists, when you click **Next** in the security page, the following message is displayed: :

```
Warning
The group; name "DB2ADMS" already exists
in the system and will be granted complete
access to the DB2 folders and objects through
the operating system.
Click OK to continue installation, or click
Cancel to input another group name.
```

17. In the Start copying files and create response file page, review the summary information, and click **Finish**.
18. Read the information about the page and click **Next**.
19. In the Install additional products page, do not install additional products, and click **Finish**.
20. The setup wizard closes and you return to the DB2 Welcome screen that now has "Welcome to First Steps" in the left pane. Click **Product Updates** to see whether there are any product updates.
21. (Optional) Click **Product Updates** and install the necessary updates and click **Exit**.

Attention: In the dialog box that is displayed, do not upgrade to DB2 9.5 fix pack 6 or DB2 9.7 fix pack 5. Tivoli Storage Productivity Center does not support DB2 9.5 fix pack 6 or DB2 9.7 fix pack 5 and installing these versions can create installation errors.

22. On the Tivoli Storage Productivity Center download image or DVD package, find the `db2ese_o.lic` file, which is the DB2 Enterprise Server Edition, Version 9.7 license key.
23. In the `/db2/license/Windows` directory, read and agree to the license conditions for DB2.
24. Register the license key in one of the following ways:

From the command line:

Enter the `db2licm` command:

```
db2licm -a x:\db2\license\license_filename.lic
```

where `x` represents the DVD drive that contains the DB2 license key file and `license_filename.lic` is the license key file: `db2ese_o.lic`.

From the License Center:

To register the license key:

- a. Start the DB2 Control Center and click **Tools > License Center**.
 - b. Select the system and installed product for DB2 and click **License > Add**.
 - c. In the Add License window, select the directory that contains the license file.
 - d. Select the license file and click **OK**.
25. Log off and log on to the system before you install Tivoli Storage Productivity Center.

Installing DB2 on UNIX or Linux - GUI install

This topic describes how to install DB2 on operating systems such as UNIX or Linux using a GUI installation program.

Note: Before you install DB2 using the GUI, install the X11 graphic capability.

To install DB2 on an operating system such as UNIX or Linux, complete the following steps:

1. Log in as a user ID with root authority.
2. Do one of the following:
 - If you are not using a DVD, skip to step 6 on page 185 and run `db2setup` from the directory where you have the DB2 source installation program.
 - If you are installing a DVD, create a mount point called `/cdrom` by entering `mkdir /cdrom`.

Note: If the installation program is compressed, you must use an extraction tool to decompress the downloaded file. See step 6 on page 185 and run `db2setup` that is located in the server directory. You can also select an existing mount point.

3. Insert the DB2 DVD and mount the DVD file system at the wanted mount point.
4. Run one of the following commands to add an entry to `/etc/filesystems` for the DVD mount point:

AIX

```
/usr/sbin/crfs -v cdrfs -p ro -d'cd0' -m'/cdrom' -A'no'
mount /cdrom
```


Linux

```
mkdir /cdrom
mount -o ro /dev/cdrom /cdrom
```

The **crfs** command only has to be run once for a given mount point. For each DVD you insert later, you can use **mount** and **umount**, for example, **mount /cdrom** and **umount /cdrom**.

5. Enter **cd /cdrom** to access the directory where the DVD is mounted.
6. Enter **./db2setup** to start the DB2 Setup wizard.
7. Enter **cd /cdrom/server** and then **./db2setup** to start the DB2 Setup wizard.
8. On the DB2 Setup Launchpad page, click **Install a Product**.
9. On the content page, click **Install New for DB2 Enterprise Server Edition Version 9.7 Fix Pack 4**.
10. On the Welcome to the DB2 Setup wizard page, click **Next**.
11. On the Software License Agreement page, read, click **Accept**, and click **Next**.
12. On the Select the installation type page, select **Typical** and click **Next**.
13. On the Select installation, response file creation, or both page, select **Install DB2 Enterprise Server Edition on this computer and save my settings in a response file**.
14. Enter a response file name, or accept the default, and click **Next**.
15. On the Select the installation directory page, enter a directory or use the default and click **Next**.
16. On the Install the IBM Tivoli System Automation for Multiplatforms (SA MP) page, select **Do not install SA MP**, and click **Next**.
17. On the Set user information for the DB2 Administration Server page, select **New user**, and enter the following information:
 - a. User name and password of the DB2 user account that you want to use. DB2 adds this user ID to the DB2ADMNS group, and if this user ID does not exist, DB2 creates it.
 - b. Group name
 - c. Password
 - d. Home directory
18. Select **Use default UID** and **Use default GID** and click **Next**. The DAS user is used to administer the DB2 database.
19. On the Set up a DB2 instance page, select **Create a DB2 instance**, and click **Next**.
20. On the Set up partitioning options for the DB2 instance page, select **Single partition instance**, and click **Next**.
21. On the Set user information for the DB2 instance owner page, select **New user**, and enter the following information: The DB2 instance owner user is the user you enter when installing Tivoli Storage Productivity Center.
 - User name
 - Group name
 - Password
 - Home directory
22. Select **Use default UID** and **Use default GID**, and click **Next**.
23. On the Set user information for the fenced user page, select **New user**, and enter the following information:
 - User name

- Group name
 - Password and confirm password
 - Home directory
24. Select **Use default UID** and **Use default GID**, and click **Next**.
 25. On the Prepare the DB2 tools catalog page, select **Do not prepare the DB2 tools catalog** and click **Next**.
 26. On the Set up notifications page, select **Do not set up your DB2 server to send notifications at this time**, and click **Next**.
 27. On the Start copying files and create response file page, review the summary information, and click **Finish**.
 28. Read the information about the panel and click **Finish**.
 29. From the Tivoli Storage Productivity Center download image or DVD package, find the db2ese_o.lic file for the DB2 Enterprise Server Edition, Version 9.7 license key.
 30. Ensure that you understand and agree to the license conditions for DB2.
 31. You can register the license key by using one of the following ways:
 - **From the command line:**
Enter the **db2licm** command:

```
db2licm -a x/server/db2/license/license_filename.lic
```

 where *x* represents the DVD drive that contains the DB2 license key file and *license_filename.lic* is the license key file *db2ese_o.lic*.
 - **From the License Center:**
 - a. Start the DB2 Control Center and click **Tools > License Center**.
 - b. Select the system and installed product for DB2 and click **License > Add**.
 - c. In the Add License window, select the directory that contains the license file.
 - d. Select the license file and click **OK**.

After you install DB2, edit the `/etc/group` file, add `root` to the `db2iadm1` group. The `db2iadm1` line in `/etc/group` looks like the following line: .

```
db2iadm1:x:102:root
```

You must edit the `/etc/group` file and add `db2inst1` to the group `root`. After adding `db2inst1`, the line should look like this line:

```
root:x:0:root,db2inst1
```

Installing DB2 on AIX using the command line

You can install DB2 on an operating system such as AIX using commands.

Ensure that you have:

- Two to three GB of hard disk drive space for the DB2 installation tar file and extracted files.
- Chosen a file system with at least 30 GB for the Tivoli Storage Productivity Center repository.

To install DB2 on the AIX operating system using the command line, complete the following steps:

1. Log in with a user ID that has root authority.

2. Do one of the following steps:
 - If you are not using a DVD, skip to step 5 and work from the directory where you have the DB2 source installation program.
 - If you are using a DVD, create a mount point called /cdrom by entering `mkdir /cdrom`.
3. Insert the DB2 DVD into the DVD drive and mount the file system at the mount point. On AIX, you can use the `crfs` command to add an entry to `/etc/filesystems` for the mount point.
4. Run the following commands:


```
/usr/sbin/crfs -v cdrfs -p ro -d'cd0' -m'/cdrom' -A'no'
mount /cdrom
```

The `crfs` command must be run only once for a mount point. For each DVD you insert after that, you can use `mount` and `umount`, for example, `mount /cdrom` and `umount /cdrom`.

5. Enter `cd /cdrom` and navigate to the directory where the DVD-ROM is mounted.
6. Install the DB2 Manager.
 - a. Create a temporary directory (for example, `db2temp`) to hold the DB2 installer tar file and extracted files.
 - b. Copy or download the DB2 installer into `db2temp`. The name of the file varies depending on the location from where the file is downloaded or copied and the language to which it is associated.
 - c. Optional: If the DB2 installer file is compressed, use the appropriate extracting tool. For example, if the name of the file is `v9.7fp4_aix64_server.tar.gz`, enter `gunzip v9.7fp4_aix64_server.tar.gz`. In this example, `v9.7fp4_aix64_server.tar.gz` is displayed in the `db2temp` directory.
 - d. Create an installation directory by entering, for example, `mkdir install`.
 - e. Enter `cd install` to change to the new installation directory.
 - f. Extract the DB2 installation file. For example, if the name of the installation file is `v9.7fp4_aix64_server.tar`, enter `tar xvf ../v9.7fp4_aix64_server.tar`.
 - g. Enter `cd /db2temp/installer/ese/disk1` and navigate to the directory that contains the DB2 installation program `db2_install`.
 - h. Enter `./db2_install` to run the command-line installation program.
 - i. Select **DB2.ESE**. The Installation Summary is displayed, which indicates a successful installation. DB2 is installed in `/opt/ibm/db2`.
7. Create users and groups for use with DB2.
 - a. Enter `mkgroup id=999 db2iadm1`.
 - b. Enter `mkgroup id=998 db2fadm1`.
 - c. Enter `mkgroup id=997 dasadm1`.
 - d. Enter `mkuser id=1004 pgrp=db2iadm1 groups=db2iadm1 home=/home/db2inst1 db2inst1`.
 - e. Enter `mkuser id=1003 pgrp=db2fadm1 groups=db2fadm1 home=/home/db2fenc1 db2fenc1`.
 - f. Enter `mkuser id=1002 pgrp=dasadm1 groups=dasadm1 home=/home/dasusr1 dasusr1`.
 - g. Verify the owner of the directories by entering `ls -ld /home/db2inst1`. The directory ownership is displayed as follows:

- /home/db2inst1 shows group db2iadm1 and user db2inst1
- /home/db2fenc1 shows group db2fadm1 and user db2fenc1
- /home/dasusr1 shows group dasadm1 and user dasusr1

If the directory ownership is not correct, run the following commands as appropriate:

```
chown -R db2inst1:db2iadm1 /home/db2inst1
chown -R db2fenc1:db2fadm1 /home/db2fenc1
chown -R dasusr1:dasadm1 /home/dasusr1
```

- h. Enter `passwd db2inst1 password`, where *password* represents the password you want to use for the DB2 instance.
 - i. Enter `pwdadm -f NOCHECK db2inst1`.
 - j. Enter `passwd db2fenc1 password`, where *password* represents the password you want to use for the fenced user.
 - k. Enter `pwdadm -f NOCHECK db2fenc1`.
 - l. Enter `passwd dasusr1 password`, where *password* represents the password you want to use for the DB2 administration server (DAS) user.
 - m. Enter `pwdadm -f NOCHECK dasusr1`.
 - n. Enter `chgroup users=db2inst1,root db2iadm1`.
8. Create a DB2 Administrative Server (DAS) and enter `/opt/db2/V9.5/instance/dascrt -u dasusr1`.
 9. Create a DB2 instance:
 - a. Enter `/opt/db2/V9.5/instance/db2icrt -a server -u db2fenc1 db2inst1`.
 - b. Enter `./home/db2inst1/sqllib/db2profile`.
 10. Change the default location for database repositories. By default, this location is /home/db2inst1. /home is typically not large enough for database repositories.
 11. To change the default location:
 - a. Enter `db2 update dbm cfg using DFTDBPATH new_repository_path IMMEDIATE`, where *new_repository_path* represents the new location for the repository.
 - b. Enter `chown -R db2inst1:db2iadm1 new_repository_path` to assign ownership to db2inst1 and permission to anyone in db2iadm1 (same as the ownership for /home/db2inst1).
 12. Configure DB2 communication:
 - a. Edit /etc/services and verify or add the following line at the end of the file: `db2c_db2inst1 50000/tcp`
 - b. Enter `db2 update dbm cfg using svcename db2c_db2inst1`.
 - c. Enter `db2set`. An example of the output is `DB2COMM=tcPIP`.
 13. From the Tivoli Storage Productivity Center download image or DVD package, find the `db2ese_o.lic` file for the DB2 Enterprise Server Edition, Version 9.7 license key.


```
db2licm -a x:\db2\license\license_filename.lic
```

where *x* represents the DVD drive that contains the DB2 license key file and *license_filename.lic* is the license key file: `db2ese_o.lic`.
 14. Read and agree to the DB2 license terms that are located in the /db2/license/Windows directory.
 15. Register the license key by entering the `db2licm -a x:\db2\license\license_filename.lic` command.
 16. Enter the following commands to restart DB2:

```
db2stop force
db2 terminate
db2start
```

17. In a command window, do the following to test your DB2 instance:
 - a. Enter `./home/db2inst1/sql1lib/db2profile` to source the environment.
 - b. Enter `db2level` to get information about the instance and DB2.
 - c. Enter `db2 create db test` to create a test database.
 - d. Enter `db2 list db directory` to list information about all the databases created for this instance.
 - e. Enter `db2 connect to test user db2inst1 using password` to connect to the test database, where *password* is the password you defined in step 7 for the DB2 instance.
 - f. Enter `db2 disconnect test` to disconnect from the test database.
 - g. Enter `db2 drop db test` to drop the test database.

Installing DB2 on AIX 6.1 POWER7 Systems

This topic describes how to install DB2 on AIX 6.1 POWER7 Systems.

Ensure that you have run `export JAVA_COMPILER=none` from a command window. This command is necessary as a result of APAR IC63450.

To install DB2, complete the following steps:

For more information about installing DB2, see <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=/com.ibm.db2.luw.qb.server.doc/doc/t0006833.html>.

1. Log on using a user ID with root authority.
2. Do one of the following:
 - If you are not using a CD, you can skip to step 5 and run **db2setup** from the directory where you have the DB2 source installation program.
 - If you are using the CD, enter, for example, `mkdir /cdrom` to create a mount point called `/cdrom`,
3. Insert the DB2 CD into the CD drive.

On AIX, you can use the **crfs** command to add an entry to `/etc/filesystems` for the CD mount point.
4. Mount the CD file system at the wanted mount point.
5. Run the following commands:

```
/usr/sbin/crfs -v cdrfs -p ro -d'cd0' -m'/cdrom' -A'no' mount /cdrom
```

The **crfs** command only must be run once for a given mount point. For each CD or DVD you insert after that, you can use **mount** and **umount**, for example, `mount /cdrom` and `umount /cdrom`.
6. Enter `cd /cdrom` and change to the directory where the CD is mounted.
7. Run the following command from the installation image:

```
./db2_install
```

Enter the necessary information.

Note: If the following message is displayed: :
WARNING: A minor error occurred while installing

check the installation log for errors. If you do not see any errors, ignore this warning message.

8. Create DB2 groups, users, and instances.
 - a. For the DB2 groups, run the following commands:

```
mkgroup id=999 db2iadm1
mkgroup id=998 db2fadm1
mkgroup id=997 dasadm1
```
 - b. For the DB2 users, run the following commands:

```
mkuser id=1004 pgrp=db2iadm1 groups=db2iadm1
home=/home/db2inst1 db2inst1
mkuser id=1003 pgrp=db2fadm1 groups=db2fadm1
home=/home/db2fenc1 db2fenc1
mkuser id=1002 pgrp=dasadm1 groups=dasadm1
home=/home/dasusr1 dasusr1
```
 - c. Set the password for db2inst1, db2fenc1, and dasusr1 by running the following commands:

```
passwd db2inst1
passwd db2fenc1
passwd dasusr1
```
 - d. Setting the new passwords might require that you log in once to change the initial password.

```
login db2inst1
```

Change the initial password and replace with a new password. Repeat this step for db2fenc1 and dasusr1 users.
 - e. Create the db2inst1 instance by running the following command:

```
/opt/IBM/db2/V9.7/instance/db2icrt -a server -u db2fenc1 db2inst1
```
 - f. Create the dasusr1 instance by running the following command:

```
/opt/IBM/db2/V9.7/instance/dascrt dasusr1
```
9. Assign a communication port for the DB2 instance in the `/etc/services` file.
 - a. Select a name that is related to the instance name.
 - b. Select a free port, for example, 50000. Add the following line to the `/etc/services` file:

```
db2inst1c 50000/tcp
```
10. Configure the communication port in the database manager configuration for the instance:

```
db2 update dbm cfg using SVCENAME db2inst1c
```
11. Enable TCP/IP for the instance:

```
db2set DB2COMM=TCPIP
```
12. From the Tivoli Storage Productivity Center download image or DVD package, find the `db2ese_o.lic` file for the DB2 Enterprise Server Edition, Version 9.7 license key:

```
db2licm -a x/db2/license/license_filename.lic
```

where `x` represents the DVD drive that contains the DB2 license key file and `license_filename.lic` is the license key file: `db2ese_o.lic`.
13. Restart the database instance:
14. Read and agree to the license conditions that are located in the `/db2/license/Windows` directory.
15. To register the license key, enter the **db2licm** command from the command line:

```
su - db2inst1
db2stop
db2start
```

Note: After you install DB2, edit the file `/etc/group` and add **root** to the `db2iadm1` group. The `db2iadm1` line in `/etc/group` looks like the following line:

```
db2iadm1:x:102:root
```

Verifying that DB2 is installed correctly

You can verify that DB2 has been installed properly using the command-line processor (CLP) or the First Steps GUI.

To verify that DB2 has been installed, complete the following steps:

1. Create the SAMPLE database.
2. Connect to the SAMPLE database.
3. Run a query against the SAMPLE database.
4. Drop the SAMPLE database.

Verifying DB2 installation using the command-line processor (CLP)

You can verify that DB2 has been successfully installed using the command-line processor (CLP).

To verify that DB2 is installed using the command line processor, complete the following steps:

1. Log on to the system as a user with SYSADM authority.
2. Enter `db2start` to start the database manager.

Note: In operating systems such as UNIX and Linux, you must source the `db2profile` before you run the `db2start` command. For more information sourcing the profile, go to the Tivoli Storage Productivity Center information center and search for **Using the command line on UNIX and Linux**.

3. Enter the `db2sampl` command to create the SAMPLE database. This command might take a few minutes to process. There is no completion message. When the command prompt returns, the process is complete. The SAMPLE database is automatically cataloged with the database alias SAMPLE when it is created.
4. Enter the following commands in a DB2 command window to connect to the SAMPLE database, retrieve a list of all the employees who work in Department 20, and reset the database connection:

```
db2 connect to sample
db2 "select * from staff where dept = 20"
db2 connect reset
```

5. After verifying the installation, remove the SAMPLE database to free up disk space. For more information about verifying DB2 installation, refer to the DB2 documentation for your operating system.
6. Enter `db2 drop database sample` to drop the SAMPLE database.

Verifying DB2 installation using First Steps

You can use First Steps to verify that DB2 has been installed successfully.

- Ensure that the domain user account you use to create the sample database has SYSADM or SYSCTRL authority.
- Install Control Center and the First Steps component.

First Steps are part of the getting started component grouping in the DB2 Setup wizard. It is installed as part of a typical installation or can be selected during a Custom installation.

To verify the DB2 installation using First Steps, complete the following steps:

1. Log on to the system with the user account that you want to use to verify the installation.
2. Start First Steps by doing one of the following:
 - On operating systems such as UNIX or Linux, type `db2fs`.
 - On Windows, type `db2fs.bat`.
3. On the First Steps launchpad, select **Create Sample Databases**.
4. On the Create Sample Databases page, select the databases you want to create and click **OK**. The DB2 UDB sample database is used to verify the installation. By default, the SAMPLE database is created on the computer where DB2 is installed. When the SAMPLE database has been created, you receive a completion message.
5. Click **OK**.
6. Open the DB2 Control Center. You can perform administration tasks on different instance and database objects through the Control Center.
7. Expand the Control Center object tree to view the SAMPLE database and SAMPLE database objects.
8. Select the Tables object to view the SAMPLE database tables in the content pane.
9. After verifying the installation, remove the SAMPLE database to free up disk space. For more information about verifying DB2 installation, refer to the DB2 documentation for your operating system.
10. Enter `db2 drop database sample` to drop the SAMPLE database.

Reinstalling the software if a failure occurs

If an installation failure occurs, you do not have to uninstall components that were successfully installed. Tivoli Storage Productivity Center provides an option to partially, or fully, roll back the installation.

A full rollback uninstalls all the components (even if they were installed successfully), and a partial rollback uninstalls only the components that were not successfully installed. The partial rollback option helps you resolve the problems that are causing the installation failure. After you resolve the problem, you can install the remaining components. For example, if the correct libraries were not installed for the Linux operating system, and the installation stops. You can install the correct libraries and resume the installation.

Tip: If you install Tivoli Storage Productivity Center on a Windows operating system and select **Full Rollback**, you must restart the server after the rollback is complete. The installation is complete only after the rollback errors are resolved.

If you originally selected a partial rollback, but decided to completely uninstall, you must run the uninstallation program before you reinstall Tivoli Storage Productivity Center.

The following table shows the results from a partial rollback. The components are listed in the order in which they were installed.

Table 32. Results from a partial rollback

If an installation failure occurs when you install this component...	Partial rollback results are...
Tivoli Storage Productivity Center common files and Java Runtime Environment	There is no partial rollback. The installation program performs a full rollback.
Tivoli Integrated Portal	Tivoli Integrated Portal is uninstalled.
Database repository	The database repository is uninstalled.
Data server	The Data server is uninstalled.
WebSphere Application Server	The WebSphere Application Server is uninstalled.
Device server	The Device server is uninstalled.
Replication server	The Replication server is uninstalled.
Storage Resource agent	The Storage Resource agent is uninstalled.
stand-alone GUI	The stand-alone GUI is uninstalled.
web-based GUI	The web-based GUI is uninstalled.
Command-line interface	The command-line interface is uninstalled.

Chapter 3. Configuring Tivoli Storage Productivity Center

After you installed IBM Tivoli Storage Productivity Center, you must configure it. You can use the Configuration Utility to learn how to configure and use Tivoli Storage Productivity Center. This section also describes each node under **Administrative Services > Configuration**.

The Data server hosts the control points for product scheduling functions, configuration, event information, reporting, and graphical user interface support. The Data server coordinates communication with the agents for data collection. The agents scan file systems and databases to gather storage demographics and populate the database with results. Automated actions can be defined to drive functions like data deletion, IBM Tivoli Storage Manager backup or archiving, or event reporting when defined thresholds are encountered. The Data server is the primary contact point for all user interface functions. The user interface functions also include actions that schedule data collection and discovery for the Device server.

The Device server discovers objects, gathers information from, analyzes performance of, and controls storage subsystems and SAN fabrics. The Device server coordinates communication with the agents to collect data about the SAN fabrics.

Starting Tivoli Storage Productivity Center

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

The stand-alone GUI and the web-based GUI provide different functions and offer different advantages. Start the one that is best suited to your current needs.

- The stand-alone GUI contains functions for monitoring the condition of storage resources, and all the tools for managing data, disk, fabric, tape, and replication. Start the stand-alone GUI if you must have the full administrative capabilities of Tivoli Storage Productivity Center. How you start the stand-alone GUI depends on whether the stand-alone GUI component is installed on your system.
- The web-based GUI provides quick access to pages that you can use to monitor the condition, capacity, and relationships of the resources within your storage environment. Start the web-based GUI if you are interested in monitoring your storage environment but do not require the full set of administrative tools available in the stand-alone GUI.

Starting the Tivoli Storage Productivity Center stand-alone GUI

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

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

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

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

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

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

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

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

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

In the preceding address, specify the following values:

hostname

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

port

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

2. The Tivoli Storage Productivity Center GUI requires an IBM Java Runtime Environment 1.6.0. Links are provided on the Java Web start page for you to download the IBM Java Runtime Environment 1.6.0 for Windows, Linux, or AIX. If an IBM Java Runtime Environment 1.6.0 is not already installed on the system, click the link to download the IBM Java Runtime Environment 1.6.0 for your operating system. On Windows, download a self-extracting executable. On Linux, download an RPM file. On AIX, download an executable JAR file. Install

the JRE from the file that you downloaded before you start the stand-alone GUI. If you are prompted to install the JRE as the system JRE or to overwrite the current system JRE, click **No**.

3. From the Java Web Start page, click **TPC GUI (Java Web Start)**

The JNLP file is downloaded.

Important: Depending on your browser and operating system and their default settings for your browser and operating system) you might need to:

- Verify to your browser that the `tpcgui.jnlp` file is safe to download
- Indicate that JNLP files are to be opened by the Java Web Start Launcher

For more information, go to the Tivoli Storage Productivity Center information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for *Cannot start the Tivoli Storage Productivity Center stand-alone GUI remotely*.

4. The stand-alone Tivoli Storage Productivity Center GUI opens in a separate window. To log on to Tivoli Storage Productivity Center, enter the following information and click **OK**.
 - a. Enter your user ID and password.

Important: If you are logging on immediately after installation and, see “Overview of required user IDs for initial log on to the Tivoli Storage Productivity Center family of products” on page 201 for information on the user ID you must use.

- b. If the server field does not contain the address and port number of for the server, specify the following values for the *hostname* and *port*.

hostname

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

port

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

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

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

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

1. Start the Tivoli Storage Productivity Center GUI by using the Start menu or batch file on a Windows operating system or a shell script on an AIX or Linux operating system.

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

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

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

hostname

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

port

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

- c. Click **OK**.

Starting the Tivoli Storage Productivity Center web-based GUI

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

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

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

1. On a server running the Windows operating system, click **Start > Programs > IBM Tivoli Storage Productivity Center > TPC Web-based GUI** to open a browser to the correct address. If you are not on a server running the Windows operating system, start a web browser and enter the following address in the address field:

```
http://hostname:port/srm
```

In the preceding address, specify the following values:

hostname

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

port

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

You might be redirected from the address that you enter to another address and port that provides secure access using the HTTPS protocol. This page is the Tivoli Storage Productivity Center logon page.

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

Starting Tivoli Storage Productivity Center for Replication

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

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

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

1. On a server running the Windows operating system, click **Start > Programs > IBM Tivoli Storage Productivity Center > TPC Replication Manager GUI** to open a browser to the correct address. If you are not on a server running the Windows operating system, start a web browser and enter the following address in the address field. The address is case-sensitive.

`https://hostname:port/CSM`

In the preceding address, specify the following values:

hostname

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

port

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

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

You can also find the port number in the *install_root/eWAS/profiles/profile_name/properties/portdef.props* file. The default port for the HTTP protocol is defined by the **WC_defaulthost_secure** property within the file.

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

Starting the Tivoli Storage Productivity Center family of products through Tivoli Integrated Portal

You can start the Tivoli Storage Productivity Center stand-alone GUI, the Tivoli Storage Productivity Center web-based GUI, and the Tivoli Storage Productivity Center for Replication web-based GUI within IBM Tivoli Integrated Portal. Tivoli Integrated Portal provides the advantages of the single sign-on capability and launch-in-context integration.

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

These instructions assume that you have previously started the Tivoli Storage Productivity Center stand-alone GUI, Tivoli Storage Productivity Center web-based GUI, or the Tivoli Storage Productivity Center for Replication web-based GUI outside of Tivoli Integrated Portal and that you already configured your system. For example, to start the stand-alone Tivoli Storage Productivity Center GUI by using Java Web Start, an IBM Java Runtime Environment 1.6.0 is required.

By logging on to Tivoli Integrated Portal, you eliminate the need to log on to Tivoli Storage Productivity Center or any other Tivoli application that uses the common portal framework. In addition to the convenience of single sign-on, Tivoli Integrated Portal provides launch-in-context integration. By using these functions, you can move between Tivoli Storage Productivity Center and another Tivoli application at specific points within the other application interface.

1. Start a web browser and enter the following address in the address field:

`http://hostname:port/ibm/console`

In the preceding address, specify the following values:

hostname

The Tivoli Integrated Portal server. You can specify the host name as an IP address or a Domain Name System (DNS) name.

port

The port number for Tivoli Integrated Portal. If you did not upgrade your system from an earlier version of the product, the default port number for connecting to Tivoli Integrated Portal using the HTTP protocol is 9568. However, this port number might be different for your site. For example, the port number might be different if you upgraded your system from an earlier version of the product or if the default port range was not accepted during installation. If the default port number does not work, see “Tivoli Integrated Portal port numbers” on page 35 for alternative default ports.

You might be redirected from the address you enter to another address and port that provides secure access using the HTTPS protocol. This page is the Tivoli Integrated Portal logon page.

2. From the Tivoli Integrated Portal logon page, type your Tivoli Storage Productivity Center ID and password and click **Log in**.
3. From the Tivoli Integrated Portal navigation tree, click **Tivoli Storage Productivity Center**.
4. From the Tivoli Storage Productivity Center portlets page, you can start the Tivoli Storage Productivity Center stand-alone GUI, the Tivoli Storage Productivity Center web-based GUI, or the Tivoli Storage Productivity Center for Replication web-based GUI. Select one of the following options:
 - **Start the stand-alone GUI for Tivoli Storage Productivity Center.**
 - **Start the web-based GUI for Tivoli Storage Productivity Center.**
 - **Start the web-based GUI for Tivoli Storage Productivity Center for Replication.**

Related tasks:

“Starting the Tivoli Storage Productivity Center stand-alone GUI by using Java Web Start” on page 196

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

“Starting the Tivoli Storage Productivity Center web-based GUI” on page 198

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

“Starting Tivoli Storage Productivity Center for Replication” on page 199

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

Overview of required user IDs for initial log on to the Tivoli Storage Productivity Center family of products

All the graphical user interfaces (GUIs) in the Tivoli Storage Productivity Center family require a user ID and password. If you are logging on to a GUI immediately after you installed the software, the user ID that you must use differs depending on the type of installation and the GUI interface.

Immediately after you install the software, you must use the user ID as described in the following tables. After you log on, you can assign roles for users in Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication. When users log on, their user roles determine their authorization level and the components that they can view in each GUI.

For example, if a user is in a group that is assigned to the Tivoli Storage Productivity Center Superuser role, the user has full access to all Tivoli Storage Productivity Center functions. However, if a user is in a group that is assigned to the Tape Administrator role, the user has access to only tape-related functions.

Required user ID for initial log on after installation on a single server when only the common user is defined

After you install the software on a single server, if you defined only the common user, and did not define a separate user for Tivoli Integrated Portal and Tivoli Common Reporting, the required user ID for initial log is the common user name that was defined for the installation of Tivoli Storage Productivity Center.

You must use this user ID to log on to Tivoli Storage Productivity Center stand-alone GUI, the Tivoli Storage Productivity Center web-based GUI, the Tivoli Storage Productivity Center for Replication web-based GUI, or the Tivoli Integrated Portal.

Required user ID for initial log on after installation on a single server when the common user and a separate user for Tivoli Integrated Portal and Tivoli Common Reporting are defined

After you install the software on a single server if you defined a common user and a separate user for Tivoli Integrated Portal and Tivoli Common Reporting, you must use the user ID as described in the following table.

Table 33. Required user ID for initial log on after installation on a single server when the common user and a separate user for Tivoli Integrated Portal and Tivoli Common Reporting are defined

To log on to:	Use this user ID:
Tivoli Storage Productivity Center stand-alone GUI	The common user name that was defined for the installation of Tivoli Storage Productivity Center
Tivoli Storage Productivity Center web-based GUI	The common user name that was defined for the installation of Tivoli Storage Productivity Center
Tivoli Storage Productivity Center for Replication web-based GUI	The common user name that was defined for the installation of Tivoli Storage Productivity Center
Tivoli Integrated Portal	The common user name that was defined for the installation of Tivoli Integrated Portal and Tivoli Common Reporting

Required user ID for initial log on after installation on multiple servers with a remote database schema

After you install the software on multiple servers with a remote database schema, you must use the user ID as described in the following table.

Table 34. Required user ID for initial log on after installation on multiple servers with a remote database schema

To log on to:	Use this user ID:
Tivoli Storage Productivity Center stand-alone GUI	The user name that was defined for the installation of Tivoli Storage Productivity Center
Tivoli Storage Productivity Center web-based GUI	The user name that was defined for the installation of Tivoli Storage Productivity Center
Tivoli Storage Productivity Center for Replication web-based GUI	The user name that was defined for the installation of Tivoli Storage Productivity Center
Tivoli Integrated Portal	The user name that was defined for the installation Tivoli Integrated Portal and Tivoli Common Reporting

Required user ID for initial log on after installation on multiple servers with remote Tivoli Integrated Portal and Tivoli Common Reporting

After you install the software on multiple servers with remote Tivoli Integrated Portal and Tivoli Common Reporting, you must use the user ID as described in the following table.

Table 35. Required user ID for initial log on after installation on multiple servers with remote Tivoli Integrated Portal and Tivoli Common Reporting

To log on to:	Use this user ID:
Tivoli Storage Productivity Center stand-alone GUI	The user name that was defined for the installation of Tivoli Storage Productivity Center. Because user authentication is conducted on the server where Tivoli Integrated Portal is installed, the user name and user group must exist in the operating system user and group registry on the Tivoli Integrated Portal server.
Tivoli Storage Productivity Center web-based GUI	The user name that was defined for the installation Tivoli Integrated Portal and Tivoli Common Reporting
Tivoli Storage Productivity Center for Replication web-based GUI	The user name that was defined for the installation of Tivoli Storage Productivity Center. The Tivoli Storage Productivity Center for Replication GUI detects only users and groups that are in the operating system user and group registry on the server where Tivoli Integrated Portal is installed. To provide access to a user or group in the Tivoli Storage Productivity Center for Replication GUI, the user or group must be in the operating system user and group registry on the Tivoli Integrated Portal server.
Tivoli Integrated Portal	The user name that was defined for the installation of Tivoli Integrated Portal and Tivoli Common Reporting

User authentication considerations for multiple server installations

These are installation considerations when you install DB2 and Tivoli Storage Productivity Center on one server, and IBM Tivoli Integrated Portal and IBM Tivoli Common Reporting on the second server.

Tivoli Storage Productivity Center authenticates users with Tivoli Integrated Portal. When the Tivoli Storage Productivity Center servers and Tivoli Integrated Portal are installed on separate servers, the user and group accounts on the Tivoli Integrated Portal server are used for authentication.

Tivoli Storage Productivity Center is installed using local OS repository. When the local OS repository is used for authentication, the user accounts on Tivoli Integrated Portal Server host must be used to log on to stand-alone GUI and to run command-line interface (CLI) commands.

Since the user accounts on the Tivoli Storage Productivity Center servers host cannot be used for authentication, in this scenario, you should change the user authentication from OS to LDAP repository. If you cannot change the user authentication to an LDAP repository, then both servers should have the same operating system. The user names, passwords, group names, and memberships used to administer Tivoli Storage Productivity Center are defined on both servers and are kept identical.

Changing the user authentication configuration

The Tivoli Storage Productivity Center installation program establishes a default authentication configuration using the federated repositories feature of the WebSphere Application Server. If you want to change the user authentication configuration, follow these procedures.

Within the federated repositories framework, the Tivoli Storage Productivity Center installation program creates two repositories:

File-based user repository

This repository contains the user "tpcFileRegistryUser" (This user password is the same as the common user password entered during the Tivoli Storage Productivity Center installation.)

Operating system repository

In the federated repositories framework, the Tivoli Storage Productivity Center installation program creates two repositories on all the following Tivoli Storage Productivity Center IBM® WebSphere® Application Servers:

- Tivoli Integrated Portal server
- Device server
- Replication server

The Tivoli Integrated Portal server is used as the primary WebSphere Application Server for user authentication in Tivoli Storage Productivity Center. If the Tivoli Integrated Portal server is down, the Device Server and Replication Server are used as the back up WebSphere Application Servers to perform the user authentication and allow the authorized local OS users to log in to the Tivoli Storage Productivity Center stand-alone GUI and Tivoli Storage Productivity Center for Replication web-based GUI. You can add an LDAP repository as a post-install task. This configuration is completed on the Tivoli Integrated Portal server. The LDAP repository configuration settings are not propagated to the Device and Replication

servers. So, if the Tivoli Integrated Portal server is down, then the authorized LDAP users cannot log in to Tivoli Storage Productivity Center stand-alone GUI and Tivoli Storage Productivity Center for Replication web-based GUI. The back up user authentication mechanism based on Device and Replication servers allows the authorized local OS users to log in to the Tivoli Storage Productivity Center stand-alone GUI and the Tivoli Storage Productivity Center for Replication web-based GUI.

This repository is referred to as the localOS repository and contains the users and groups that are managed by the local operating system. On computers that are members of a Windows domain, this local OS repository also contains the domain users and groups that are managed by the Windows domain, provided that the computer is correctly configured with the Windows domain.

Use of the Tivoli Storage Productivity Center single sign-on feature, especially with storage system element managers, is limited with this default authentication configuration. Storage system element managers do not support using the localOS repository or the file-based repository for single sign-on, even when the element manager is installed on the same system as Tivoli Storage Productivity Center.

The LDAP repositories supported by Tivoli Storage Productivity Center depend on the Tivoli Integrated Portal support. For more information about the LDAP repositories supported, see one of the following topics:

- For the Windows operating system: <http://www.ibm.com/support/docview.wss?rs=180&uid=swg27012421> .
- For the AIX operating system: <http://www.ibm.com/support/docview.wss?rs=180&uid=swg27012389>.
- For the Linux operating system: <http://www.ibm.com/support/docview.wss?rs=180&uid=swg27012415>.

Search for *LDAP Servers using Federated Repository Configuration*.

Whenever you change the user authentication configuration (by adding or removing an LDAP repository within the federated repositories framework), you must first back up the existing Tivoli Integrated Portal WebSphere configuration files. The WebSphere backup files provide a restoration point in the event of a failure when performing these procedures. You must also back up the WebSphere configuration files after you add an LDAP repository to the user authentication configuration and want to later change the LDAP authentication settings.

The WebSphere Application Server cannot resolve duplicated users or groups if these users or groups are present in more than one repository in the federated repositories framework. For example, an Administrator user exists in both the local OS and LDAP repository. You must ensure that the duplicated users (or groups) are not used during the configuration or to manage Tivoli Storage Productivity Center.

The following topics describe how to select and configure each user authentication method.

Adding an LDAP repository to the federated repositories

You can configure 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 add an LDAP repository to the federated repositories in Tivoli Storage Productivity Center, you must use Tivoli Integrated Portal.

If you installed Tivoli Storage Productivity Center on a correctly configured computer that is a member of a Windows domain, Active Directory users and groups also exist in the localOS repository. To determine whether the Active Directory users exist in the localOS repository, log on to Tivoli Integrated Portal and click **Users and Groups > Manage users**. You should not add the Active Directory repository if the Active Directory users exist. If you plan to add the Active Directory repository for performance reasons, you must remove the localOS repository from the federated repositories.

The WebSphere Application Server cannot resolve duplicated users or groups if these users or groups are present in more than one repository in the federated repositories framework. For example, an Administrator user exists in both the localOS and LDAP repository. You must ensure that the duplicated users (or groups) are not used during the configuration or to manage Tivoli Storage Productivity Center.

To add an LDAP repository to the federated repositories in Tivoli Storage Productivity Center, complete the following steps. Contact your LDAP server administrator for assistance.

1. Back up the Tivoli Integrated Portal WebSphere Application Server configuration files. The WebSphere Application Server configuration backup files provide a restoration point in the event of a failure when you complete this procedure. This procedure uses the variable name *TIP_installation_directory* to indicate where Tivoli Integrated Portal is located.
 - a. Back up the WebSphere Application Server configuration files.
For the AIX or Linux operating systems, run the following command:

```
TIP_installation_directory/profiles/TIPProfile/bin/  
backupConfig.sh
```


For the Windows operating systems, run the following command:

```
TIP_installation_directory\profiles\TIPProfile\bin\  
backupConfig.bat
```
 - b. Back up the `soap.client.props` file for Tivoli Integrated Portal.
For the Windows operating systems, use the **copy** command to back up this file:

```
TIP_installation_directory\profiles\TIPProfile\properties\  
soap.client.props
```


For the AIX or Linux operating systems, use the **cp** command to back up this file:

```
TIP_installation_directory/profiles/TIPProfile/properties/  
soap.client.props
```
 - c. Back up the IBM Cognos 8.ear directory.
For Windows operating systems, enter the **xcopy** command to back up this directory:

```
TIP_installation_directory\profiles\TIPProfile\installedApps\  
TIPCell\IBM Cognos 8.ear
```


For the AIX or Linux operating systems, enter the **cp** command to back up this directory:

```
TIP_installation_directory/profiles/TIPProfile/installedApps/  
TIPCell/IBM Cognos 8.ear
```

- d. The WebSphere Application Server **backupConfig** command generates a compressed file: `WebSphereConfig_YYYY-DD-MM.zip`. When you run the **backupConfig** command, the Tivoli Integrated Portal server will shut down. The Tivoli Integrated Portal server must be restarted after the **backupConfig** command is finished.

To start the Tivoli Integrated Portal server, follow these steps:

- 1) For Windows operating systems, click **Start > Administrative Tools > Services**. On the Windows Services page, right-click **Tivoli Integrated Portal - V2.2_TIPProfile_Port_xxxx** where *xxxx* indicates the port specified during installation and select **Start**.
- 2) For AIX or Linux operating systems, open a command prompt window and enter the following command:

```
TIP_installation_directory/profiles/TIPProfile/bin/  
startServer.sh
```

2. To restore the WebSphere Application Server and IBM Cognos configuration, complete the following steps:
 - a. Restore the WebSphere configuration by running the `restoreConfig.sh` or `restoreConfig.bat` command. The **restoreConfig** command is located here:

- For Windows operating systems:

```
TIP_installation_directory\profiles\TIPProfile\bin\  
restoreConfig.bat
```

- For AIX or Linux operating systems:

```
TIP_installation_directory/profiles/TIPProfile/bin/  
restoreConfig.sh
```

- b. Restore the IBM Cognos 8.ear directory that you backed up in step 1c. For the Windows operating systems, enter the **xcopy** command to replace the contents of this directory:

```
TIP_installation_directory\profiles\TIPProfile\installedApps\TIPCell\IBM Cognos 8.ear
```

For the AIX or Linux operating systems, use the **cp** command to replace the contents of this directory:

```
TIP_installation_directory/profiles/TIPProfile/installedApps/  
TIPCell/IBM Cognos 8.ear
```

3. Open an Internet Explorer or Mozilla Firefox web browser and enter the following information in the address bar:

```
http://hostname:port/ibm/console/logon.jsp
```

or

```
https://hostname:port/ibm/console/logon.jsp
```

Where *hostname* is the server that is running Tivoli Integrated Portal such as the server name or IP address and *port* is the port number for Tivoli Integrated Portal.

The port number differs depending on the protocol that you use (http or https) and the options that you selected during the installation of Tivoli Storage Productivity Center. For more information about port numbers, see “Tivoli Integrated Portal port numbers” on page 35.

You can also contact your Tivoli Storage Productivity Center administrator to verify the host name and port number.

4. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have Tivoli Integrated Portal administrator authorization.

5. In the Tivoli Integrated Portal navigation tree, click **Settings > WebSphere Administrative Console**.
6. In the Tivoli Integrated Portal window, click **Launch WebSphere administrative console**.
7. In the WebSphere Application Server Administrative Console navigation tree, click **Security > Global security**.
8. On the Global security page, in the **User account repository** section, click the **Configure** button next to the **Available realm definitions** list.

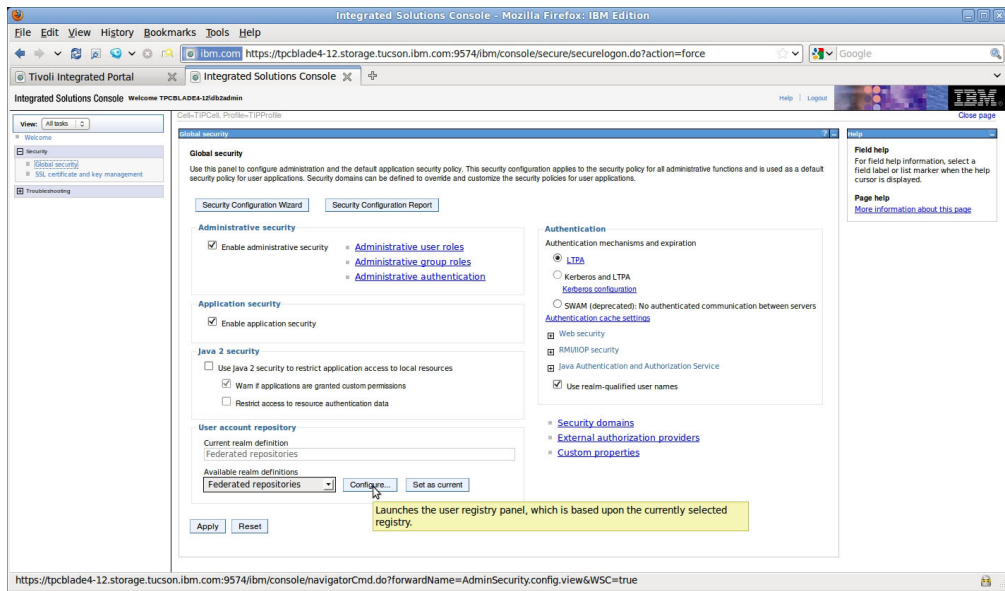


Figure 1. Global security page, configure federated repositories

9. On the Global security > Federated repositories page, under "Related Items," click **Manage repositories**.

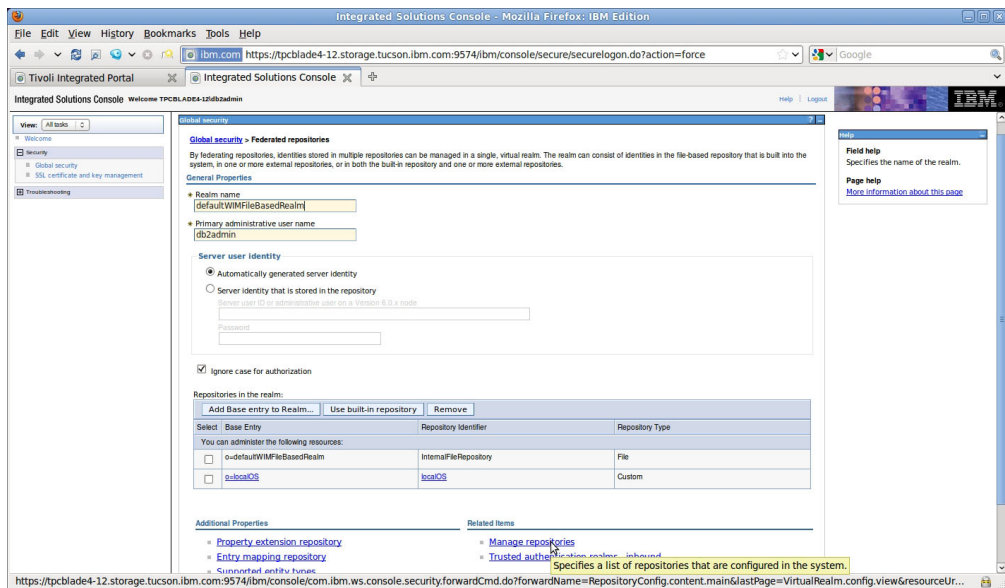


Figure 2. Federated repositories page, Manage repositories

In this window, you must ensure that the **Primary administrative user name** exists only in one repository. The WebSphere Application Server will not start if this user name exists in multiple repositories.

10. On the Global security > Federated repositories > Manage repositories page, add the LDAP repository that you want to use for authentication. Follow these steps:
 - a. Click **Add** to add a new repository.

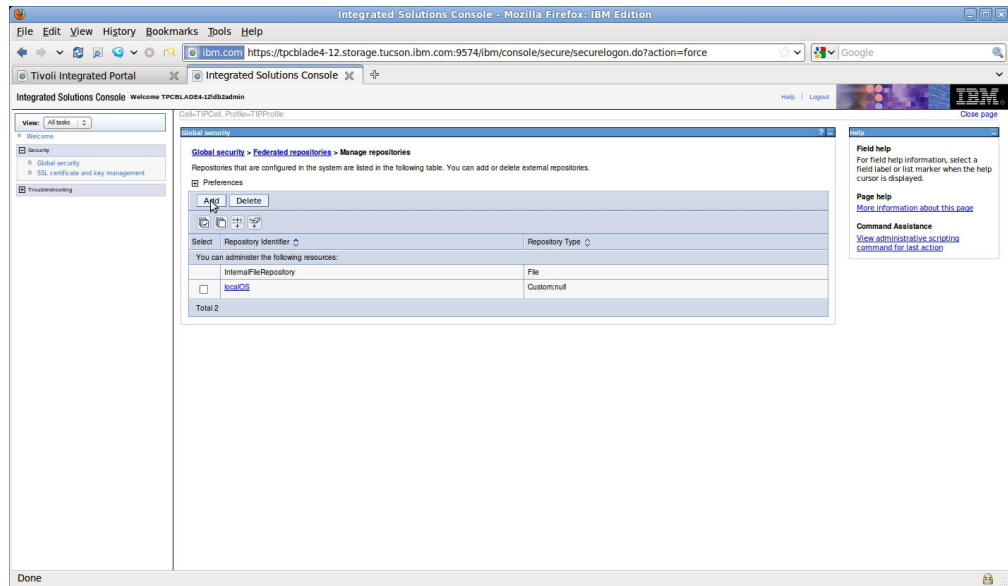


Figure 3. Manage repositories page, Add a new repository

- 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 realm, for example, LDAP1.

Directory type

The type of LDAP server that you want to connect to.

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. You can specify a different port depending on your LDAP server configuration. If you do not know what port to use, contact your LDAP server administrator.

Bind distinguished name

The distinguished name (DN) for WebSphere Application Server to use when binding to the LDAP repository. If no name is specified, WebSphere Application Server binds anonymously. In most cases, bind DN and bind password are required. However, when an anonymous bind can satisfy all the required functions, bind DN and bind password are not required.

Bind password

The password for WebSphere Application Server to use when binding to the LDAP repository.

Login properties

The authentication properties that are used to log on to WebSphere Application Server. Type uid;cn in this field. This value enables WebSphere Application Server to use the property that is required for the directory type.

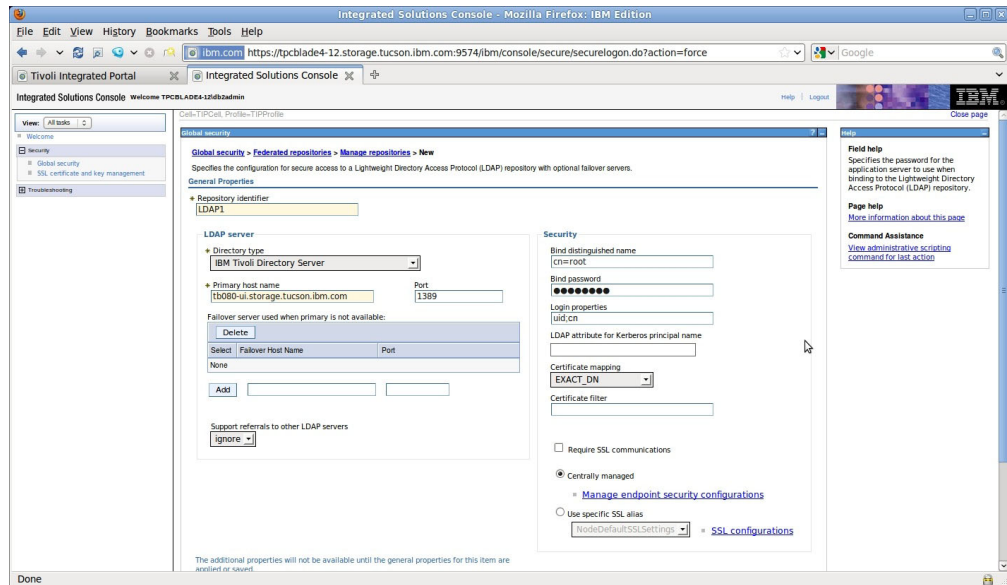


Figure 4. Manage repositories > New page

- c. Click **OK**.
- d. In the messages dialog box that is displayed on the Manage repositories page, click the **Save** link in **Save directly to the master configuration**.

Important: On the Global security > Federated repositories > Manage repositories page, do not delete the "localOS" repository.

11. From the Manage repositories page, return to the Global security > Federated repositories page.
12. Under "Repositories in the realm," click **Add Base entry to Realm**.

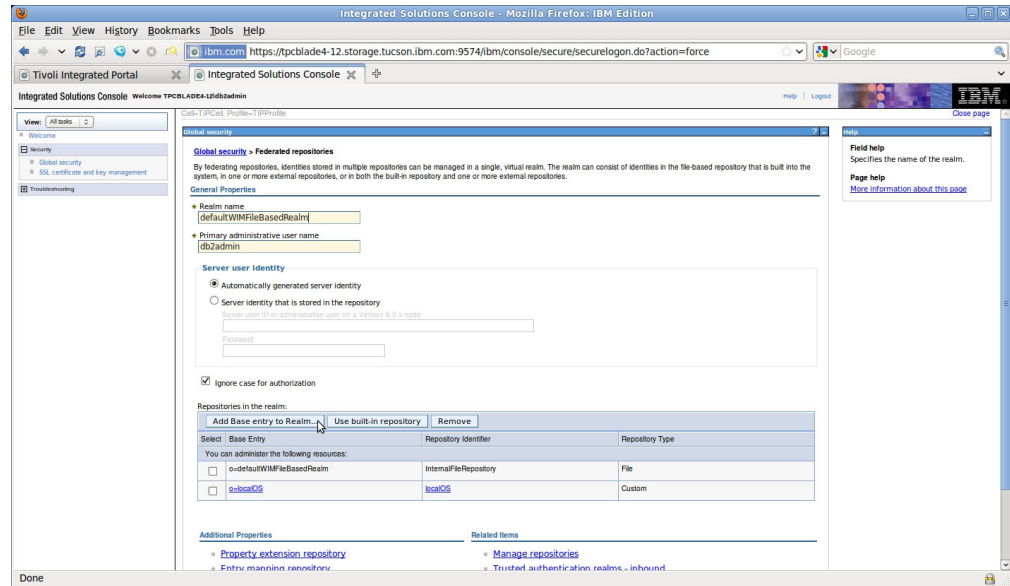


Figure 5. Global security > Federated repositories page, Add base entry to realm

13. If the Messages dialog box is displayed on the Repository reference page, click the **Save** link in **Save directly to the master configuration**.
14. On the Repository reference page, configure the following items:
 - a. In the Repository list, select the repository that you created in step 10 on page 209.
 - 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.
 - c. In the **Distinguished name of a base entry in this repository** field, enter the distinguished name of the base entry in the repository. In most instances, this value is 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. Click **OK**.
 - d. In the messages dialog box that is displayed, click the **Save** link in **Save directly to the master configuration**.

Important: In the **Repositories in the realm** table, do not remove the `InternalFileRepository` entry. Do not remove the `localOS` entry, unless you are planning to add an Active Directory repository to Tivoli Storage Productivity Center that is installed on a computer that is a member of a Windows domain and have verified that the Active Directory users exist in the `localOS` repository.




Figure 6. Global security > Federated repositories > Repository reference page, Add base entry to realm

15. Log out from the WebSphere Application Server Administrative Console.
16. Log out from the Tivoli Integrated Portal Administrative Console.
17. Stop and restart the Tivoli Integrated Portal server. For information about stopping and starting the Tivoli Integrated Portal server, see the Tivoli Storage Productivity Center information center. Search for *Starting and stopping the Tivoli Integrated Portal server*.

To verify that the LDAP federated repository is configured correctly, complete the following steps:

1. Log on to Tivoli Integrated Portal by using the same user ID and password from step 4 on page 207.
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 repositories. The list of users includes users from the local file repository, the operating system repository, and the LDAP repository.
5. In the Tivoli Integrated Portal navigation tree, click **Users and Groups > Manage Groups**.
6. In the Search by list, select **Group name**.
7. Click **Search** to search for groups in the federated repositories. The list of groups includes groups from the operating system repository and the LDAP repository.

Important: Before you establish the desired authorization configuration for Tivoli Integrated Portal, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication, you must ensure that there are no duplicated user names or group names within the local file repository, the operating system repository, and the LDAP repository.

After you have completed and verified the addition of an LDAP repository to the Tivoli Storage Productivity Center federated repositories authentication configuration, you must establish the desired authorization configuration for Tivoli Integrated Portal. Tivoli Integrated Portal administrators require the `iscadmins` and `administrator` roles. Operators viewing Tivoli Storage Productivity Center reports require the `tcrPortalOperator` role. For Tivoli Storage Productivity Center Version 5.1 single sign-on and launch in context features to work correctly, the LDAP user or group requires the following Tivoli Integrated Portal roles:

- administrator
- iscadmins
- tcrPortalOperator
- chartAdministrator
- monitor

For information about user roles and groups, go to the Tivoli Integrated Portal administrative console online help. On the Welcome page, open the online help by clicking `?`. On the Help welcome page, click **Using the Tivoli Integrated Portal > Users and Groups** in the navigation pane.

You should also assign Tivoli Storage Productivity Center roles to groups using the **Role-to-Group Mapping** node in the Tivoli Storage Productivity Center stand-alone GUI. You can map Tivoli Storage Productivity Center roles, such as Tape Operator or Fabric Administrator, to user groups that you create in your operating system repository and your LDAP repository. For information about the **Role-to-Group Mapping** node, go to the Tivoli Storage Productivity Center information center. Search for *Role-to-Group Mappings*.

For information about assigning roles to users and roles to groups for Tivoli Storage Productivity Center for Replication, go to the Tivoli Storage Productivity Center information center. Search for *Managing replication*. On the Managing replication page, click **Security**.

Removing an LDAP repository from the federated repositories

To remove an LDAP repository from the federated repositories, you must use Tivoli Integrated Portal.

If removing an LDAP repository from the federated repositories in Tivoli Storage Productivity Center leaves only the localOS repository and the file-based repository, then the use of the Tivoli Storage Productivity Center single sign-on feature is limited. Storage system element managers do not support using the localOS repository for single sign-on, even if the element manager is installed on the same system as Tivoli Storage Productivity Center. This procedure uses the variable name `TIP_installation_directory` to indicate where Tivoli Integrated Portal is located.

To remove an LDAP repository from the federated repositories, complete the following steps:

1. Back up the Tivoli Integrated Portal WebSphere configuration files. The WebSphere configuration backup files provide a restoration point in the event of a failure when performing this procedure. This procedure uses the variable name `TIP_installation_directory` to indicate where Tivoli Integrated Portal is located.
 - a. Back up the WebSphere configuration files.

For the AIX or Linux operating systems, run the following command:

```
TIP_installation_directory/profiles/TIPProfile/bin/  
backupConfig.sh
```

For the Windows operating systems, run the following command:

```
TIP_installation_directory\profiles\TIPProfile\bin\  
backupConfig.bat
```

- b. Back up the soap.client.props file for Tivoli Integrated Portal .

For the Windows operating systems, use the **copy** command to back up this file:

```
TIP_installation_directory\profiles\TIPProfile\properties\  
soap.client.props
```

For the AIX or Linux operating systems, use the **cp** command to back up this file:

```
TIP_installation_directory/profiles/TIPProfile/properties/  
soap.client.props
```

- c. Back up the IBM Cognos 8.ear directory.

- For Windows operating systems:

```
TIP_installation_directory\profiles\TIPProfile\installedApps\  
TIPCell\IBM Cognos 8.ear
```

- For AIX or Linux operating systems:

```
TIP_installation_directory/profiles/TIPProfile/installedApps/  
TIPCell/IBM Cognos 8.ear
```

- d. The WebSphere **backupConfig** command generates a compressed file: WebSphereConfig_YYYY-DD-MM.zip. When you run the **backupConfig** command, the Tivoli Integrated Portal server will shut down. The Tivoli Integrated Portal server must be restarted after the **backupConfig** command is finished.

To start the Tivoli Integrated Portal server, follow these steps:

- 1) For Windows operating systems, click **Start > Administrative Tools > Services**. On the Windows Services page, right-click **Tivoli Integrated Portal - V2.2_TIPProfile_Port_xxxx** where *xxxx* indicates the port specified during installation. Select **Start**.
- 2) For AIX or Linux operating systems, open a command prompt window and enter the following command:

```
TIP_installation_directory/profiles/TIPProfile/bin/  
startServer.sh
```

2. To restore the WebSphere Application Server and IBM Cognos configuration, complete the following steps:
- a. Restore the WebSphere configuration by running the restoreConfig.sh or restoreConfig.bat command. The **restoreConfig** command is located here:
- For Windows operating systems:

```
TIP_installation_directory\profiles\TIPProfile\bin\  
restoreConfig.bat
```
 - For AIX or Linux operating systems:

```
TIP_installation_directory/profiles/TIPProfile/bin/  
restoreConfig.sh
```
- b. Restore the IBM Cognos 8.ear directory that you backed up in step 1c.
- For the Windows operating systems, use the **xcopy** command to replace the contents of this directory:

```
TIP_installation_directory\profiles\TIPProfile
```
 - For the AIX or Linux operating systems, use the **cp** command to replace the contents of this directory:

TIP_installation_directory/profiles/TIPProfile/installedApps/
TIPCell/IBM Cognos 8.ear

3. Open an Internet Explorer or Mozilla Firefox web browser, and type the following information in the address bar:

`https://hostname:port/ibm/console/login.jsp`

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 range was accepted during the installation of Tivoli Integrated Portal, the port number is 9569. Contact your Tivoli Storage Productivity Center administrator if you must verify the host name and port number.

4. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator authorization.
5. In the Tivoli Integrated Portal navigation tree, click **Settings > WebSphere Administrative Console**.
6. In the Tivoli Integrated Portal content pane, click **Launch WebSphere administrative console**.

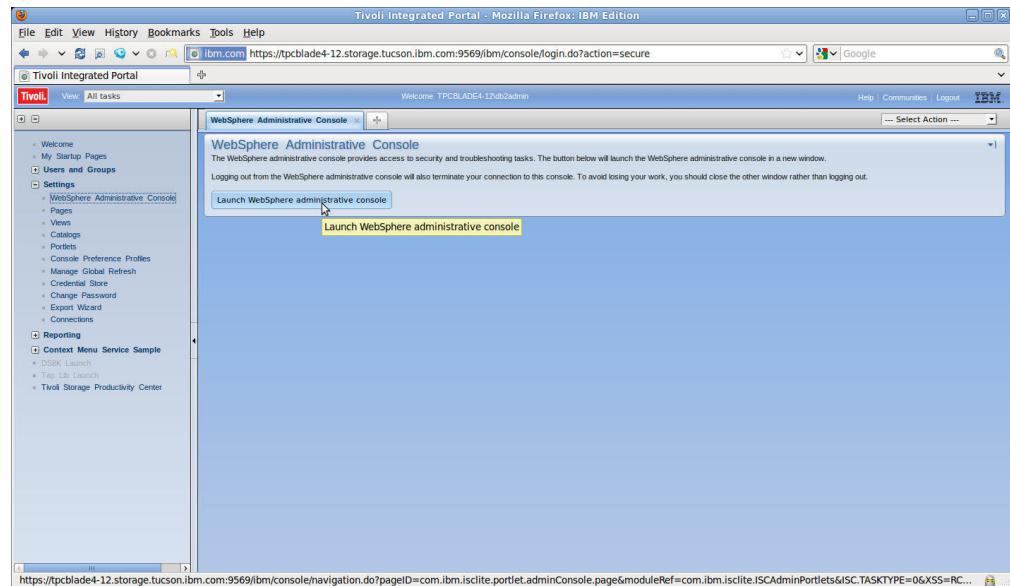


Figure 7. Launch the WebSphere administrative console

7. In the WebSphere Administrative Console navigation tree, click **Security > Global Security**.

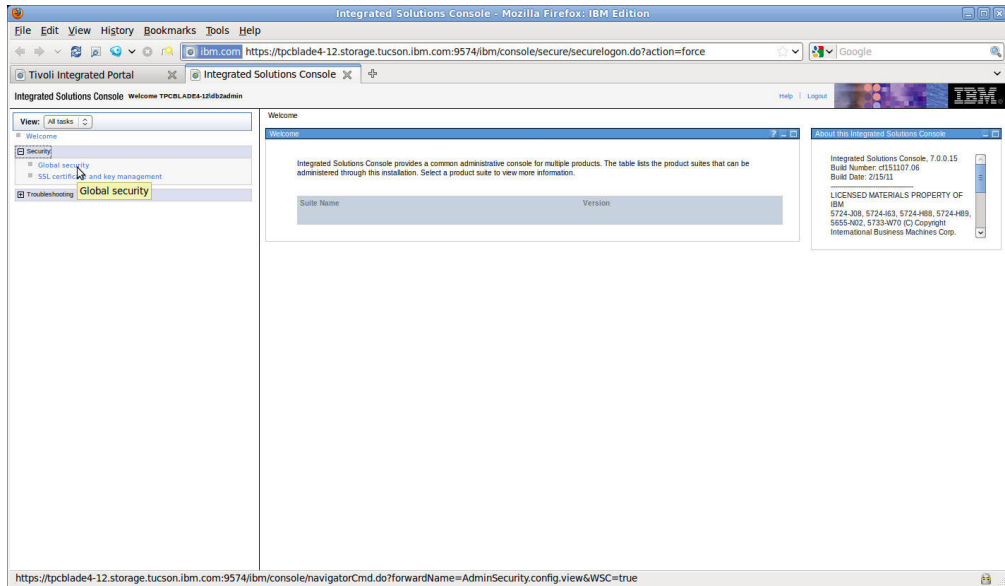


Figure 8. Opening the Global security page

- On the Global security page, in the "User account repository" section, click **Configure** next to the "Available realm definitions" drop-down list.

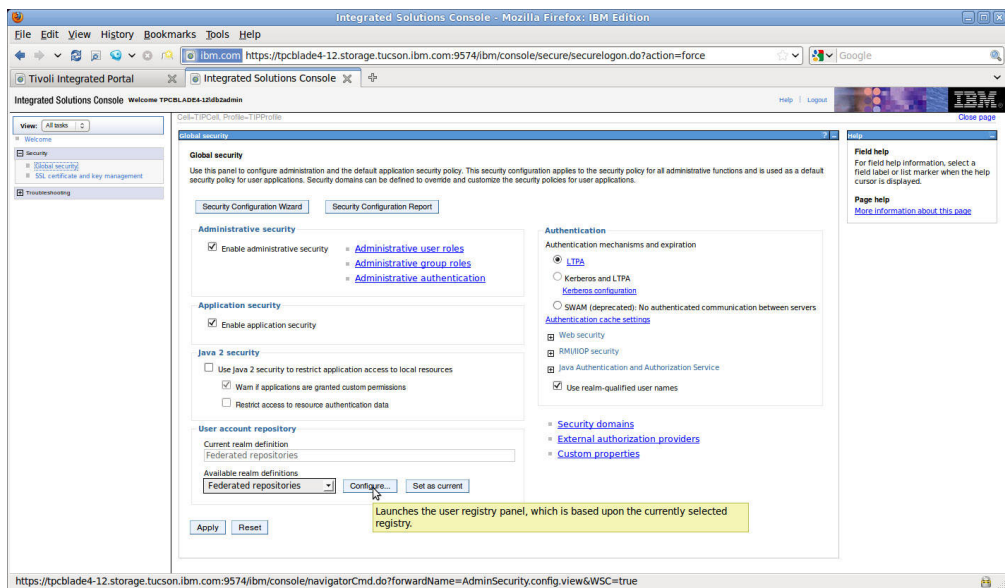


Figure 9. Configuring the available realm definitions

- On the Global security > Federated repositories page, in the "Repositories in the realm" table, select the entry for the LDAP repository you previously added. Click **Remove**.

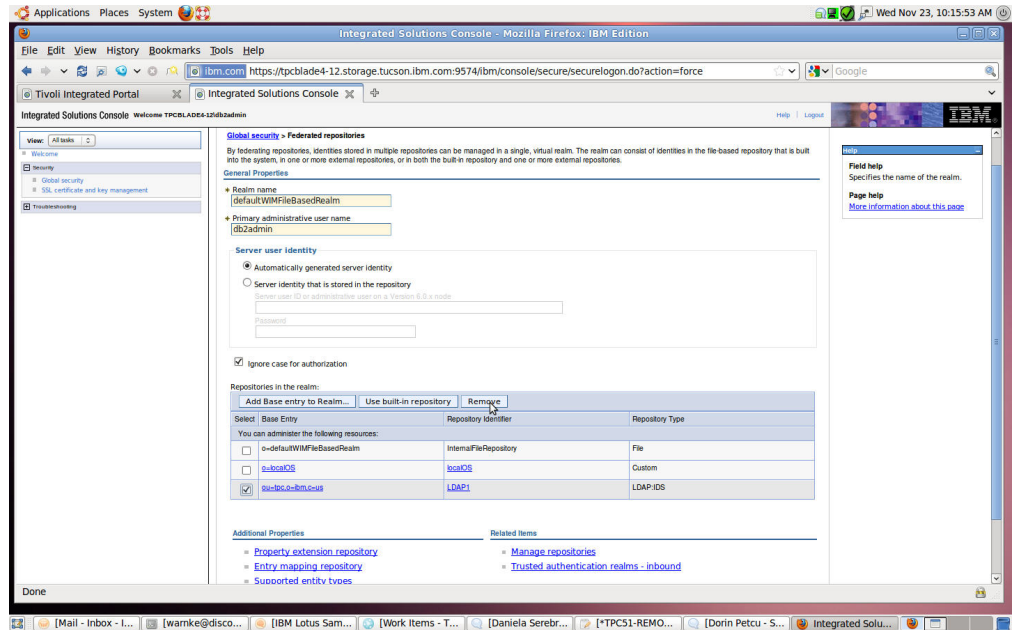


Figure 10. Removing the LDAP repository

10. In the message dialog box displayed on the Federated repositories page, click the **Save** link in **Save directly to the master configuration**.
11. On the Global security > Federated repositories page, under "Related items," click **Manage repositories**.

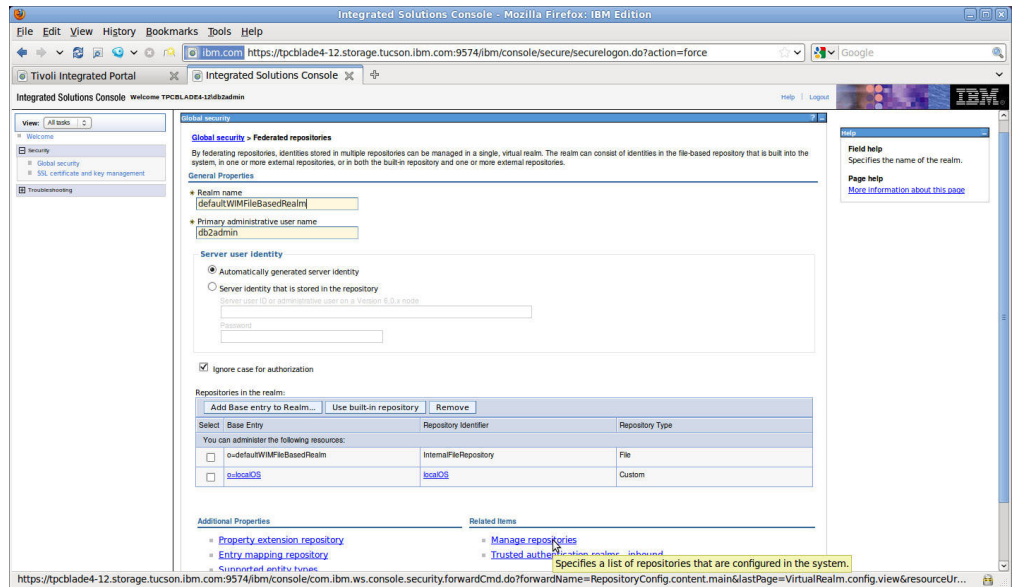


Figure 11. Managing the repositories

12. On the Global security > Federated repositories > Manage repositories page, select the entry for the LDAP repository that you previously added and click **Delete**.

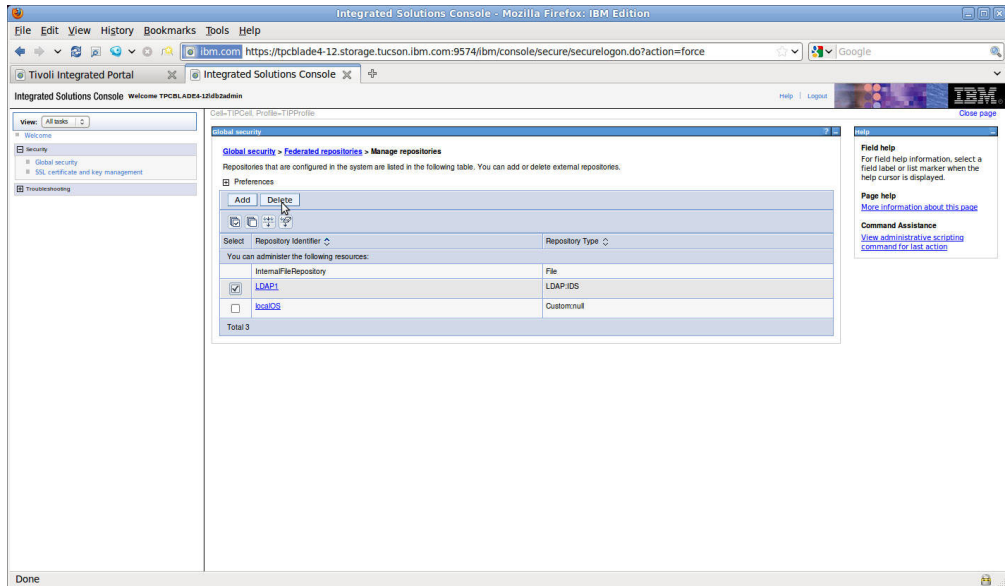


Figure 12. Deleting the LDAP repository

13. In the message dialog box displayed on the Manage repositories page, click the **Save** link in **Save directly to the master configuration**.
14. Log out from the WebSphere Administrative Console.
15. Log out from the Tivoli Integrated Portal Administrative Console.
16. Stop and restart the Tivoli Integrated Portal server. For information about stopping and starting the server, go to the Tivoli Storage Productivity Center information center. Search for *starting and stopping the Tivoli Integrated Portal server*.

Important: Before establishing the desired authorization configuration for Tivoli Integrated Portal, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication, you must ensure that there are no duplicated user names or group names within the local file-based repository and the localOS repository.

After you have completed the removal of the LDAP repository from the Tivoli Storage Productivity Center federated repositories authentication configuration, you should establish the desired authorization configuration for Tivoli Integrated Portal. Tivoli Integrated Portal administrators require the `iscadmins` and `administrator` roles. Operators viewing Tivoli Common Reporting reports require the `tcrPortalOperator` role. For information about user roles and groups, go to the Tivoli Integrated Portal administrative console online help. On the Welcome page, open the online help by clicking `?`. On the help welcome page, click **Using the Tivoli Integrated Portal > Users and Groups** in the navigation pane.

You should also assign Tivoli Storage Productivity Center roles to groups using the **Role-to-Group Mapping** node in the Tivoli Storage Productivity Center stand-alone GUI. You can map Tivoli Storage Productivity Center roles, such as `Tape Operator` or `Fabric Administrator`, to user groups that you create in your local operating system repository. For information about the **Role-to-Group Mapping** node, go to the Tivoli Storage Productivity Center information center. Search for *Role-to-Group Mappings*.

For information about assigning roles to users and roles to groups for Tivoli Storage Productivity Center for Replication, go to the Tivoli Storage

Productivity Center information center. Search for *Managing replication*. On the Managing replication page, click **Security**.

Tivoli Integrated Portal and Tivoli Common Reporting Roles

This topic provides a list of the predefined roles in IBM Tivoli Integrated Portal and IBM Tivoli Common Reporting.

Note: The Tivoli Storage Productivity Center installation program adds the administrator, iscadmins, tcrPortalOperator, chartAdministrator, and monitor roles to the Tivoli Storage Productivity Center installation user.

The Tivoli Storage Productivity Center administrator needs to assign the appropriate Tivoli Integrated Portal roles to any other Tivoli Storage Productivity Center users who are added later to the system.

Table 36. Roles in Tivoli Integrated Portal and Tivoli Common Reporting

Role	Description
iscusers	This role is set to All Authenticated users, and all users have this role by default. Users who are assigned this role can access the Welcome page and Credential Store portlets.
operator	Legacy role with no special privileges.
monitor	This Tivoli Integrated Portal role is a legacy role with no special privileges.
configurator	This Tivoli Integrated Portal role is a legacy role with no special privileges.
administrator	This Tivoli Integrated Portal role is a legacy role with no special privileges.
iscadmins	Users assigned to this role are the super users and have administrative access to all pages and portlets defined in Tivoli Integrated Portal.
chartAdministrator	Users assigned to this Tivoli Integrated Portal role create and delete charting connections to data sources, download the BIRT Designer, upload charts, and clear the charting cache (useful for troubleshooting).
chartCreator	Users assigned to this role can download the BIRT Designer, upload charts, view, and edit them. They cannot create or delete chart connections nor can they clear the charting cache.
chartViewer	Users assigned to this role can select and view charts, but cannot modify them or their preferences. They cannot download the BIRT Designer, upload charts, create connections, or clear the charting cache.
tcrPortalOperator	Users assigned to this Tivoli Common Reporting role are authorized to the Tivoli Common Reporting navigation items in the Tivoli Integrated Portal console.

Assigning Tivoli Integrated Portal roles to users

Assigning LDAP users and operating system users a Tivoli Integrated Portal role enables the users to take advantage of all the features in the Tivoli Integrated Portal Version 5.1 web-based GUI and IBM Cognos® Business Intelligence reports. Users can also use all of the launch-in-context and single sign-on features in Tivoli Storage Productivity Center 5.1.

To assign a Tivoli Integrated Portal role, complete the following steps:

1. Log on to Tivoli Integrated Portal. For example, you can enter `http://<machine where Tivoli Integrated Portal is installed>:9568/ibm/console`.
2. Click **Users and Groups > User Roles**.
3. Search for a user using an ID or a first and last name.
4. Select a user.
5. Under **Available Roles**, select a role and click **Save**.

The selected role is assigned to the user.

Configuration overview

After you install Tivoli Storage Productivity Center, you must configure your system to get the information you want. Some of the tasks associated with configuring the system include knowing which agents are required, the jobs that must be run, and the alerts you can set, and changing passwords.

To collect data from the devices and servers that are managed and monitored, different configuration steps are required to enable these different functions in Tivoli Storage Productivity Center.

Tivoli Storage Productivity Center uses four different types of agents to gather data about the devices and servers that are managed and monitored. Different combinations of these agents are required to effectively enable the functions of Data Manager, Fabric Manager, Disk Manager, and Tape Manager. In addition to these manager functions, the topology viewer is affected by the discovery of all the managed entities in the management scope of Tivoli Storage Productivity Center.

Agents

The agents are:

Storage Management Initiative Specification (SMI-S) providers

SMI-S providers or Common Information Model (CIM) agents are provided by the vendor of the storage device, fabric switch, or tape library. For storage, they are needed for storage asset information, provisioning, alerting, and performance monitoring.

CIM agents are not required for the DS8000, SAN Volume Controller, Storwize V7000 Unified, Storwize V7000, the XIV system, or IBM SONAS. Tivoli Storage Productivity Center uses native interfaces for these storage systems.

For the fabric switch, CIM agents are used for performance monitoring. The CIM agents are also used for certain switch models to collect topology information, port state information, zoning information, zone control, and alerting. For tape libraries, they are used for asset and inventory information. Each vendor of the storage, fabric switch, or tape library

supplies unique CIM agent code for their family of devices. These agents implement an SMI-S provider to provide a communication transport between Tivoli Storage Productivity Center and the managed devices. After a CIM agent is installed and configured, Tivoli Storage Productivity Center can be configured to communicate with it.

When setting up your CIM agents for managing devices, read the CIM agent documentation first. This documentation can provide guidelines for how many devices the CIM agent can be configured to manage. If the CIM agent documentation does not include such guidelines, limit three subsystems per CIM agent.

Storage Resource agents

The Storage Resource agents perform the functions of the Data agents and Fabric agents. Tivoli Storage Productivity Center does not use Data agent or the Fabric agent.

For computers that have storage that must be monitored, the Storage Resource agents are used to monitor that storage. These agents collect information from the server they are installed on. Asset information, file and file system attributes, and any other information needed from the computer system is gathered. The Storage Resource agents can also gather information about database managers installed on the server, and NAS device information. In Tivoli Storage Productivity Center, you can create ping, probe, and scan jobs to run against the servers that have Storage Resource agents installed.

For fabric information, the Storage Resource agents are installed on computer systems that have fiber connectivity (through HBAs) into the SAN fabrics you want to manage and monitor. The Storage Resource agents use scanners to collect information. The scanners communicate through the HBA to collect fabric topology information, port state information, and zoning information. They also can identify other SAN-attached devices (if they are in the same zone). Storage Resource agents are discovered during the agent installation process and are not required to be discovered separately. A best practice is to have two agents connected to each switch for redundancy.

Out-of-band agents

The out-of-band agents are used to collect topology information from fabric switches through the IP network by using SNMP queries to the switches. A best practice is to have Tivoli Storage Productivity Center communicating with the out-of-band agent (SNMP agent) on each switch in each SAN fabric you are managing.

Jobs

Tivoli Storage Productivity Center uses data collection jobs to provide information about your environment. The types of data collection jobs are:

Discovery jobs

Discovery jobs locate data sources and collect basic information about these data sources.

When the discovery job collects basic information for the DS8000, SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, the XIV system, or IBM SONAS, Tivoli Storage Productivity Center communicates directly with the storage system.

For other storage systems, the discovery job communicates with a storage subsystem CIM agent or tape library CIM agent. The discovery job locates all the storage systems or tape libraries for the CIM agent and retrieves basic information that the CIM agent holds for these systems. When a discovery job communicates with a fabric switch CIM agent, the job retrieves basic available information about the fabric and switches.

Probe jobs

Probe jobs collect detailed statistics on all the assets of the managed infrastructure, such as computers, disk controllers, hard disks, clusters, fabrics, storage subsystems, LUNs, tape libraries, and file systems. Probe jobs can also discover information about new or removed disks and file systems. Probe jobs can be directed against any element in the managed infrastructure.

There are four types of probe jobs: for computers, fabrics, storage subsystems, and tape libraries.

Computer probe jobs collect local host information such as file systems, directories, clusters, and hard disks. Storage subsystem probe jobs collect information about the storage subsystems, such as disk, volumes, and LUNs. Fabric probe jobs collect information about the fabric. Tape library probe jobs collect information about the tape libraries.

Computer probe jobs for a Storage Resource agent also perform a fabric discovery. The information collected includes fabric WWNs of the fabrics connected to the host on which the Storage Resource agent is running, fabric to switch relationships, and some key switch attributes. Computer probe jobs for a Storage Resource agent also collect HBA information.

When probing storage subsystems that are registered with the same CIM agent, do not probe more than three storage subsystems within the same probe job because the increased load on the agent would increase the likelihood of timeouts. Instead, spread the storage subsystems across multiple probe jobs with different starting times.

Scan jobs

Scan jobs collect statistics about the actual storage consumption. Scan jobs are always directed against a Storage Resource agent and deliver detailed information about the file systems, files, and databases of a computer.

Ping jobs

Ping jobs gather statistics about the availability of the managed computers. Ping jobs generate TCP/IP pings and consider the computer available if it gets an answer. Like scan jobs, ping jobs can be directed against computers only.

Performance monitoring jobs

Performance monitoring jobs collect statistics about the performance of storage subsystems and switches. Performance monitoring jobs can be run against storage subsystems, SAN Volume Controllers, Storwize V7000 Unified, Storwize V7000 systems, the XIV system, IBM SONAS systems, and switches.

Restriction: You cannot use storage subsystem groups with the performance monitor.

Alerts

You can define alerts so that Tivoli Storage Productivity Center notifies you when a specified event occurs. Such events are the triggering conditions for the alert. The specific triggering conditions that you can use to define an alert depends on the storage resource that you are monitoring.

The alerts you can define are:

Computer

Use this alert to be notified when a condition is detected on a computer during a probe job.

This alert is a Data Manager alert.

Examples of computer alerts are when RAM is increased or decreased, virtual storage is increased or decreased, or when a new disk is detected or not found.

Filesystem

Use this alert to be notified when a condition is detected on a file system during a probe job or upon an SNMP event.

Examples of file system alerts are when the file system is not found or reconfigured or is low on free space.

This alert is a Data Manager alert.

Directory

Use this alert to be notified when a condition is detected on a directory during a scan job.

Examples of directory alerts are when the directory is not found or when a user or directory storage quota is exceeded.

This alert is a Data Manager alert.

Hypervisor

Use this alert to be notified when a hypervisor is discovered or missing, or a virtual machine is added or deleted from a VMware system.

This alert is a Data Manager alert.

Other NAS

Use this alert to be notified when a condition is detected on a network-attached storage (NAS) device during a probe job or upon an SNMP event. This alert type is for NAS devices that are not included in other alert types.

This alert is a Data Manager alert.

Instance

Use this alert to be notified when a condition is detected for a database.

Examples of instance alerts are when a new database or table space is discovered for a specific RDBMS.

This alert is a Data Manager for Databases alert.

Database-tablespace

Use this alert to be notified when a condition is detected for a database-tablespace.

Examples of database-tablespace alerts are when Tivoli Storage Productivity Center detects that free space is low or the database or table space is offline or dropped for a specific RDBMS.

This alert is a Data Manager for Databases alert.

Table Use this alert to be notified when a condition is detected for a table in a database.

Examples of table alerts are when a managed table is dropped or exceeds its storage quota.

This alert is a Data Manager for Databases alert.

Storage subsystem

Use this alert to be notified when a condition is detected on a storage subsystem during a probe job.

This alert is a Disk Manager alert.

Storage subsystem alerts include the following examples:

- When a storage subsystem is discovered or not found
- A subsystem cache is increased or decreased
- A disk is detected or not found
- A subsystem goes offline or online
- A port is discovered or missing
- A back-end controller goes offline or online
- A volume goes offline or online
- A pool is missing or discovered
- A node goes offline or online
- A SAN Volume Controller Storwize V7000 Unified, Storwize V7000, or an IBM SONAS node is missing or discovered
- A disk utilization percentage threshold is reached

Restriction: You cannot use storage subsystem groups with threshold alerts.

Fabric

Use this alert to be notified when a condition is detected on a fabric during a probe job.

Examples of fabric alerts are if a fabric is missing or discovered, a fabric goes online or offline, a zone is missing or discovered, or a zone set is missing or discovered.

This alert is a Fabric Manager alert.

Switch

Use this alert to be notified when a condition is detected on a switch during a probe job.

Examples of switch alerts are if a switch is missing or discovered, a switch goes online or offline, a switch blade is missing or discovered, a switch blade goes online or offline, a total port data rate threshold is reached, or a total port packet rate threshold is reached.

This alert is a Fabric Manager alert.

Endpoint device

Use this alert to be notified when a condition is detected on an endpoint

device during a probe job. Endpoint devices are devices discovered by the Fabric Manager but cannot be identified as a specific device type, such as a computer or subsystem.

Examples of endpoint device alerts are when an endpoint device is missing or discovered, a peripheral entity to node association is missing or discovered, or the version for an endpoint device changes.

This alert is a Fabric Manager alert.

After you define alerts, you can choose to be notified through different methods.

SNMP trap

Generate an SNMP trap message to any console or terminal to indicate the occurrence of an alert.

TEC/OMNIbus event

Send alerts to the IBM Tivoli Enterprise Console or OMNIbus console. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Data Manager events and perform responses such as send further notification, create or update trouble tickets, or run programs.

Login notification

Send alerts to a specified user upon logging on to the system.

Windows Event Log, UNIX Syslog

Record alerts to the OS Event Log. If you already have an administrator monitoring OS logs, this action is an easy way to have all of your priority messages centralized for quick notification and viewing.

Run script

Run a script in response to an alert. (This option is not valid for the Storage Resource agent.)

Email Send alerts through to email addresses that you specify.

Change password tool

The password tool changes the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal so that they can continue to authenticate to one another whenever you change a password. For more information about the password tool, see “Changing passwords using the password tool.”

Changing passwords using the password tool

The password tool changes the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal so that they can continue to authenticate to one another whenever you change a password.

Single server installation where components use the same logon credentials

Use the password tool to change the password for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on the same server, and use the same user name and password.

Before you use the password tool, ensure that you know the existing password or passwords that you want to change.

1. Complete the following steps:

For Windows operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

- a. From the desktop, click **Start > Control Panel**.
- b. Double-click **User Accounts**.
- c. Click **Change your password**.
- d. Enter the current password, and then enter and confirm the new password. Click **Change password**. Close the User Accounts window.

For AIX or Linux operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

Log in with the DB2 user ID, and run the **passwd** command.

Tip: To change other user credentials, run the **passwd *username*** command.

Tip: To find out how to change the LDAP user password, refer to your LDAP server documentation.

2. Open a command prompt window and change the directory to the following default directory:

On Windows operating systems:

C:\Program Files\IBM\TPC\service

On Linux or UNIX operating systems:

/opt/IBM/TPC/service

3. Start the password tool by running the following program:

On Windows operating systems:

changepasswords.bat

On Linux or AIX operating systems:

changepasswords.sh

4. From the "Choose the password to change" window, click **Change Tivoli Integrated Portal password** and click **OK**.
5. Complete the following steps:
 - a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1.
 - b. Ensure that the **Restart servers** option is selected. Click **OK**.
 - c. In response to the following message, click **Yes**:
Are you sure you want to change the password?
 - d. When the process is completed, click **Back to Main**.
6. To change the password for Tivoli Storage Productivity Center, click **Change Tivoli Storage Productivity Center server and DB2 passwords** and click **OK**.
 - a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1. Click **OK**.
 - b. In response to the following message, click **Yes**:
Are you sure you want to change the password?

Tip: After you click **Yes**, if the DB2 servers are already started, the following errors might be displayed in a command prompt window. Ignore them and continue with the procedure.

```
Changing DB2 User password for Data Server...
Changing DB2 User password for Device Server...
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
```

- c. When the process is completed, click **Back to Main**.

Tip: To verify that the password changes were completed successfully, review the `PWTool.log` file that is in the `TPC_installation_directory\service directory`, where `TPC_installation_directory` is the directory where Tivoli Storage Productivity Center is installed.

7. Exit the password tool.

Related tasks:

“Single server installation where components use different logon credentials”
Use the password tool to change the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on the same server, but use different user names and passwords.

“Multiple server installation where Tivoli Integrated Portal is remote” on page 229
Change the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on two servers, with Tivoli Integrated Portal on server *A*, and DB2 and Tivoli Storage Productivity Center on server *B*.

“Multiple server installation where DB2 is remote” on page 231
Change the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on two servers, with Tivoli Integrated Portal and Tivoli Storage Productivity Center on server *A*, and DB2 on server *B*.

Single server installation where components use different logon credentials

Use the password tool to change the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on the same server, but use different user names and passwords.

Before you use the password tool, ensure that you know the existing password or passwords that you want to change.

1. Complete the following steps:

On Windows operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

- a. From the desktop, click **Start > Control Panel**.
- b. Double-click **User Accounts**.
- c. Click **Change your password**.
- d. Enter the current password, and then enter and confirm the new password. Click **Change password**. Close the User Accounts window.

On AIX or Linux operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

When you are logged in with the DB2 user ID, run the `passwd` command. To change other user credentials, run the `passwd username` command.

Tip: To find out how to change the LDAP user password, refer to your LDAP server documentation.

2. Open a command prompt window and change the directory to the following default directory:

On Windows operating systems:

C:\Program Files\IBM\TPC\service

On Linux or UNIX operating systems:

/opt/IBM/TPC/service

3. Start the password tool by running the following program:

On Windows operating systems:

changepasswords.bat

On Linux or AIX operating systems:

changepasswords.sh

4. From the "Choose the password to change" window, click **Change Tivoli Integrated Portal password** and click **OK**.
5. To change the password, complete the following steps:
 - a. Enter and confirm the new password and click **OK**.
 - b. In response to the following message, click **Yes**:
Are you sure you want to change the password?
 - c. When the process is completed, click **Back to Main**.
6. To change the password for Tivoli Storage Productivity Center, click **Change Tivoli Storage Productivity Center server passwords** and click **OK**.
7. Complete the following steps:
 - a. Enter and confirm the new password and click **OK**.
 - b. In response to the following message, click **Yes**:
Are you sure you want to change the password?
 - c. When the process is completed, click **Back to Main**.
8. To change the password for DB2, click **Change DB2 password** and click **OK**.
9. Complete the following steps:
 - a. Enter and confirm the new password. Ensure that the **Restart servers** check box is selected and click **OK**.
 - b. In response to the following message, click **Yes**:
Are you sure you want to change the password?

Tip: After you click **Yes**, if the DB2 servers are already started, the following errors might be displayed in a command prompt window. Ignore them and continue with the procedure.

```
Changing DB2 User password for Data Server...
Changing DB2 User password for Device Server...
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
```

- c. When the process is completed, click **Back to Main**.

Tip: To verify that the password changes were completed successfully, review the `PWTool.log` file that is in the `TPC_installation_directory\service` directory, where `TPC_installation_directory` is the directory where Tivoli Storage Productivity Center is installed.

Important: For Windows operating systems, if Tivoli Integrated Portal fails to stop and restart, restart your system manually.

Related tasks:

“Single server installation where components use the same logon credentials” on page 225

Use the password tool to change the password for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on the same server, and use the same user name and password.

“Multiple server installation where Tivoli Integrated Portal is remote”

Change the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on two servers, with Tivoli Integrated Portal on server *A*, and DB2 and Tivoli Storage Productivity Center on server *B*.

“Multiple server installation where DB2 is remote” on page 231

Change the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on two servers, with Tivoli Integrated Portal and Tivoli Storage Productivity Center on server *A*, and DB2 on server *B*.

Multiple server installation where Tivoli Integrated Portal is remote

Change the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on two servers, with Tivoli Integrated Portal on server *A*, and DB2 and Tivoli Storage Productivity Center on server *B*.

Before you use the password tool, ensure that you know the existing password or passwords that you want to change.

1. On the server where DB2 and Tivoli Storage Productivity Center are installed, complete the following steps to change the DB2 and Tivoli Storage Productivity Center passwords:

On Windows operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

- a. From the desktop, click **Start > Control Panel**.
- b. Double-click **User Accounts**.
- c. Click **Change your password**.
- d. Enter the current password, and then enter and confirm the new password. Click **Change password**. Close the User Accounts window.

For AIX or Linux operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

When you are logged in with the DB2 user ID, run the **passwd** command. To change other user credentials, run the **passwd *username*** command.

Tip: To find out how to change the LDAP user password, refer to your LDAP server documentation.

2. On the server where Tivoli Integrated Portal is installed, complete the following tasks:

On Windows operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

- a. From the desktop, click **Start > Control Panel**.
- b. Double-click **User Accounts**.
- c. Click **Change your password**.

- d. Enter the current password, and then enter and confirm the new password. The new password can be different from the password that you created in step 1. Click **Change password**. Close the User Accounts window.

For AIX or Linux operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

When you are logged in with the DB2 user ID, run the **passwd** command. To change other user credentials, run the **passwd *username*** command.

3. On the same server, open a command prompt window and change the directory to the following default directory:

On Windows operating systems:

C:\Program Files\IBM\TPC\service

On Linux or UNIX operating systems:

/opt/IBM/TPC/service

4. Start the password tool by running the following program:

On Windows operating systems:

changepasswords.bat

On Linux or AIX operating systems:

changepasswords.sh

5. From the “Choose the password to change” window, click **Change Tivoli Integrated Portal password** and click **OK**.
6. Complete the following steps:
 - a. Enter and confirm the new password and click **OK**.
 - b. In response to the following message, click **Yes**:
Are you sure you want to change the password?
 - c. When the process is completed, click **Back to Main**.
7. On the same server, from the “Choose the password to change” window, click **Change DB2 password** and click **OK**.
8. Complete the following steps:
 - a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1. Ensure that the **Restart servers** check box is selected and click **OK**.
 - b. In response to the following message, click **Yes**:
Are you sure you want to change the password?

Tip: After you click **Yes**, if the DB2 servers are already started, the following errors might be displayed in a command prompt window. Ignore them and continue with the procedure.

```
Changing DB2 User password for Data Server...
Changing DB2 User password for Device Server...
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
```

- c. When the process is completed, click **Back to Main**.
9. On the server where Tivoli Storage Productivity Center and DB2 are installed, from the “Choose the password to change” window, click **Change Tivoli Storage Productivity Center server passwords** and click **OK**.
10. Complete the following steps:

- a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1. Click **OK**.
 - b. In response to the following message, click **Yes**:
Are you sure you want to change the password?
 - c. When the process is completed, click **Back to Main**.
11. On the same server, click **Change DB2 password** and click **OK**.
 12. Complete the following steps:
 - a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1. Click **OK**.
 - b. In response to the following message, click **Yes**:
Are you sure you want to change the password?

Tip: After you click **Yes**, if the DB2 servers are already started, the following errors might be displayed in a command prompt window. Ignore them and continue with the procedure.

```
Changing DB2 User password for Data Server...
Changing DB2 User password for Device Server...
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
```

- c. When the process is completed, click **Back to Main**.
13. Exit the password tool.

Tip: To verify that the password changes were completed successfully, review the `PWTool.log` file that is in the `TPC_installation_directory\service` directory, where `TPC_installation_directory` is the directory where Tivoli Storage Productivity Center is installed.

Related tasks:

“Single server installation where components use the same logon credentials” on page 225

Use the password tool to change the password for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on the same server, and use the same user name and password.

“Single server installation where components use different logon credentials” on page 227

Use the password tool to change the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on the same server, but use different user names and passwords.

“Multiple server installation where DB2 is remote”

Change the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on two servers, with Tivoli Integrated Portal and Tivoli Storage Productivity Center on server *A*, and DB2 on server *B*.

Multiple server installation where DB2 is remote

Change the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on two servers, with Tivoli Integrated Portal and Tivoli Storage Productivity Center on server *A*, and DB2 on server *B*.

Before you use the password tool, ensure that you know the existing password or passwords that you want to change.

1. On the server where DB2 is installed, do the following steps:

For Windows operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

- a. From the desktop, click **Start > Control Panel**.
- b. Double-click **User Accounts**.
- c. Click **Change your password**.
- d. Enter the current password, and then enter and confirm the new password. Click **Change password**. Close the User Accounts window.

For AIX or Linux operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

When you are logged in with the DB2 user ID, run the **passwd** command. To change other user credentials, run the **passwd username** command.

Tip: To find out how to change the LDAP user password, refer to your LDAP server documentation.

2. On the server where Tivoli Integrated Portal and Tivoli Storage Productivity Center are installed, do the following steps:

For Windows operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

- a. From the desktop, click **Start > Control Panel**.
- b. Double-click **User Accounts**.
- c. Click **Change your password**.
- d. Enter the current password, and then enter and confirm the new password. The new password can be different from the password that you created in step 1. Click **Change password**. Close the User Accounts window.

For AIX or Linux operating systems where Tivoli Storage Productivity Center is configured for OS authentication:

When you are logged in with the DB2 user ID, run the **passwd** command. To change other user credentials, run the **passwd username** command.

3. On the same server, open a command prompt window and change the directory to the following default directory:

On Windows operating systems:

C:\Program Files\IBM\TPC\service

On Linux or UNIX operating systems:

/opt/IBM/TPC/service

4. Start the password tool by running the following program:

On Windows operating systems:

changepasswords.bat

On Linux or AIX operating systems:

changepasswords.sh

5. From the “Choose the password to change” window, click **Change Tivoli Integrated Portal password** and click **OK**.

6. Complete the following steps:

- a. Enter and confirm the new password and click **OK**.
- b. In response to the following message, click **Yes**:
Are you sure you want to change the password?

- c. When the process is completed, click **Back to Main**.
7. On the same server, from the “Choose the password to change” window, click **Change Tivoli Storage Productivity Center server passwords** and click **OK**.
8. Complete the following steps:
 - a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 2. Click **OK**.
 - b. In response to the following message, click **Yes**:
Are you sure you want to change the password?
 - c. When the process is completed, click **Back to Main**.
9. On the server where DB2 is installed, from the “Choose the password to change” window, click **Change DB2 password** and click **OK**.
10. Complete the following steps:
 - a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1. Ensure that the **Restart servers** check box is selected, and click **OK**.
 - b. In response to the following message, click **Yes**:
Are you sure you want to change the password?

Tip: After you click **Yes**, if the DB2 servers are already started, the following errors might be displayed in a command prompt window. Ignore them and continue with the procedure.

```
Changing DB2 User password for Data Server...
Changing DB2 User password for Device Server...
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
```
 - c. When the process is completed, click **Back to Main**.
11. Click **Exit program**.
12. On the server where Tivoli Integrated Portal and Tivoli Storage Productivity Center are installed, from the “Choose the password to change” window, click **Change DB2 password**.
13. Complete the following steps:
 - a. Enter and confirm the new password, ensuring that it is the same as the new password that you entered in step 1. Ensure that the **Restart servers** check box is selected, and click **OK**.
 - b. In response to the following message, the following message, click **Yes**:
Are you sure you want to change the password?

Tip: After you click **Yes**, if the DB2 servers are already started, the following errors might be displayed in a command prompt window. Ignore them and continue with the procedure.

```
Changing DB2 User password for Data Server...
Changing DB2 User password for Device Server...
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
OpenServiceW failed, error code = 5
```
 - c. When the process is completed, click **Back to Main**.
14. Exit the password tool.

Important: For the changes to take effect, you might have to manually restart the Data server, the Device server, or Tivoli Integrated Portal, or all three.

Tip: To verify that the password changes were completed successfully, review the `PWTool.log` file that is in the `TPC_installation_directory\service` directory, where `TPC_installation_directory` is the directory where Tivoli Storage Productivity Center is installed.

Related tasks:

“Single server installation where components use the same logon credentials” on page 225

Use the password tool to change the password for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on the same server, and use the same user name and password.

“Single server installation where components use different logon credentials” on page 227

Use the password tool to change the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on the same server, but use different user names and passwords.

“Multiple server installation where Tivoli Integrated Portal is remote” on page 229
Change the passwords for DB2, Tivoli Storage Productivity Center, and IBM Tivoli Integrated Portal when they are installed on two servers, with Tivoli Integrated Portal on server *A*, and DB2 and Tivoli Storage Productivity Center on server *B*.

Configuration Utility

Use the Configuration Utility to help you learn how to configure your Tivoli Storage Productivity Center system.

Configuration utility tasks

The configuration utility is a tool that guides you through the steps and prerequisite steps required to configure your system.

All of the functions in the configuration utility exist on their own within the IBM Tivoli Storage Productivity Center GUI. However, the configuration utility provides you with a single place for new users to find all these common functions and information to help you decide what agents you need and which jobs to run to get the information you want. It provides links into the Tivoli Storage Productivity Center GUI to display reports, add agents, create jobs, create alerts, and so on.

If you are a new user of Tivoli Storage Productivity Center, use the configuration utility to learn what the common functions are and to help you with the task of deciding what agents you need and which jobs to run to get the information you want.

If you are an experienced Tivoli Storage Productivity Center user, you can go to the GUI directly to perform all of the configuration tasks. However, you might also want to use the configuration utility to provide you with a quick summary review of your system.

Depending on what information you want to see and collect, you will go to a different tab in the configuration utility for information and configuration.

How to manage the servers and agents

This section describes how to use the configuration utility to manage your servers and agents.

To manage your servers, storage subsystems, and agents, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center > Configuration Utility > Services** tab.
2. You can see the following items to manage:
 - Data Server - indicates whether the server is up or down. Provides detailed information about the Data server.
 - Device Server - indicates whether the server is up or down. Provides detailed information about the Device server.
 - Data/Storage Resource Agents - lists the Data agents and Storage Resource agents. Provides detailed information about the agents and their status.
 - CIM Agents - lists the CIM Agents installed and connected to Tivoli Storage Productivity Center. You can add CIM agents, edit CIM agent details, test the CIM agent connections, run a CIMOM discovery, and show the managed devices for the CIM agent.
 - Inband Fabric Agents - lists the Inband Fabric agents installed. Provides detailed information about the agents and their status.
 - Out of Band Fabric Agents - lists the Out of Band Fabric agents. Provides detailed information about the agents and their status.
 - Storage Subsystem Interfaces - lists the storage systems. You can add, modify, or remove a storage system. You can also update or configure the credentials for the storage system, or test the connection to the storage system.

Show managed devices

This section describes how to use the configuration utility to show the devices (storage subsystems, switches, or tape libraries) that are managed by the CIM agent.

To show the devices, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center > Configuration Utility > Services** tab.
2. Under the CIM Agents section, click **Show Managed Devices**. A dialog box displays the managed devices.

How to display availability reports

This section describes how to use the configuration utility to display availability reports.

To display availability reports of storage assets, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center > Configuration Utility > Data Manager** tab.
2. In the Ping section, you can run, create, or display information about the ping job. You can then display or define the storage availability system report.

How to monitor your storage assets, storage systems, and fabrics

This section describes how to use the configuration utility to monitor your storage assets, storage systems, and fabrics in your enterprise.

To monitor your storage assets, storage systems, and fabrics, complete the following steps:

1. For information collected from the Data Manager, expand **Tivoli Storage Productivity Center > Configuration Utility > Data Manager** tab. Use the Ping

section to collect information about the availability of the hosts managed by Storage Resource agents or Data agents in your enterprise. The ping jobs provide information for the availability reports.

Use the Computer Probes section to collect statistics about the storage assets in your enterprise. The computer probe jobs provide information for different reports.

Use the Scan section to collect statistics about the files on the hosts where Storage Resource agents or Data agents are installed. The scan jobs provide information for capacity, usage, usage violations, and backup reports.

2. For information collected from the Disk Manager, expand **IBM Tivoli Storage Productivity Center > Configuration Utility > Disk Manager** tab. Use the Storage Subsystems Probes section to collect information about the configuration of the storage system. Some of the information collected are storage pools, storage volumes, FC ports, disks, and host-to-volume assignments. A probe job must be run before volumes can be created or monitored for performance on a storage system.

Use the Storage Subsystem Performance Management section to create and edit a performance monitor job. The performance monitor job collects information about the performance statistics for the storage system.

3. For information collected from the Fabric Manager, expand **IBM Tivoli Storage Productivity Center > Configuration Utility > Fabric Manager** tab.

You can have Storage Resource agents installed on systems in each fabric that you are managing, CIM agents that are configured to manage each fabric, and out-of-band Fabric agents that communicate with each switch in the fabric. It is a good practice to have multiple agent types configured. CIM agents are the preferred agent type for Brocade fabrics. For Cisco switches, out-of-band agents are required for VSAN information. Zoning information is provided only by the Storage Resource agents for Cisco and QLogic switches.

Before you can run a switch performance monitoring job in IBM Tivoli Storage Productivity Center, you must have a switch CIM agent installed and configured in Tivoli Storage Productivity Center. This CIM agent must be SMI-S V1.1 compliant.

For more information about the supported agent types for switch performance management and fabric zone configuration, see “Collecting data with the Fabric Manager” on page 91.

Use the Fabric Probes section to collect information about the configuration of the fabric. A probe job must be run before information about performance for a fabric can be collected.

The discovery of Storage Resource agents is performed automatically when the agent associated with your Tivoli Storage Productivity Center is installed or started. Storage Resource agents perform fabric discovery during a computer probe and in response to events. The information gathered during a fabric discovery performed by the Storage Resource agent includes fabric and fabric to switch relationships.

A Storage Resource agent discovery gathers the WWN of the fabric and WWN of the switches. For topology information, the agent discovers the topology for the entire fabric. A fabric probe with the Storage Resource agent collects other information such as zoning, topology, and so on. For more information about what information is gathered, see “Planning for the Fabric Manager” on page 85.

How to create alert definitions

This section describes how to use the configuration utility to create alert definitions.

You can create alert definitions to notify you of an alert condition. When you receive an alert notification, you can take the appropriate action when a certain condition occurs on storage subsystems, computers, file systems, directories, fabrics, switches, or endpoint devices.

To create an alert definition, complete the following steps:

1. If you want to see the alerts for Data Manager, expand **IBM Tivoli Storage Productivity Center > Configuration Utility > Data Manager** tab to create alerts for a computer resource.

In the Alerting section, you can create alerts for computers, file systems, and directories. You can also see the alert logs for computers, file systems, and directories.

2. If you want to see the alerts for Disk Manager, expand **IBM Tivoli Storage Productivity Center > Configuration Utility > Disk Manager** tab

In the Alerting section, you can create alert definitions for the storage system and view the alert logs.

There are two categories of alerts: device alerts and performance alerts. Device alerts use triggering conditions such as actions on the device itself. These alerts include discovery or state changes. Performance alerts use triggering conditions that are thresholds on measured metrics.

Performance thresholds are only available with alert definitions when the IBM Tivoli Storage Productivity Center package contains the performance capability.

Performance alerts must be defined before the performance monitor job is run. Because these alerts are based on thresholds violated from measured metrics, they can only be triggered when defined in an alert and when those metrics are being collected by a monitoring job.

3. If you want to see the alerts for Fabric Manager, expand **IBM Tivoli Storage Productivity Center > Configuration Utility > Fabric Manager** tab. In the Alerting section, you can create fabric, switch, and endpoint point device alerts. You can also view the logs for these alerts.

There are two categories of alerts: the device alert and performance alert. Device alerts use triggering conditions which are actions on the device itself. Examples of these alerts are discovery, state change, and so on. Performance alerts use triggering conditions that are thresholds on measured metrics. Device alerts are composed of three categories: fabric alerts, switch alerts, and endpoint alerts. A best-practice procedure would be to define more than the three default alerts that are present. In addition, the three default alerts should be modified to include a "Triggered Action".

Performance thresholds are only available for alert definitions when the IBM Tivoli Storage Productivity Center package contains the performance capability. For Fabric Manager, the only performance alerts are switch.

Performance alerts must be defined before a performance monitor job is run. Because these alerts are based on thresholds violated from measured metrics, they can only be triggered when they are defined in an alert and when those metrics are being collected in a monitoring job.

How to manage tape library reports

This section describes how to use the configuration utility to manage tape library reports.

A CIM agent must be installed and configured to collect information from the tape libraries. After the CIMOM discovery job completes, you can create a tape library probe job that collects asset information.

To manage your tape library reports, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center > Configuration Utility > Tape Manager** tab.
2. Use the CIMOM Discovery section to run a CIMOM Discovery job.
3. Use the Tape Library Probes section to create and run a probe job for the tape library.
4. After you have run the probe job, you can see information in the tape library reports. The probe job collects information about the configuration of the tape library. Some of the information collected are drives, changers, cartridges, slots, and I/O ports.

Managing element managers and external tools

In the **Configuration Utility** node of the stand-alone GUI, you can add and manage element managers and external tools.

Element managers helps you configure and maintain your storage devices. An example of an element manager is DS8000 Storage Manager, which is the default element manager that is provided with IBM System Storage DS8000. External tools are tools that are web-based user interfaces and local executable files. For example, you might use an external tool like a text editor or spreadsheet application that you often use. You can add and manage connections to element managers and external tools from the Configuration Utility.

To add and manage connections to element managers and external tools:

1. Log in to the IBM Tivoli Storage Productivity Center stand-alone GUI.
2. Click **IBM Tivoli Storage Productivity Center > Configuration Utility > Element Manager**.
3. On the **Element Manager** tab, configure the following panels:

DS8000 Element Manager

Displays the element managers that were added for DS8000.

Other Device Element Manager

displays the element managers that were added for devices other than DS8000. For example, element managers for other DS series systems, IBM TotalStorage Enterprise Storage Server, SAN Volume Controller, Storwize V7000, and IBM XIV Storage System.

External Tools

Displays external tools such as web-based interfaces and local executable files that were added.

Advanced Settings

Defines default settings for element managers and external tools. For example, you can define the default web browser to start web-based applications.

- a. In the list, select **Configure Launch** and enter information in the following fields:
 - **AIX/UNIX/Linux application launch command**, enter the command that is used to start applications on an AIX, UNIX, or Linux operating system.

- **AIX/UNIX/Linux URL launch command**, enter the command that is used to start the web browser for web-based applications on AIX, UNIX, or Linux operating systems.
For example, to start the Mozilla Firefox browser, if the executable file is in the `/usr/bin` directory, enter `/usr/bin/firefox`.
- **Windows application launch command**, enter the command that is used to start applications on a Windows operating system.
- **Windows URL launch command**, enter the command that is used to start the web browser for web-based applications on a Windows operating system. For example, to start Mozilla Firefox, and the executable file can be in any directory, you enter:
`cmd /c start firefox`
- **URL parsing regular expression**, enter the regular expressions used for Uniform Resource Locator (URL) commands. Regular expressions are sets of symbols and characters that are used to match patterns of text. You can use regular expressions to search specific URLs across your web environment. The default regular expression in this field matches the URLs that include the protocol `http`, `https`, or `ftp`.

Accessing Tivoli Storage Productivity Center for Replication through Tivoli Storage Productivity Center

To access Tivoli Storage Productivity Center for Replication through the Tivoli Storage Productivity Center graphical user interface (GUI), you must configure the Replication Manager and specify a web browser in the Element Manager.

To access Tivoli Storage Productivity Center for Replication through Tivoli Storage Productivity Center:

1. Click **IBM Tivoli Storage Productivity Center > Configuration Utility > Element Manager**.
 - a. Expand **Advanced Settings**.
 - b. Select **Configure Launch**.
 - c. Enter the URL launch command for your operating system.
 - In the **AIX/Linux/UNIX URL launch command** field, specify a web browser for your operating system. For example, to start Mozilla Firefox, if the executable file is in `/usr/bin`, change `{DEFAULT_BROWSER}` to `/usr/bin/firefox`.
 - In the **Windows URL launch command** field, enter the command that is used to start the web browser for web-based applications on a Windows operating system. For example, to start Mozilla Firefox, where the executable file can be in any directory, enter `cmd /c start firefox`

Restriction: You must define a web browser for operating systems to start Tivoli Storage Productivity Center for Replication from Tivoli Storage Productivity Center. If a browser is not specified in this field, or an incorrect browser is defined, Tivoli Storage Productivity Center for Replication cannot start.

2. Click **IBM Tivoli Storage Productivity Center > Configuration Utility** and click the following buttons to configure **Replication Manager**:

Replication Health Overview

Provides the following statuses and connections for Tivoli Storage Productivity Center for Replication:

- Sessions

- Connections

If you are logged on to the local server, the status of the standby server is displayed, and if you are logged on to the standby server, the status of the local server is displayed.

Replication Sessions Overview

Lists the sessions that are defined in the Tivoli Storage Productivity Center for Replication environment, including state and status. You must click **Create Session** to define a session.

Replication Storage Systems Overview

Lists all known storage systems and, if enabled, indicates whether the storage systems are communicating with the active and remote servers.

Replication Paths Overview

Summarizes all known ESS/DS series paths and lists them by storage system.

Replication Management Servers Overview

Displays the status of the management servers configuration and lists up to operating two management servers. Using the list, you can define a standby server or define the local server as a standby server to an alternate server. If you defined a standby server, this page has more than one variation.

Replication Administration

If you belong the Administrator user role in Tivoli Storage Productivity Center for Replication, displays a list of Tivoli Storage Productivity Center for Replication users and groups and their access privileges. Use this option to manage users and groups.

Replication Advanced Tools

Use this option to create a diagnostic package, change the automatic refresh rate of the GUI, and in IBM TotalStorage Enterprise Storage Server and DS, enable or disable the heartbeat.

Configure Devices wizard

Use the Configure Devices wizard to set up storage devices for monitoring by IBM Tivoli Storage Productivity Center. The wizard guides you through the steps for adding a device as a data source, running a discovery, including devices in groups, specifying alerts, and setting up data collection schedules.

In previous versions of IBM Tivoli Storage Productivity Center, setting up a storage device for monitoring required performing tasks in different locations of the user interface, and each type of device required different steps. In Tivoli Storage Productivity Center v4.2 and later, all the tasks for setting up and configuring a device are included in the Configure Devices wizard.

Use the wizard to configure the following device types for monitoring:

- Storage subsystems
- Fabrics and switches
- Computers
- Tape libraries

You can add multiple devices of the same type using a single session of the wizard. For example, you can add an IBM DS8000, IBM XIV, IBM SAN Volume Controller, IBM Storwize V7000, IBM Storwize V7000 Unified, IBM SONAS, and

storage subsystems associated with CIM agents by using a single session of the wizard. You cannot configure multiple devices of different types at the same time. For example, you cannot add a storage subsystem and a fabric using a single session of the wizard.

Note: The Configure Devices wizard queries all schedules to identify whether a subsystem or entity is part of a scheduled probe. If a probe schedule has been defined for a global entity such as **All Computers** or **All Storage subsystems**, the Configure Devices wizard will not display a new device as an entity requiring configuration.

The pages in the wizard are divided into two sections:

Steps Displays the steps for completing in the wizard. An icon is displayed next to the current step. The wizard includes the following steps for configuring a device:

1. Select the device type that you want to configure.
2. Select the specific devices that you want to configure.
3. Determine how data is collected from the devices.
4. Review your selections and run the wizard.
5. Review the results of the wizard.

How you access the wizard determines on what step you begin. For example, if you click the **Configure Devices** icon in the toolbar or select the **Configure Devices** node in the navigation tree, you start by selecting the device type that you want to configure. If you click the **Add Storage Subsystem** button on the Disk Manager > Storage Subsystem panel, you begin by entering connection information about the storage systems that you want to configure.

Options

Displays the options and fields that you must select and complete for each step.

Configuring storage subsystems

Use the Configure Devices wizard to configure storage systems so that these storage systems can be monitored by Tivoli Storage Productivity Center. Storage subsystems that you configure using the wizard are automatically included in probe schedules and associated with alert conditions.

You can configure more than one storage system at a time by using the Configure Devices wizard. All the storage systems that you configure by using the wizard share probe schedule and alert conditions.

Depending on the type of storage system that you want to configure, you must define connection information for its native interface or associated CIMOM. Enter information in the wizard about the native interfaces for the following storage systems:

- System Storage DS8000
- SAN Volume Controller
- The XIV system
- Storwize V7000
- Storwize V7000 Unified
- IBM SONAS

For all other storage systems, you must enter information about the CIM agents that manage them.

To configure storage systems to be monitored by Tivoli Storage Productivity Center, complete the following steps:

1. Select one of the following ways to access the wizard:
 - To configure existing storage systems, click the **Configure Devices** icon on the toolbar. You cannot configure existing storage systems if they are already included in a probe schedule.
If you select this method, you must also complete the following steps:
 - a. On the Select Device Type page, select **Storage Subsystem**.
 - b. Click **Configure existing storage subsystems**. The page displays the storage systems that are known to Tivoli Storage Productivity Center. Continue to next step.
 - To only add storage systems, select one of the following methods and proceed to step 4:
 - In the navigation tree, expand **Disk Manager > Storage Subsystems** and click **Add Storage Subsystem**.
 - In the navigation tree, expand **IBM Tivoli Storage Productivity Center > Configuration Utility > Services > Storage Subsystem Interfaces** and click **Add**.
2. Select the existing storage systems that you want to configure. Storage systems that are already included in probe schedules are not displayed.
If the Credentials column indicates Missing or New Credentials Required for the storage systems that you select, you must update the storage system connection properties before you continue.
3. Enter valid connection properties for each storage system that you want to configure, and continue to step 6.
4. Configure the storage system or CIMOM connections. Tivoli Storage Productivity Center uses this information to connect to a storage system or CIMOM.
Enter information for a storage system or CIMOM, and then click **Add**. Repeat steps 1 - 3 to configure multiple storage systems or CIMOMs in the same session of the wizard.
Tivoli Storage Productivity Center will discover the storage systems managed by each of the CIMOM connections that you entered.
5. Select the storage systems that you want to configure. Any storage systems that were added in previous steps are automatically selected. Deselect storage systems that you do not want to configure.
6. Specify how data is collected for storage systems by including them in a monitoring group or template. Use monitoring groups and templates to determine when data is collected and what alert conditions are checked for the storage systems that you are configuring.
 - Select a monitoring group if you want to manage a collection of storage systems in the same manner. By default, the monitoring group that you select is associated with an existing probe schedule and set of alerts. When you select the group, its data collection schedule and alerts are automatically applied to all the storage systems that you are configuring. After completing the wizard, the storage systems remain in the group and you can use that group when working in other parts of the user interface, such as in reports.

Tip: Tivoli Storage Productivity Center provides a number of alerts and schedules associated with monitoring groups for storage systems. If you want to include a storage system in a user-defined group rather than in a default group, that group must be associated with a data collection schedule. If the user-defined group is not associated with a data collection schedule, it is not displayed in the list of monitoring groups from which you can select.

- Select a template if you want to create a unique probe schedule and set of alerts for a storage system rather than using the settings of an existing group.

The option that you select in this step determines the next step in the wizard.

- If you select a monitoring group, the following actions occur:
 - The storage systems are automatically included in a probe schedule to which the group is associated. For example, if you select Subsystem Basic Group, the storage systems are included in the Subsystem Basic Probe schedule. The Subsystem Basic Probe schedule runs one time per week on Tuesday at 1:00 am.
 - A set of alert conditions are applied to the storage systems in a group. When any of these conditions are detected during data collection, an alert is triggered. The group that you select determines the alerts that are associated with the storage systems. For example, if you select Subsystem Basic Group, 14 alert conditions are automatically checked for the storage systems.
 - The storage systems are included in the group. You can use that group in other parts of the user interface, such as in reports.
- If you select a monitoring template, the following actions occur:
 - The storage systems are automatically included in a probe schedule that is created for the template. For example, if you select the Basic Subsystem Monitoring template, the storage systems are automatically included in a new probe that runs one time per week on Tuesday. The probe is created after the wizard is completed. Tivoli Storage Productivity Center determines the time when the probe is run based on server load to ensure optimal performance.
 - A set of alert conditions is applied to storage systems in the template. When any of these conditions are detected during data collection, an alert is triggered. The template you select determines the alerts that are associated with the storage systems. For example, if you select to the Basic Subsystem Monitoring template, 14 alert conditions are automatically checked for the storage systems.
 - You must enter a prefix for the probe schedule and alerts that are created for the template. This prefix is used to identify the probe and alerts in the user interface. For example, if you select the Basic Subsystem Monitoring template and enter the prefix "ABC", a probe schedule named "*userID*.ABCssBasicProbe" and alerts such as "*userID*.ABCPool Discovered (Basic)" are created, where *userID* represents the ID of the user that is logged on to Tivoli Storage Productivity Center.

The new probe and alerts are displayed under the corresponding node in the navigation tree for you to view and edit. For example, a probe schedule created from a template in the wizard appears under the **IBM Tivoli Storage Productivity Center > Monitoring > Probes** node; new storage system alerts appear under the **Disk Manager > Alerting** node.

See Alerts and schedules associated with monitoring groups in the information center for a list of groups and templates with their associated schedules and alerts.

7. On the summary page of the wizard, review the selections that you made.

View the results of the device configuration and close the wizard. The devices that you configured are probed.

Configure storage subsystem connections page

Use this page in the Configure Devices wizard to enter or update connection properties for the storage systems and CIMOMs that you are configuring. Tivoli Storage Productivity Center uses this information to collect data about the storage systems.

Use the Configure storage system connections page to enter or update connection properties for the following storage entities. Depending on the type of storage system that you want to configure, you must define connection information for its native interface or associated CIMOM. Enter information on the page about the native interfaces for the following storage systems:

- System Storage DS8000
- XIV system
- SAN Volume Controller / Storwize V7000
- Storwize V7000 Unified
- IBM SONAS

The entity that is selected from the **Device Type** list determines the information that Tivoli Storage Productivity Center requires to connect to that entity. The following sections describe the fields for each entity.

Requirement: If you are adding or updating connection properties for a SAN Volume Controller or Storwize V7000 storage system, select **IBM SAN Volume Controller / IBM Storwize V7000** in the **Device Type** list. If you are adding or updating connection properties for a Storwize V7000 Unified storage system, select **IBM Storwize V7000 Unified**.

IBM DS8000

Enter or update the following information to have Tivoli Storage Productivity Center connect to a System Storage DS8000 storage system.

HMC Address

Enter the IP address or host name for the Hardware Management Console (HMC) that is used to manage the System Storage DS8000 storage system.

HMC2 Address (Optional)

Enter the IP address or host name for a second HMC that is used to manage the storage system.

Username

Enter the user name for logging in to the System Storage DS8000 element manager. The default user name is admin.

Note: This user name is the same as the ESSNI user name.

Password

Enter the password for the user name.

IBM XIV

Enter or update the following information to have Tivoli Storage Productivity Center connect to the XIV system.

IP Address

Enter the IP address or host name for the XIV system that you want to configure.

Username

Enter the user name for the XIV system.

Password

Enter the password for the user name.

IBM SAN Volume Controller / IBM Storwize V7000

Enter or update the following information to have Tivoli Storage Productivity Center connect to a SAN Volume Controller or Storwize V7000 storage system.

Software Version

Enter the version of the SAN Volume Controller or Storwize V7000 software that you want to configure.

IP Address

Enter the IP address or host name of the SAN Volume Controller or Storwize V7000. You can enter an IPv4 or IPv6 address, for example:

- 127.0.0.1
- [2001:DB8:0:0:0:0:0:0/32]

Select Key

A Secure Shell (SSH) key is required to access the storage system. Select one of the following actions for the SSH key:

- **Upload new key:** Select this option to upload a new key. Public key content is extracted from the private key that is provided in the **Private SSH Key** field and is uploaded to the SAN Volume Controller or Storwize V7000 cluster. If you select this option, the **Admin Username** and **Admin Password** fields are required.
- **Use existing uploaded key:** Select this option if you want to use a public key that was already uploaded by using a method other than through Tivoli Storage Productivity Center, such as the SAN Volume Controller or Storwize V7000 web interface. Enter the location of the private key file that corresponds to the public key in the **Private SSH Key** field. The **Admin Username** and **Admin Password** fields are not required if you select this option.

Administrator User Name

If you selected **Upload new key** in the **Select Key** list, enter the name of a user that has the Security Administrator role for the SAN Volume Controller or Storwize V7000 cluster that contains the storage system. Tivoli Storage Productivity Center uses this value during the setup process to contact the cluster and configure the SSH keys for the user. This user name must have privileges to modify other user accounts (Security Administrator role); otherwise, Tivoli Storage Productivity Center cannot configure the SSH keys.

If you selected the **Use existing uploaded key** option, this field is optional.

Note: The **Administrator User Name** applies to only SVC5.x+; it does not apply to SVC 4.x.

Administrator Password

Enter the password for the user name that you entered in the **Administrator User Name** field.

User Name

Enter an authorized user name for the storage system. Tivoli Storage Productivity Center uses this value when monitoring and configuring the storage system. If the user name does not exist it is created. This user name must be assigned the Administrator role.

This field is active if you selected **Upload new key** in the **Select Key** list and entered a user name and password in the **Administrator User Name** and **Administrator Password** fields.

Important: If you upload a new key for the existing SAN Volume Controller or Storwize V7000 user by using Tivoli Storage Productivity Center, that new key overrides the existing key for the user. If you have another application that is using that user (for example Tivoli Storage Productivity Center for Replication), the application cannot log on to SAN Volume Controller or Storwize V7000 because the key associated with the user is overwritten.

If you have this situation, create a user ID specifically for Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication on the SAN Volume Controller or Storwize V7000 cluster. Add the SAN Volume Controller or Storwize V7000 to Tivoli Storage Productivity Center for Replication using the new user ID and password (or change the user ID and password while the storage system is in the disconnected state).

Select User

Click this button to select from a list of existing users that are defined on the SAN Volume Controller or Storwize V7000.

Private SSH Key

By default, the private key entry includes `${device.conf}\tpc_svc.pem`, which represents the Tivoli Storage Productivity Center default key `tpc_svc.pem`. The `tpc_svc.pem` key is in the `conf` directory where the Device server is installed. If you want to use your own key, the key must exist on the system where you are running the Tivoli Storage Productivity Center user interface. The key is uploaded to the Tivoli Storage Productivity Center server system. To identify a new key, click **Browse**.

The private SSH key must be in OpenSSH format or in PuTTY (.ppk) format that is not password protected. If you have a password protected PuTTY .ppk key, use the PuTTYgen utility on a Windows operating system to convert that key to OpenSSH format by following these steps:

1. Start the PuTTYgen utility.
2. Load the .ppk file by clicking **File > Load private key**.
3. From the menu bar, click **Conversion > Export OpenSSH key conversion**.

This action saves the key in OpenSSH format.

Passphrase (Optional)

Enter the passphrase for the private SSH key. The passphrase protects your private SSH key file. If you do not have a passphrase, leave this field blank.

Storwize V7000 Unified

Enter or update the following information to have Tivoli Storage Productivity Center connect to a Storwize V7000 Unified storage system.

Storwize V7000 Unified is composed of Storwize V7000 and Storwize V7000 File Module. You can connect to only Storwize V7000 or to both Storwize V7000 and Storwize V7000 File Module. Storwize V7000 provides block-level data and Storwize V7000 File Module provides file-level data.

Storwize V7000

Use the information in this section to connect Storwize V7000.

IP Address

Enter the IP address or host name of the Storwize V7000 storage system. You can enter an IPv4 or IPv6 address, for example:

- 127.0.0.1
- [2001:DB8:0:0:0:0:0/32]

Select Key

An SSH key is required to access Storwize V7000. Select one of the following actions for the SSH key:

- **Upload new key:** Select this option to upload a new key. Public key content is extracted from the private key that is provided in the **Private SSH Key** field and is uploaded to the Storwize V7000 cluster. If you select this action, the **Administrator User Name** and **Administrator Password** fields are required.
- **Use existing uploaded key:** Select this option if you want to use a public key that was already uploaded by using a method other than through Tivoli Storage Productivity Center, such as the Storwize V7000 web interface. Enter the location of the private key file that corresponds to the public key in the **Private SSH Key** field. The **Administrator User Name** and **Administrator Password** fields are not required if you select this option.

Administrator User Name

If you selected **Upload new key** in the **Select Key** list, enter the name of a user that has the Security Administrator role for the Storwize V7000 cluster that contains the storage system. Tivoli Storage Productivity Center uses this value during the setup process to contact the cluster and configure the SSH keys for the user. This user must belong to the Security Administrator role, which has privileges to modify other user accounts. Otherwise, Tivoli Storage Productivity Center cannot configure the SSH keys.

If you selected **Use existing uploaded key**, this field is optional.

Administrator Password

Enter the password for the user name that you entered in the **Administrator User Name** field.

If you selected **Use existing uploaded key**, this field is optional.

User Name

Enter an authorized user name for the storage system. Tivoli Storage Productivity Center uses this value to monitor and configure the storage system. If the user name does not exist, it is created. This user name must belong to the Administrator role.

This field is active if you selected **Upload new key** in the **Select Key** list and entered a user name and password in the **Administrator User Name** and **Administrator Password** fields. **Important:** If you upload a new key for the existing Storwize V7000 user by using Tivoli Storage Productivity Center, that new key overrides the existing key for the user name. If you have another application that is using that user name (for example Tivoli Storage Productivity Center for Replication), the application cannot log on to Storwize V7000 because the key associated with the user is overwritten.

If this situation occurs, create a user ID specifically for Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication on the Storwize V7000 cluster. Add the Storwize V7000 storage system to Tivoli Storage Productivity Center for Replication using the new user ID and password (or change the user ID and password while the storage system is in the disconnected state).

Select User

Select a user name from the list. The names are defined on Storwize V7000.

Private SSH Key

By default, the private key entry includes `${device.conf}\tpc_svc.pem`, which represents the Tivoli Storage Productivity Center default key `tpc_svc.pem`. The `tpc_svc.pem` key is in the `conf` directory where the Device server is installed. If you want to use your own key, the key must exist on the system where you are running the Tivoli Storage Productivity Center user interface. The key is uploaded to the Tivoli Storage Productivity Center server system. To identify a new key, click **Browse**.

The private SSH key must be in OpenSSH format or in PuTTY (.ppk) format that is not password protected. If you have a password protected PuTTY .ppk key, use the PuTTYgen utility on a Windows operating system to convert that key to OpenSSH format by following these steps:

1. Start the PuTTYgen utility.
2. Load the .ppk file by clicking **File > Load private key**.
3. From the menu bar, click **Conversion > Export OpenSSH key conversion**.

This action saves the key in OpenSSH format.

Passphrase (Optional)

Enter the passphrase for the private SSH key. The passphrase protects your private SSH key file. If you do not have a passphrase, leave this field blank.

Storwize V7000 File Module (Optional)

To connect to the Storwize V7000 Unified system, select the check box.

Depending on the configuration of Storwize V7000 Unified, the information that is required to connect to this system can be the same as or different from the information for Storwize V7000. The fields and field descriptions are the same as the fields in the Storwize V7000 pane with the following exceptions:

- The **IP Address** field is not displayed because the IP address is retrieved automatically.
- The **User Name** field is always required.
- The private SSH key must be in OpenSSH format. The PuTTY format is not supported. You can use the PuTTYgen utility to convert the key to the OpenSSH format.
- The **Select Key** list contains the option **Use user password** in addition to the **Upload new key** and **Use existing uploaded key** options. If you select **Use user password**, the **User Password** field is enabled and you can use a user name and password to connect to the file module.
If you select **Use user password**, the **Administrator User Name** and **Administrator Password** fields are optional.

IBM SONAS

Enter or update the following information so that Tivoli Storage Productivity Center can connect to an IBM SONAS system.

IP Address

Enter the IP address or host name of the IBM SONAS system. You can enter an IPv4 or IPv6 address, for example:

- 127.0.0.1
- 2001:DB8:0:0:0:0:0/32

Select Key

You must use a Secure Shell (SSH) key to access the storage system. Select one of the following actions for the SSH key:

- **Upload new key:** Select this option to upload a new key. Public key content is extracted from the private key that is provided in the **Private SSH Key** field and is uploaded to the IBM SONAS system. If you select this action, you must specify a value for the **Administrator User Name** and **Administrator Password** fields.
- **Use existing uploaded key:** Select this option if you want to use a public key that was already uploaded by using a method outside of Tivoli Storage Productivity Center, such as the IBM SONAS system web interface. Enter the location of the private key file that corresponds to the public key in the **Private SSH Key** field. You do not have to specify a value in the **Administrator User Name** and **Administrator Password** fields.
- **Use user password:** Select this option if you want to use a user name and password to connect to the file module. If you select this option, the **User Password** field is enabled and you can use a user name and password to connect to the file module.

If you select **Use user password**, the **User Password** field is enabled and the **Administrator User Name** and **Administrator Password** fields are optional.

Administrator User Name

If you selected **Upload new key** in the **Select Key** list, enter the name of a user that has the Security Administrator role for the IBM SONAS system. Tivoli Storage Productivity Center uses this value during the setup process to contact the system and configure the SSH keys for the user. This user name must have privileges to modify other user accounts (Security Administrator role); otherwise, Tivoli Storage Productivity Center cannot configure the SSH keys.

If you selected **Use existing uploaded key**, this field is optional.

Administrator Password

Enter the password for the user name that you entered in the **Administrator User Name** field.

User Name

Enter an authorized user name for the storage system. Tivoli Storage Productivity Center uses this value when it configures and monitors the storage system. If the user name does not exist it is created. This user name must have the Administrator role.

This field is active if you selected **Upload new key** in the **Select Key** list and entered a user name and password in the **Administrator User Name** and **Administrator Password** fields.

Select User

Select from a list of existing user names that are defined on the IBM SONAS system.

Private SSH Key

By default, the private key entry includes `${device.conf}\tpc_svc.pem`, which represents the Tivoli Storage Productivity Center default key `tpc_svc.pem`. The `tpc_svc.pem` key is in the `conf` directory where the Device server is installed. If you want to use your own key, the key must exist on the system where you are running the Tivoli Storage Productivity Center user interface. The key is uploaded to the Tivoli Storage Productivity Center server system. To identify a new key, click **Browse**.

The private SSH key must be in OpenSSH format or in PuTTY (.ppk) format that is not password protected. If you have a password protected PuTTY .ppk key, use the PuTTYgen utility on a Windows operating system to convert that key to OpenSSH format by following these steps:

1. Start the PuTTYgen utility.
2. Load the .ppk file by clicking **File > Load private key**.
3. From the menu bar, click **Conversion > Export OpenSSH key conversion**.

This action saves the key in OpenSSH format.

Passphrase (Optional)

Enter the passphrase for the private SSH key. The passphrase protects the private SSH key file. If you do not have a passphrase, leave this field blank.

User Password

Enter the password for the user name that you entered in the **User Name** field.

Other (CIMOM)

Enter or update connection properties for a CIMOM. Tivoli Storage Productivity Center uses this information to discover the storage systems that are managed by a CIMOM. You can then configure those storage systems as part of the wizard session.

The connection information that you enter for a CIMOM is similar to the information that you enter when adding a CIMOM on the Add CIMOM page. *Adding a CIM agent connection* in the information center for descriptions of the

fields on this page. For information about the fields on this page, go to the information center. Search for **Adding a CIM agent connection**.

Discovered Storage Systems

After you run a discovery job, a table is displayed on the "Configure storage subsystem connections" page. This table provides information about the storage systems that were discovered for the selected storage system type.

You can remove the connection to a storage system by clicking the row for the system in the table, and then clicking **Remove**.

Configure fabrics and switches

Use the Configure Devices wizard to set up fabrics and switches for monitoring by Tivoli Storage Productivity Center. Fabrics and switches that you configure in the wizard are automatically included in probe schedules and associated with alert conditions.

You can use the wizard to configure the following fabric and switch types for monitoring:

- Brocade
- Cisco
- QLogic

You can configure one or more fabrics using a single session of the Configure Devices wizard. All the fabrics and switches that you configure in a single pass through the wizard share the same probe schedule and alert conditions.

To configure fabrics and switches using the wizard, complete the following steps:

1. Access the wizard based on the task you want to perform:
 - If you want to add new or configure existing fabrics and switches, click the configure devices icon in the toolbar or select the **Configure Devices** node in the navigation tree. If you use this method, you must also complete the following steps:
 - a. Select **Fabrics/Switches** on the Select Device Type page.
 - b. Determine whether to add new fabrics or configure existing fabrics that are known to Tivoli Storage Productivity Center.
 - To add fabrics and switches, select **Add and configure new fabrics** and go to step 3.
 - To configure existing fabrics and switches, select **Configure existing fabrics**. The page displays the fabrics that are known to Tivoli Storage Productivity Center. Go to step 2.
 - If you only want to add new fabrics and switches, go to **Fabric Manager > Fabrics** and click the **Add Fabric** button. Go to step 3.
2. Select the existing fabrics that you want to configure. Fabrics that are already included in probe schedules are not displayed. Go to step 6 after selecting the fabrics that you want to configure.
3. Configure CIMOM or agent connections. The fields for entering connection information for a fabric changes based on the type of fabric you are adding. Tivoli Storage Productivity Center uses the information you enter in this step to connect to a fabric or CIMOM. Click **Add** to configure multiple fabrics in a session of the wizard.

Note: Depending on the type of fabric you are adding, the wizard recommends whether to enter CIMOM or agent information. For example, if you are configuring a Brocade fabric, the recommended way to administer it is with a CIMOM. If a CIMOM is not available, enter connection information for an out of band fabric agent.

4. Discover the fabrics and fabrics managed by the CIMOMs that you entered in the previous step. Tivoli Storage Productivity Center collects initial configuration information for these entities.
5. Select the fabrics that you want to configure for monitoring. Any fabrics that were added in previous steps are automatically selected. Clear the checkboxes for the fabrics that you do not want to configure.
6. Specify how data is collected for fabrics by including them in a monitoring group or template. Use monitoring groups and templates to determine when data is collected and what alert conditions are checked for the fabrics that you are configuring.
 - Select a monitoring group if you want to manage a collection of fabrics in the same manner. Each monitoring group that you can select is associated with an existing probe schedule and set of alerts. When you select a group, its data collection schedule and alerts are automatically applied to all the fabrics that you are configuring. After completing the wizard, the fabrics remain in the group and you can use that group when working in other parts of the user interface, such as in reports.

Tip: Tivoli Storage Productivity Center provides a number of for fabrics. If you want to include a fabric in a user-defined group rather than in a default group, that group must be associated with a data collection schedule. If the user-defined group is not associated with a data collection schedule, it is not displayed in the list of monitoring groups from which you can select.

- Select a template if you want to create a unique probe schedule and set of alerts for a fabric rather than using the settings of an existing group.

The option that you select in this step determines the field and pages that appear next in the wizard.

- If you select a monitoring group, the following actions occur:
 - The fabrics are automatically included in a probe schedule to which the group is associated. For example, if you select Fabric Basic Group, the fabrics are included in the Fabric Basic Probe schedule. The Fabric Basic Probe schedule runs two times per week on Monday and Wednesday at 3:00 am.
 - A set of alert conditions it applied to the fabrics in a group. When any of these conditions are detected during data collection, an alert is triggered. The group that you select determines the alerts that are associated with the fabrics. For example, if you select Fabric Basic Group, nine alert conditions are automatically checked for the fabrics.
 - The fabrics are included in the group. You can use that group in other parts of the user interface, such as in reports.
- If you select a monitoring template, the following actions occur:
 - The fabrics are automatically included in a probe schedule that is created for a template. For example, if you select the Basic Fabric Monitoring template, the fabrics are automatically included in a new probe that runs two times per week on Monday and Wednesday. The probe is created after the wizard is completed. Tivoli Storage Productivity Center determines the time when the probe is run based on server load to ensure optimal performance.

- A set of alert conditions is applied to the fabrics in a template. When any of these conditions are detected during data collection, an alert is triggered. The template you select determines the alerts that are associated with the fabrics. For example, if you select the Basic Fabric Monitoring template, nine alert conditions are automatically checked for the fabrics.
- You must enter a prefix for the probe schedule and alerts that are based on a template. This prefix is used to identify the probe and alerts in the user interface. For example, if you select the Basic Fabric Monitoring template and enter the prefix "ABC", a probe schedule named "*userID*.ABCfabBasicProbe" and alerts such as "*userID*.ABCFabric State Change (Basic)" are created, where *userID* represents the ID of the user that is logged on to Tivoli Storage Productivity Center.

The new probe and alerts are displayed under the corresponding node in the navigation tree for you to view and edit. For example, a probe schedule created from a template in the wizard appears under the **Tivoli Storage Productivity Center >Monitoring > Probes** node; new storage system alerts appear under the **Fabric Manager > Alerting** node.

See for a list of groups and templates with their associated schedules and alerts.

7. On the summary page of the wizard, review the selections that you made.

View the results of the device configuration and close the wizard. The devices that you configured are probed.

Configure computers

Use the Configure Devices wizard to set up computers for monitoring by Tivoli Storage Productivity Center. Computers that you configure in the wizard are automatically included in probe and ping schedules and associated with alert conditions.

Use the Configure Devices wizard to configure computers that meet the following conditions:

- A Storage Resource agent was deployed on the computer. For information about deploying Storage Resource agents, see Storage Resource Agent Deployments.
- The computer is not included in an existing probe schedule. Devices that are included in existing probe schedules are not displayed in the Configure Devices wizard because they are considered to be already configured for monitoring.

Important: When computers are added to Tivoli Storage Productivity Center, they are automatically included in the default computer group, and that group is included in the default probe schedule (TPCUser.Default Probe). Because the computers are included in TPCUser.Default Probe, they are considered to be configured for monitoring and are not displayed in the wizard.

To use the Configure Devices wizard to configure computers that were included in TPCUser.Default Probe, move them into a group that is not associated with an existing probe schedule. To move them into a new group, complete the following steps:

1. In the navigation tree, go to **Data Manager > Monitoring > Groups > Computer**.
2. Right click **Computer** and select **Create Computer Group**.
3. Select the computers that you want to configure in the Configure Devices wizard and click the right arrow icon to move them into the new group.

4. Click the save icon in the toolbar to save the group. Do not associate this group with a probe schedule.
5. Start the Configure Devices wizard to the configure computers. The computers in the new group are displayed in the wizard.

You can configure one or more computers using a single session of the Configure Devices wizard. All the computers that you configure in a single pass through the wizard share the same probe schedule, ping schedule, and alert conditions.

To configure computers using the wizard, complete the following steps:

1. Click the configure devices icon in the toolbar or select the **Configure Devices** node in the navigation tree.
2. Select **Computers** on the Select Device Type page.
3. Select the computers that you want to configure. Computers that are already included in an existing probe schedule are not displayed.
4. Specify how data is collected for computers by including them in a monitoring group or template. Use monitoring groups and templates to determine when data is collected and what alert conditions are checked for the computers that you are configuring.
 - Select a monitoring group if you want to manage a collection of computers in the same manner. Each monitoring group that you can select is associated with existing probe and ping schedules and a set of alerts. When you select a group, its data collection schedules and alerts are automatically applied to all the computers that you are configuring. After completing the wizard, the computers remain in the group and you can use that group when working in other parts of the user interface, such as in reports.

Tip: Tivoli Storage Productivity Center provides a number of default monitoring groups for computers. If you want to include a computer in a user-defined group rather than in a default group, that group must be associated with a data collection schedule. If the user-defined group is not associated with a data collection schedule, it is not displayed in the list of monitoring groups from which you can select.

- Select a template if you want to create a unique probe schedule, ping schedule, and set of alerts for a computer rather than using the settings of an existing group.

The option that you select in this step determines the field and pages that appear next in the wizard.

- If you select a monitoring group, the following actions occur:
 - The computers are automatically included in probe and ping schedules to which the group is associated. For example, if you select Computer Basic Group, the computers are included in the Computer Basic Probe and Computer Basic Ping schedules. The Computer Basic Probe schedule runs two times per week on Monday and Wednesday at 3:00 am. The Computer Basic Ping schedule runs every 3 hours.
 - A set of alert conditions is applied to the computers in a group. When any of these conditions are detected during data collection, an alert is triggered. The group that you select determines the alerts that are associated with the computers. For example, if you select Computer Basic Group, seven alert conditions are automatically checked for the computers.
 - The computers are included in the group. You can use that group in other parts of the user interface, such as in reports.

- If you select a monitoring template, the following actions occur:
 - The computers are automatically included in a probe and ping schedule that are created for the template. For example, if you select the Basic Computer Monitoring template, the computers are automatically included in a new probe that runs two times per week on Monday and Wednesday and a ping that runs every three hours. The probe and ping are created after the wizard is completed. Tivoli Storage Productivity Center determines the time when the probe is run based on server load to ensure optimal performance.
 - A set of alert conditions is applied to the computers in a template. When any of these conditions are detected during data collection, an alert is triggered. The template you select determines the alerts that are associated with the computers. For example, if you select to the Basic Computer Monitoring template, seven alert conditions are automatically checked for the computers.
 - You must enter a prefix for the probe schedules, ping schedules, and alerts that are based on a template. This prefix is used to identify these definitions in the user interface. For example, if you select the Standard Computer Monitoring template and enter the prefix "ABC", the following definitions are created in the user interface after the wizard is completed:
 - A probe schedule named "*userID*.ABCcompStandardProbe"
 - A ping schedule named "*userID*.ABCcompStandardPing"
 - Seven alerts that follow this naming convention: *userID*.ABCStatus Change Online (Standard)

Where *userID* represents the ID of the user that is logged in to Tivoli Storage Productivity Center.

The new schedule definitions and alerts are displayed under the corresponding node in the navigation tree for you to view and edit. For example, a probe schedule created from a template in the wizard is displayed under the **Tivoli Storage Productivity Center > Monitoring > Probes** node; a ping schedule is displayed under the **Data Manager > Monitoring > Pings** node; new storage system alerts appear under the **Data Manager > Alerting** node.

See "Alerts and schedules associated with monitoring groups and templates" for a list of groups and templates with their associated schedules and alerts.

5. Review the selections that you made in the wizard. The Review user selections page includes the following information:
 - A list of devices that you are configuring.
 - The name of the monitoring group that you selected. This value is not displayed if you selected a template.
 - The name of the probe and ping schedules that are created based on the template you selected. This value is not displayed if you selected a monitoring group.
 - Information about the schedules created for a template.
 - The names of the alerts that are created based on the template you selected. This value is not displayed if you selected a monitoring group.

View the results of the device configuration and close the wizard. The devices that you configured are probed.

Configure tape libraries

Use the Configure Devices wizard to set up tape libraries for monitoring by IBM Tivoli Storage Productivity Center. Tape libraries that you configure in the wizard are automatically included in probe schedules.

You can configure one or more tape libraries using a single session of the Configure Devices wizard. All the tape libraries that you configure in a single pass through the wizard share the same probe schedule.

To configure tape libraries using the wizard, complete the following steps:

1. Access the wizard based on the task you want to perform:
 - If you want to add new or configure existing tape libraries, click the configure devices icon in the toolbar or select the **Configure Devices** node in the navigation tree. If you use this method, you must also complete the following steps:
 - a. Select **Tape Libraries** on the Select Device Type page.
 - b. Determine whether to add new tape libraries or configure existing tape libraries that are known to Tivoli Storage Productivity Center.
 - To add tape libraries, select **Add and configure new tape libraries** and go to step 3.
 - To configure existing tape libraries, select **Configure existing tape libraries**. The page displays the tape libraries that are known to Tivoli Storage Productivity Center. Go to step 2.
 - If you only want to add new tape libraries, go to **Tape Manager > Tape Libraries** and click the **Add Tape Library** button. Go to step 3.
2. Select the existing tape libraries that you want to configure. Tape libraries that are already included in probe schedules are not displayed. Go to step 6 after selecting the tape libraries that you want to configure.
3. Configure CIMOM connections for tape libraries. Tivoli Storage Productivity Center uses the information you enter in this step to connect to a CIMOM. Click **Add** to configure multiple CIMOMs in a session of the wizard.

The connection information that you enter for a CIMOM is like the information that you enter on the Add CIMOM panel.
4. Discover the tape libraries that are managed by the CIMOMs you entered in the previous step. Tivoli Storage Productivity Center uses the information you enter in this step to connect to a CIMOM. Click **Add** to configure multiple CIMOMs in a session of the wizard.
5. Select the tape libraries that you want to configure for monitoring. Any tape libraries that were added in previous steps are automatically selected. Deselect tape libraries that you do not want to configure.
6. Specify how data is collected for tape libraries by including them in a monitoring group or template. Use monitoring groups and templates to determine when data is collected for the tape libraries that you are configuring.
 - Select a monitoring group if you want to manage a collection of tape libraries in the same manner. Each monitoring group that you can select is associated with an existing probe schedule. When you select a group, its data collection schedule is automatically applied to all the tape libraries that you are configuring. After completing the wizard, the tape libraries remain in the group and you can use that group when working in other parts of the user interface, such as in reports.

Tip: Tivoli Storage Productivity Center provides a number of "Alerts and schedules associated with monitoring groups and templates" for tape libraries. If you want to include a tape library in a user-defined group rather than in a default group, that group must be associated with a data collection schedule. If the user-defined group is not associated with a data collection schedule, it is not displayed in the list of monitoring groups from which you can select.

- Select a template if you want to create a unique probe schedule for a tape library rather than using the settings of an existing group.

The option that you select in this step determines the field and pages that appear next in the wizard.

- If you select a monitoring group, the following actions occur:
 - The tape libraries are automatically included in a probe schedule to which the group is associated. For example, if you select Tape Basic Group, the tape libraries are included in the Tape Basic Probe schedule. The Tape Basic Probe schedule runs one time per week on Sunday at 4:00 am.
 - The tape libraries are included in the group. You can use that group in other parts of the user interface, such as in reports.
- If you select a monitoring template, the following actions occur:
 - The tape libraries are automatically included in a probe schedule that is created for a template. For example, if you select the Basic Tape Monitoring template, the tape libraries are automatically included in a new probe that runs one time per week on Sunday. The probe is created after the wizard is completed. Tivoli Storage Productivity Center determines the time when the probe is run based on server load to ensure optimal performance.
 - You must enter a prefix for the probe schedule and alerts that are based on a template. This prefix is used to identify the probe and alerts in the user interface. For example, if you select the Basic Tape Monitoring template and enter the prefix "ABC", a probe schedule named "*userID*.ABCtapeBasicProbe" is created, where *userID* represents the ID of the user that is logged in to Tivoli Storage Productivity Center.

The new probe is displayed under the corresponding node in the navigation tree for you to view and edit. For example, a probe schedule created from a template in the wizard appears under the **Tivoli Storage Productivity Center >Monitoring > Probes** node.

See "Alerts and schedules associated with monitoring groups and templates" for a list of groups and templates with their associated schedules and alerts.

7. On the summary page of the wizard, review the selections that you made.

View the results of the device configuration and close the wizard. The devices that you configured are probed.


Managing schedules, jobs, and runs

Use the Job Management panel to view and manage all the schedules, runs, and jobs related to the storage devices that are monitored by IBM Tivoli Storage Productivity Center.

Specifically, you can use the Job Management panel to perform the following tasks:

- View information about all schedules, runs, and jobs in one central location. In previous versions of the product, this information was available only through nodes under each schedule type in the navigation tree.
- Identify all the schedules that include a particular storage device.
- Identify all the schedules, runs, and jobs that are encountering problems when executing.
- Ensure that a storage device is fully monitored by Tivoli Storage Productivity Center.
- Edit, rename, delete, and run schedules. You must be logged in as the user who created a schedule to perform some of these actions.

Use either of the following methods to access the Job Management panel:

- Click the job management icon  in the toolbar.
- Open the definition of a schedule and click the **View Job History** button to view the runs and jobs related to the schedule. For example, in the navigation tree, click **IBM Tivoli Storage Productivity Center > Monitoring > Probes > TPCUser.Default Probe** to view the probe definition panel and click **View Job History**.

The following terms are used on this panel:

- **Schedule:** A schedule is a definition of a monitoring task that collects data about storage resources. Monitoring tasks include scans, probes, pings, quotas, discovery, and performance monitors. For example, define a probe schedule to collect data about storage subsystems; define a scan schedule to collect file-level information about hosts.
- **Run:** A run is a single invocation of a schedule.
- **Job:** A job is a unit of work within a run for a schedule. One run might have multiple jobs. For example, a subsystem probe schedule has one job created per subsystem for which it collects data.

The Job Management panel is divided into the following sections:

Filters section

Use this section to perform the following tasks:

- Determine the type of storage entities whose schedules are displayed on the panel.
- View recommendations for storage entities that are not fully monitored by Tivoli Storage Productivity Center.

Filtering the contents of this panel based on entity type is optional. You select **Show All** from the **Entity Type** menu to display all the schedules defined in Tivoli Storage Productivity Center.

Schedules section

Use this section to manage the schedules related to the entities you selected in the **Filters** section. You can perform the following actions for each schedule:

- View and edit the definition of a schedule.
- Run a schedule immediately.
- Rename a schedule.
- Delete a schedule.
- Display schedules that have a certain status, such as Success, Running, or Failed.

- View basic information about a schedule, including the ID of the user who created it, description, job type (for example, ping, probe, or scan), the last time the schedule was run, and the next time the schedule is run.

Jobs for Selected Schedule section

Use this section to manage each run of a schedule and the jobs associated with a run. You can perform the following tasks for jobs and runs:

- View the status of a run and the date and time when it was executed.
- View information about the jobs associated with each run, including the name of the computer on which a job was run, the type of the monitored device, and the status of the job.
- View the log files for one or more jobs at a time.
- Display jobs that have started or completed within a specified length of time. For example, display jobs that have completed in the last two weeks.
- Display jobs that have a certain status, such as Success, Running, or Failed.

Click and drag the horizontal lines on the panel to resize each section. Click the names of columns in the **Schedules** and **Jobs for Selected Schedule** sections to sort the contents of lists based on the clicked column.

Filters

Refresh All

Click this button to update all the lists that are displayed on the panel.

Entity Type

Select the type of storage entity whose related schedule information you want to view. By default all schedules are displayed when you first access this panel. When you select a specific entity type, the following actions take place:

- The schedules, runs, and jobs related to the entity type are displayed in the appropriate sections on the panel.
- The names of the devices that belong to that entity type are displayed in the **Entity Name** list box.
- The text that is displayed above the **Show Recommendations** button indicates if the devices are fully monitored.

You can select from the following entity types in the **Entity Type** menu:

Show All

(default) View all the schedules defined in Tivoli Storage Productivity Center.

Storage Subsystem

View the schedules for storage subsystems that are monitored by Tivoli Storage Productivity Center.

Fabric View the schedules for fabrics that are monitored by Tivoli Storage Productivity Center.

Computer

View the schedules for computers that are monitored by Tivoli Storage Productivity Center.

Switch

View the schedules for switches that are monitored by Tivoli Storage Productivity Center.

Tape Library

View the schedules for tape libraries that are monitored by Tivoli Storage Productivity Center.

Monitoring Group

View the schedules that contain monitoring groups.

Storage Resource Group

View the schedules for all the devices contained within a particular storage resource group.

Entity Name Filter

Enter text to filter the names of entities that are displayed in the **Entity Name** list box. You can enter a combination of letters, numbers, dashes, underscores, and spaces. The filter text is not case sensitive. You can also enter the following wildcard characters:

- Use an asterisk (*) to represent zero or more unknown characters of any value. For example, enter **DS8*** to display devices of the selected type that begin with **DS** in their names, such as DS6000-A12-IBM, DS8000-B34-IBM, and DS5000-B35-IBM.
- Use a question mark to represent a single unknown character of any value. For example, enter **D?000-B3?-IBM** to include devices such as DS8000-B34-IBM, DS8000-B36-IBM, and DS5000-B35-IBM in the **Entity Name** list box.

Text that is displayed above the Show Recommendations button

Warning text is displayed above the **Show Recommendations** button if the devices selected in the **Entity Name** list box are not fully monitored.

A device is considered not fully monitored if it is not included in all data collection schedules that are available in the licensed edition of Tivoli Storage Productivity Center that is installed. For example, if using the Advanced edition of Tivoli Storage Productivity Center to monitor a host, that host must appear in the following schedule definitions to be considered fully monitored: ping, probe, and scan. If using the Disk edition of Tivoli Storage Productivity Center to monitor a storage subsystem, that storage subsystem must appear in the following schedule definitions to be considered fully monitored: probe.

If a device is considered not fully monitored, click **Show Recommendations** to view a list of recommended actions you can take to address the warning.

Show Recommendations

Click this button to view recommended actions for a selected device that is not fully monitored. From the Recommendations panel, you can perform the following tasks:

- View the schedules that are recommended for a device.
- Click the **Take Action** button to create the recommended schedule for a device.

Schedules**View Detail**

View the definition of a schedule that is selected in the schedule list. The

schedule definition panel opens when you click this button. Use this panel to view and edit the definition of a schedule.

Run Now

Run the selected schedule immediately. This option is only available if the job is not currently running.

Refresh

Update the status of the schedule that is selected in the schedule list. The process to refresh the status of a schedule might take a few moments based on the current server load and network speed.

action drop down menu

Use the options in this menu to perform actions against the selected schedule. The menu includes **View Detail**, **Run Now**, **Stop Now** (for performance monitor schedules) and the following actions:

Rename

Rename a schedule. You must be the owner of a schedule or the Tivoli Storage Productivity Center superuser to perform this action. This action is not available for the default schedules that are owned by TPCUser.

Delete Delete a schedule. When you delete a schedule, all the runs and jobs associated with the schedule are also deleted. You must be the owner of a schedule or the Tivoli Storage Productivity Center superuser to perform this action. This action is not available for the default schedules that are owned by TPCUser.

Only show schedules with status

Filter the schedules that appear in the schedules list based on their status. For example, use this menu to indicate that only schedules with a status of **Running** are displayed in the schedules list. The default setting displays all schedules.

You can filter schedules using the following statuses:

- Show All
- Show All Unsuccessful
- Success
- Running
- Warnings
- Failed
- The job is waiting
- The job is running

Tivoli Storage Productivity Center automatically updates the schedules list if you change the value selected in this menu.

schedules list

Use this list to view the schedules for the storage entities that are selected in the **Filters** section of the panel. The schedules that are displayed must also meet the filter selected in the **Only show schedules with status** menu. Click the name of a column to sort the contents of this list based on that column.

Some jobs that run in Tivoli Storage Productivity Center do not have a schedule associated with them. For example, jobs that are run through the command-line interface (CLI) do not have a schedule defined for them. To ensure that these job types are displayed on the Job Management panel, a

new schedule called **CLI and Event Driven Jobs** was created. Select this schedule to view the runs and jobs that do not have user-created schedules associated with them. The job log files for the CLI and Event Driven Jobs schedule are automatically retained for 90 days and then deleted.

The following information is displayed for each schedule:

Creator

The ID of the Tivoli Storage Productivity Center user who created the schedule.

Name The name of a schedule.

Description

The description of a schedule.

Job Type

The type of the schedule, such as Probe, Ping, Scan, Aggregator Job, CIMOM Discovery, or Subsystem Performance Monitor.

Last Run Finished

The date and time when the most recent run of a schedule completed. This column is blank if the schedule has never been run.

Last Run Status

The status of the most recent run of a schedule.

Next Run

The date and time when the next run of the schedule is planned. The value **On Demand** is displayed for schedules that are defined with the **Run Now** option. The value **Never** is displayed for schedules that are defined with the **Run Once** option and have already run.

row number

Appears in the lower right corner of the schedules list and indicates the number of rows in the schedules list. One row is displayed for each schedule. Use the horizontal lines on the panel to resize the section and increase or decrease the number of rows that are displayed.

Jobs for Selected Schedule

View Log File(s)

View the log files of jobs that are selected in the runs list. Use the Shift + click and Ctrl + click keyboard shortcuts to select multiple jobs. You can view up to 10 logs at one time. The Log File Viewer panel opens when you click this button or when you double-click a row for a job. Use this panel to view all the processing performed by the Tivoli Storage Productivity Center server and agents for jobs.

Note: The number of job log files that are retained for a schedule is determined by the values defined on the Log File and Cached Batch Report Retention panel under **Administrative Services > Configuration**. However, for CLI and Event Driven Jobs schedules, job log files are automatically retained for 90 days.

Show Jobs, in the last

Use these fields to determine the runs that are displayed for the selected schedule. Use the following conditions in the filter:

- Select **Started** or **Completed** from the **Show Jobs** menu to display runs based on whether they started or completed in the time frame specified on the **last** field.
- Specify a time frame for started or completed runs by entering a number in the **last field** and selecting a unit of measurement from the menu.

For example, select **Started**, enter **4**, and select **weeks** to display the runs for a schedule that have started within the last four weeks. Tivoli Storage Productivity Center automatically updates the runs list if you change the value selected in these fields.

Only show jobs with status

Determine the runs that are displayed in the runs list based on their status. For example, use this menu to indicate that only runs with a status of **Completed with warning(s)** are displayed. The default setting displays all runs for the selected schedule.

You can filter the display of runs using the following statuses:

- Show All
- Show All Unsuccessful
- The job is waiting
- Running
- The job is running
- Completed successfully
- Running with warning(s)
- Completed with warning(s)
- Completed with error(s)
- System error occurred
- The agent did not respond (to a request from the server)
- The agent is being upgraded
- The agent is not licensed
- Agent not found
- Agent could not be reached
- The job delay limit was exceeded.
- The agent disconnected
- Host was shutdown or recycled
- Cluster resource group unreachable: no IP
- Scripts are disabled for this agent
- Script was not found or has an invalid extension
- The agent disk space is low

Tivoli Storage Productivity Center automatically updates the runs list if you change the value selected in this menu.

runs list

Use this list to view the runs and jobs of the schedule selected in the **Schedules** section of the panel. This list displays runs based on the criteria specified in the **Show jobs** and **in the last** fields.

Click the + symbol next to an entry in the **Run** field to view the jobs associated with a run. Click the name of a column to sort the contents of this list based on that column.

The following information is displayed for each run or job:

Run Status

The status of a run.

Run The run number assigned to a run and the date and time when the run was started.

Computer

(jobs only) The name of the computer where a job is executed. For database probes, this column displays the name of the Storage Resource agent that is performing the probe. For storage subsystem probes, this column displays **localhost** because the probe is performed by the Tivoli Storage Productivity Center server.

Device

(jobs only) The name of the device to which a job applies. This column is blank if the job is not related to a specific device.

Status (jobs only) The status of a job.

Start Time

(jobs only) The date and time when a job started.

Finish Time

(jobs only) The data and time when a job completed.

Log File Name

(jobs only) The name of the log file for a job. Click the **View Log File** to view the contents of a job log.

row number

Appears in the lower right corner of the runs list and indicates the number of rows in the runs list. One row is displayed for each run. You can expand each run row to view the jobs related to that run. Use the horizontal lines on the panel to resize the section and increase or decrease the number of rows that are displayed.

Job Management - Recommendations

Use this window to view and apply the recommendations for storage devices that are not fully monitored by IBM Tivoli Storage Productivity Center.

A device is not fully monitored if it is not included in all the data collection schedules that are available in the licensed edition of Tivoli Storage Productivity Center that is installed. For example, if using the IBM Tivoli Storage Productivity Center Advanced license to monitor a host, that host must appear in the following schedule definitions to be considered fully monitored: ping, probe, and scan. If using IBM Storage Productivity Center Select Edition to monitor a storage subsystem, that storage subsystem must appear in a probe schedule to be considered fully monitored.

Recommendation Severity

Indicates whether a recommendation is a **Warning** or **Informational**.

- A **Warning** message indicates an action that you can take to more fully use Tivoli Storage Productivity Center to monitor a device. For example, when using a probe schedule in the IBM Tivoli Storage Productivity Center Advanced license to collect data about a storage system, a warning message is displayed. This message recommends that you also include the storage subsystem in a performance monitor schedule. You can then use the data collected by performance monitor to identify the bottlenecks in storage usage and performance for the storage subsystem.

- An **Informational** message is a recommendation to upgrade Tivoli Storage Productivity Center to an edition that includes additional monitoring functions. For example, when using probe schedules in IBM Storage Productivity Center Select Edition to collect data about a fabric, an informational message is displayed. This message indicates that if you upgrade to IBM Tivoli Storage Productivity Center Advanced, performance data can be collected about the fabric.

For **Warning** messages, select the corresponding row and click **Take Action** to perform the recommended action.

Entity Displays the name of a storage entity to which the recommendation applies.

Recommendation

Indicates the recommended actions for a storage entity. One row is displayed per recommendation. Click **Take Action** to resolve recommendations that have a **Warning** severity.

If multiple recommendations are displayed for a device, sometimes performing one of the recommended actions might resolve the other recommendations. For example, consider if these two recommendations are displayed for a fabric:

- In order to collect information about device configuration, define a probe schedule for fabric {0}
- To enable zone control, add a data source that is capable of zone control for fabric {0}

If you click **Take Action** for the first recommendation and include the fabric in a probe, the second recommendation is also resolved. Both recommendations are then automatically removed from the **Recommendation** list

Take Action

Perform the recommended action for a warning. When you click this button, the corresponding panel for the action is displayed. For example, if the recommended action is to include a storage subsystem in a performance monitor, the Create Subsystem Performance Monitor panel is displayed.

Cancel

Exit the window without taking action.

Configuration nodes

This section is organized to help you associate the information in the IBM Tivoli Storage Productivity Center GUI with an explanation of each node.

Expand **Administrative Services > Configuration**. You see these nodes.

Role-to-Group Mappings

One of the first tasks to perform after installing IBM Tivoli Storage Productivity Center is to assign roles to individuals who use the product. From the **Role-to-Group Mapping** node, you can map Tivoli Storage Productivity Center roles, such as Tape Operator or Fabric Administrator, to user groups that you create in your operating system or in your LDAP-compliant repository.

For example, on a Windows computer you can create user groups using the Administrative Tools control panel. When a user name is used to authenticate with Tivoli Storage Productivity Center, the group membership for the user determines the authorization level.

Role-based authorization

Operating system groups or LDAP groups (for example, groups contained in your LDAP-compliant repository) are associated with predefined roles. When a user ID is authenticated to Tivoli Storage Productivity Center through the GUI, CLI, or APIs, membership in an operating system or LDAP group determines the authorization level of the user.

Table 37 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 the Windows operating system, to the system group on the AIX operating system, or to the root group on the Linux operating system.

Table 37. Roles and authorization levels

Role	Authorization level
Superuser	Has full access to all Tivoli Storage Productivity Center functions.
Productivity Center administrator	Has full access to operations in the Administration section of the GUI.
Disk administrator	Has full access to Tivoli Storage Productivity Center disk functions.
Disk operator	Has access to reports only for Tivoli Storage Productivity Center disk functions. These reports include information about tape devices.
Fabric administrator	Has full access to Tivoli Storage Productivity Center Fabric Manager functions.
Fabric operator	Has access to reports only for Tivoli Storage Productivity Center Fabric Manager functions.
Data administrator	Has full access to Tivoli Storage Productivity Center Data Manager functions.
Data operator	Has access to reports only Tivoli Storage Productivity Center Data Manager functions.
Tape administrator	Has full access to Tivoli Storage Productivity Center tape functions.
Tape operator	Has access to reports only for tape functions.

Notes:

1. If a user has multiple roles, the authorization level is a combination of the levels for each of the roles.
2. If a user is not a member of any of the roles listed, no access is granted to that user.
3. For enterprise-rollup reports, you need superuser or Tivoli Storage Productivity Center administrator authority to do the following actions:

- Add, remove, or modify the Tivoli Storage Productivity Center subordinate server that the master server is monitoring.
- Create or run probe jobs that include Tivoli Storage Productivity Center subordinate servers.

Any Tivoli Storage Productivity Center user can generate enterprise-rollup reports.

4. You need to assign the appropriate Tivoli Integrated Portal roles to other Tivoli Storage Productivity Center users and groups who are added later to the system.
5. For more information about Tivoli Integrated Portal roles, see “Tivoli Integrated Portal and Tivoli Common Reporting Roles” on page 219.

How to associate a user and group with IBM Tivoli Storage Productivity Center

This section provides information about how to associate a user and group with IBM Tivoli Storage Productivity Center.

Creating a system group on Windows:

You must first create an operating system group on your computer, and then assign the users to the group. Then you can use the GUI to associate the group with the IBM Tivoli Storage Productivity Center role.

For Windows, first create the users you want to have access to the various Tivoli Storage Productivity Center roles. Complete the following steps:

1. Go to **Start > Settings > Control Panel > Administrative Tools > Computer Management > Local Users and Groups**.
2. Right-click on **Users** and click **New Users**.
3. The New User panel is displayed. Enter a user name, description (optional), and password (twice). Click **Create**.
4. Repeat steps 3 and 4 for each new user you want to create. Then create a system group on Windows:
 - a. Go to **Start > Settings > Control Panel > Administrative Tools > Computer Management > Local Users and Groups**.
 - b. Right-click on **Groups** and click **New Group**.
 - c. The New Group panel is displayed. Enter a Group Name and Description (optional). Click **Add**.
 - d. The Select Users or Groups panel is displayed. Enter or select the users or groups you want to add to the new group you have created. Click **Add**. Then click **Create**.

Creating a system group on UNIX or Linux:

For UNIX or Linux, to add new users to the system, follow the steps in this topic.

Complete the following steps:

1. Create a name for the user.
2. The new user must have a home directory allocated for them.
3. That user must be placed in the group file. Refer to your UNIX or Linux documentation for information about how to create a system group.

Associating the user groups to IBM Tivoli Storage Productivity Center roles:

After you have created the operating system group and assigned users to that group, you must assign the groups to IBM Tivoli Storage Productivity Center roles.

Complete the following steps:

1. Open the Tivoli Storage Productivity Center GUI.
2. Go to **Administrative Services > Configuration > Role-to-Group Mappings**.
3. On the content pane, click **Edit** for the role you want to associate with the group you have created.
4. The Edit Group dialog is displayed. Enter the group you want to associate with the Tivoli Storage Productivity Center role. Click **OK**. This associates the operating system group with the Tivoli Storage Productivity Center role.
5. Click **File > Save** to save the mappings.

License Keys

You must set database permissions to monitor databases with the Data Manager.

You must have the following permissions set to monitor the databases with Data Manager:

Table 38. Database permissions

Database	Permissions
DB2	db2admin
Microsoft SQL Server	public
Oracle	DBA
Sybase	SA level

To monitor Oracle databases with Tivoli Storage Productivity Center, the Oracle user needs DBA authority. A non-DBA Oracle user with “create session”, “select any dictionary”, “analyze any”, and “analyze any dictionary” roles can still monitor the Oracle database through Tivoli Storage Productivity Center. However, Tivoli Storage Productivity Center is not able to obtain free space information for database objects. During the database registration, when a non-DBA Oracle user is used with the roles mentioned in the preceding list, the following warning message displays:

```
RDBMS login does not have dba privilege. Freespace will not be calculated for system objects.
```

Clicking **OK** continues with the registration process. This message also is also displayed in the scan logs when scan jobs are run.

Through the License Keys node, you can administer the license keys for Data Manager for Databases. Specifically, you can:

- Assign Data Manager for Databases licenses to your Storage Resource agents.
- Edit, add, and delete the instances within your organization that you want to monitor.
- View the number of Data Manager for Databases licenses.
- View the number of Data Manager for Databases licenses that are not currently assigned (unused) to agents.

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

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

Assigning Data Manager database licenses to installed agents

This topic describes how to assign Data Manager database licenses to an agent.

To assign a Data Manager database license to an agent, complete the following steps:

1. Expand **Administrative Services > Configuration > License Keys**. The License Editor window is displayed.
2. Click the icon for the **IBM Tivoli Storage Productivity Center for Data - Databases** row. The Tivoli Storage Productivity Center for Data - Databases License Editor window is displayed. The field and button descriptions for the Licensing tab are:

Select All

Selects all the **Licensed** boxes.

Deselect All

Removes all the licenses for Data Manager - Databases.

Computer

Displays all the computers on which a Data Manager agent is installed.

OS Type

Displays the operating system of the computer where the agent is installed.

Domain

Displays the domain of a computer where the agent is installed.

Tree Name

For NetWare devices that are managed by Storage Resource agents from previous releases.

Licensed

Contains a check box that indicates whether a computer is licensed for use with Data Manager - Databases.

If you have unused Data Manager - Databases licenses, continue to the next step. To unassign licenses from agents, see "Unassigning Data Manager - Databases license" on page 273.

3. Click the check box in the **Licensed** column next to the computer with the instance you want to monitor.
4. Click **File > Save** to save the updated license settings. If an instance is registered on the machine where you licensed the agent, you are ready to set up your Data Manager - Databases jobs to monitor its storage. If an instance is not registered on the machine where you licensed the agent, see "Registering instances on machines that contain licensed agents."

Registering instances on machines that contain licensed agents

This topic describes how to register instances on machines that contain licensed agents.

1. Click the **RDBMS Logins** tab on the IBM Tivoli Storage Productivity Center for Data - Databases License Editor window.

2. The RDBMS Logins window is displayed. Use this window to edit, add, and delete the instances within your organization that you want to monitor. The field and button descriptions for the RDBMS Logins tab are:

Edit Highlight a row and click this button to edit the login information for the corresponding instance. The RDBMS Login Editor window is displayed.

Add New

Click this button to add login information for a new instance. The RDBMS Login Editor window is displayed.

Delete Highlight a row for an instance and click this button to delete that instance from Data Manager - Databases. Once you delete an instance, all the previously gathered statistics for that instance are automatically deleted from the database repository and that instance is no longer available for selection in the reporting section of the navigation tree.

3. Click **Add New** to add a new instance. The RDBMS Login Editor window is displayed. Use this window to enter information about the instance that you want Data Manager - Databases to monitor.

Note: In an Oracle Parallel Server (OPS) environment, you only need to register one of the instances within that environment. OPS is a resource-sharing system that increases availability and performance by partitioning the workload across multiple servers of a cluster (nodes). Databases installed on clustered servers or clustered database servers (for example, Oracle Real Application Cluster (RAC) environment) are not supported for monitoring.

4. Select the name of the machine where the instance is running from the **Host Name** list box.

Note: You can only register instances on machines that contain licensed agents.

5. Enter the following information in the next field:

For Oracle

Enter the Oracle SID and host for the instance.

For Microsoft SQL/Server

Enter the name of the instance you want to register in the **Instance** field.

Note: When you use the Storage Resource agent to monitor the RDBMS, the JDBC field is not shown. The Storage Resource agent does not require this field.

For Sybase

Enter the server name in the **Server** field.

For IBM UDB

Enter the name of the instance you want to register in the **Instance** field.

6. Enter a user ID that has the appropriate database privileges within the instance in the **User** field. See "License Keys" on page 268.

For Oracle

The following privileges are required for this user ID:

- CREATE SESSION
- SELECT ANY DICTIONARY
- ANALYZE ANY

Note: For Oracle 9i or 10g, specify ANALYZE ANY DICTIONARY

For Microsoft SQL Server

The login ID that Data Manager uses to log in to Microsoft SQL Server instances that you want to probe must have "permit" access.

7. Enter a password for the user ID in the **Password** field.
8. Enter the port on which the instance is listening in the **Port** field.

For Oracle

The default port is 1521.

For Microsoft SQL/Server

The default port is 1433. You must also provide the fully qualified path to the JDBC driver in the **JDBC Driver** field.

Note: When you use the Storage Resource agent to monitor the RDBMS, the JDBC field is not shown. The Storage Resource agent does not require this field.

For Sybase

The default port is 5000.

For IBM UDB

When monitoring multiple UDB instances within your environment, you must ensure that the port numbers you choose are open (unique for each instance) for JDBC and Java connections to those instances. To open up a port, run the following DB2 command on the machine where the instance is located:

```
db2jstrt
```

The default port number is 6789. You can change this default by indicating the port number you want to use when running the **db2jstrt** command. For example:

```
db2jstrt 6790
```

If you enter an incorrect port number, an error occurs. To verify the port number for an instance, do the following steps:

UNIX or Linux

View the `etc/services` file to confirm the correct port number.

AIX Run the following command:

```
ps -ef | grep db2jd
```

The output from this command indicates the port on which the instance is listening. Use this port number when running the **db2jstrt** command.

Windows

Use the information in the IBM DB2 configuration tools to confirm the port number.

9. Click **File > Save** to save the instance configuration information. You can now run a probe job against the registered instances. You must run a probe job against an instance before you can select any databases or table spaces against which you want to run a scan job.

Configuring Microsoft SQL Server 2005 or 2008:

Before you can monitor a Microsoft SQL Server 2005 or Microsoft SQL Server 2008 database, you must make some configuration changes to the Microsoft SQL Server.

Note: Before registering an instance of Microsoft SQL Server 2005 or 2008 to be monitored by a Storage Resource agent, please make sure that the directory containing the sqlcmd utility (sqlcmd.exe) is in the system PATH environment variable on the Microsoft SQL Server system. The default location of sqlcmd is C:\Program Files\Microsoft SQL Server\<version>\Tools\Binn.

Before monitoring the Microsoft SQL Server database, follow these steps:

1. Install the Microsoft SQL Server and provide the required information about the installation panels. See the Microsoft SQL Server 2005 Installation and Configuration Guide for detailed information. For the installation and configuration guide, see <http://msdn.microsoft.com/en-us/library/ms143516.aspx>.
2. For Microsoft SQL Server 2005 and 2008, make sure that the Microsoft SQL Server is using Mixed Mode authentication.
3. After the installation, go to the SQL Server Configuration Manager and make sure that you set the Dynamic TCP/IP port to the default port 1433. (See the Microsoft SQL Server 2005 Installation and Configuration Guide for detailed information.) To, configure the Microsoft SQL Server, follow these steps:
 - a. Open the SQL Server Configuration Manager.
 - b. Go to SQL Server 2005/2008 Network Configuration.
 - c. Select **Protocols for MSSQLINST** (name of the instance).
 - d. Right-click **TCP/IP**.
 - e. Select Enable: **Yes**.
 - f. Go to the **IP Addresses / IP All** and add **TCP Dynamic Ports : 1433** (default port).
4. Launch the Microsoft SQL Server setup.exe.
5. Go to **Installation/Search** for product updates. You are redirected to the Microsoft update Web site that scans the computer for the components that need to be updated.
6. When finished, select **Express Install** to install the components found.
7. Install the Microsoft CoreXML Services (MSXML) 6.0 SP 2. Reboot the system.
8. **For Microsoft SQL Server 2008 only:** Install the Microsoft SQL Server 2008 Management Objects from the following Microsoft SQL Server 2008 Feature Pack link: <http://www.microsoft.com/downloads/>. Search for **Management Objects**.
9. **For Microsoft SQL Server 2008 only:** Install the Microsoft SQL Server 2005 Backward Compatibility Components from the link specified in step 8. Search for **2005 Backward Compatibility**.
10. Install the Storage Resource agent and JDBC driver on the Microsoft SQL Server system.
11. Open the RDBMS Login Editor. Expand **Administrative Services > Configuration > License Keys**. In the content pane, click **RDBMS Logins** tab.
12. Click **Add New**. The RDBMS Login Editor opens. Enter the following information:

Database

Microsoft SQL/Server.

Instance name

Name of the instance (mssqlinst).

User User ID to logon to the Microsoft SQL Server.

Password

Password for the user ID.

Port 1433

JDBC Driver

Enter the JDBC driver. For Microsoft SQL Server 2005 and 2008, open the Microsoft SQL Server 2005/2008 JDBC Driver package and install the driver to the wanted location. After the package is unpacked, in the RDBMS logon panel, enter the absolute path to the JDBC Driver (suggested version 1.2):

`%MS_SQL_DRIVER_HOME%/sqljdbc_1.2/enu/sqljdbc.jar`

13. Click **Save**.

Unassigning Data Manager - Databases license

This topic describes how to unassign a Data Manager - Databases license.

To unassign a license, complete the following steps:

1. Expand **Administrative Services > Configuration > License Keys**.
2. The License Editor window is displayed.
3. Click the icon for the **IBM Tivoli Storage Productivity Center for Data - Databases** row. The IBM Tivoli Storage Productivity Center for Data - Databases License Editor window is displayed. Use the **Licensed** column in this window to view the agents to which licenses are currently assigned.
4. Clear the **Licensed** check box next to a machine to remove the license for the agent on that machine.
5. Click **File > Save**. When you remove the license for an agent, the following actions occur:
 - All the data gathered by that agent for the instance it monitors is removed from the database repository.
 - You can no longer run monitoring, alerting, or policy management jobs against the instance on the machine where the agent was located.
 - The number of unused licenses increases by one.

Alert disposition

You can configure IBM Tivoli Storage Productivity Center so that SNMP traps and Tivoli Enterprise Console events can be sent to other consoles.

Configuring Tivoli Storage Productivity Center to send SNMP traps and IBM Tivoli Enterprise Console events

This topic provides information about how to configure Tivoli Storage Productivity Center to send SNMP traps, Tivoli Enterprise Console events, or e-mail notifications. These actions are taken when an event specified in an alert is detected. From this node, you can also specify the number of days after which records in the alert log are deleted.

To enable SNMP traps and Tivoli Enterprise Console event alert notifications, consider the following information:

SNMP traps

System administrators must set up their SNMP trap ringer with the provided MIB files in order to receive SNMP traps from Tivoli Storage Productivity Center.

TEC events

Send an alert to the Tivoli Enterprise Console. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events. The event definitions are specified in the `tivoliSRM.baroc` and `fabric.baroc` files. It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

The `tivoliSRM.baroc` and `fabric.baroc` files must be loaded into the current active rule base of the Tivoli Enterprise Console server. This action allows Tivoli Enterprise Console events sent by Tivoli Storage Productivity Center to display on the console.

Table 39. Location of MIB and Baroc files

Component	MIB file location	Baroc file location
Data Manager and Disk Manager	<ul style="list-style-type: none">Installation media, disk1, part 1Data server: <code>data\snmp\tivoliSRM.mib</code>	<ul style="list-style-type: none">Installation media, disk1, part 1Data server: <code>data\tec\tivoliSRM.baroc</code>
Fabric Manager	<ul style="list-style-type: none">Installation media, disk1, part 1Device server: <code>device\snmp\fabric.mib</code>	<ul style="list-style-type: none">Installation media, disk1, part 1Device server: <code>device\tec\fabric.baroc</code>
Tivoli Storage Productivity Center for Replication	<ul style="list-style-type: none">Installation media, disk1, part 2<code>tpcr\TPCRM\CSM-Client\etc\ibm-TPC-Replication.mib</code>	<ul style="list-style-type: none">Installation media, disk1, part 2<code>tpcr\TPCRM\CSM-Client\etc\ibm-TPC-Replication.baroc</code>
Note: <ul style="list-style-type: none">The Tape Manager displays only indications received from the tape CIM agent in the tape alert log. Because the tape alerts cannot be configured, there is no forwarding of SNMP traps and there are no MIB or baroc files for the Tape Manager.Some alerts are generated by Tivoli Storage Productivity Center about Tivoli Storage Productivity Center for Replication. Examples of these alerts are communication failures between Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication. Other alerts generated are Tivoli Storage Productivity Center for Replication configuration changes. These alerts can be forwarded by Tivoli Storage Productivity Center as SNMP traps or through e-mail. These alerts are generated directly by Tivoli Storage Productivity Center and not by SNMP traps from Tivoli Storage Productivity Center for Replication. If Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication are running on the same system, only one program can listen on the SNMP default port. All traps received on the default port are forwarded to the other server process using a different port number.		

Port 162

Tivoli Storage Productivity Center uses port 162 to listen for SNMP traps. This port is the default port. To receive SNMP traps from switches, you must configure the switch to send SNMP traps to the Device server's IP address and port.

Note: Port 162 is also used by the Windows SNMP Trap service. If the Windows SNMP Trap service is running, stop the service or change the default port that Tivoli Storage Productivity Center uses to listen for SNMP traps. If you need to change the Device server's default port number, use the `setdscfg` command. The attribute to set is `SNMPTrapPort`.

To configure Tivoli Storage Productivity Center to send traps and Tivoli Enterprise Console events, follow these steps:

1. Open the Tivoli Storage Productivity Center GUI.
2. Go to **Administrative Services > Configuration > Alert Disposition**.
3. On the content pane, you can enter information for two SNMP destinations and one Tivoli Enterprise Console event destination. For the SNMP destination, enter the community name, host, and port information. For the Tivoli Enterprise Console event destination, enter the server and port information. You can also enter information for e-mail information. Click **File -> Save**.

Log File and Cached Batch Report Retention

You can set the retention period for log files.

You can specify the number of runs and the length of time to maintain the log files generated by IBM Tivoli Storage Productivity Center. We recommend that you set the value for **Maximum number of runs to keep of each schedule** to be less than or equal to the value in the **Maximum number of days' worth of log files to keep (regardless of schedule)** field. This setting helps avoid the situation where entries for schedule runs are displayed in the user interface, but the corresponding log files are deleted.

Note: By default, job log files for schedules are retained for 90 days, but you can change that value on the Log File and Cached Batch Report Retention panel. However, for CLI and Event Driven Jobs schedules, job log files are automatically retained for 90 days.

To set the retention period, complete the following steps:

1. Go to **Administrative Services > Configuration > Log file and Cached Batch Report Retention**.
2. In the **Maximum number of runs to keep of each schedule** field, enter the maximum number of runs for a schedule that you want to display in the user interface. The default is five runs per schedule. For example: if you enter **10** in this field and then run a probe schedule 11 times, job log and job information entries for the first run are deleted from the database repository. Additionally, only the 10 most recent runs of the schedule are displayed in the Job Management panel.
3. In the **Maximum number of days ' worth of log files to keep (regardless of schedule)** field, enter the maximum number of days to keep a log file for a run. A log file is generated for each job in a run and is stored in the appropriate log directory. Every log file that is older than the number of days you specify here is deleted. The default is 90 days.
4. In the **Cached Batch Report Retention** field, enter the maximum number of days to keep a batch report. The default is 90 days.

Note: When the batch report is created on the Tivoli Storage Productivity Center server and copied to the Storage Resource agent, the copying of the report to the agent might fail. If a failure occurs, no attempt is made to send the report again. Instead, the report remains cached on the Tivoli Storage

Productivity Center server. You can configure the time the report remains cached on the Tivoli Storage Productivity Center server.

Quota and Constraint e-mail Address Rules

You can specify rules for generating e-mail addresses of users who break the quota and constraint e-mail address rules based on their user ID, given name, or family name as they are registered within the operating system.

The user names are obtained as follows:

- For Windows: Full name field, from LDAP.
- For UNIX or Linux: User description from the Password file.

To set the rules, complete the following steps:

1. Expand **Administrative Services > Configuration > Quota and Constraint e-mail Address Rules**.
2. Click **Add After** or **Add Before** to include elements within an e-mail address template. An e-mail address template describes how to build the user ID. The user ID is notified in the event of a quota violation. This user ID represents the actual user who breaks the quota rules.
3. Select **USERNAME**, **FIRSTNAME**, **LASTNAME**, **Text**, or **Substring** from the pop-up menu to include as an element in the e-mail address rule:
 - **USERNAME**: the login ID of the user who breaks the quota or constraint rules.
 - **FIRSTNAME**: the given name of the user who breaks the quota or constraint rules.
 - **LASTNAME**: the family name of the user who breaks the quota or constraint rules.
 - **Text**: free form text that you want to appear within the e-mail address.
 - **Substring**: an element in the e-mail address that is a substring of **USERNAME**, **LASTNAME**, or **FIRSTNAME**. For example, `LASTNAME + SUBSTRING(USERNAME, 0, 3)`

includes the first three characters of the **USERNAME**. If the family name is Smith and **USERNAME** (as defined by file owner information) is 9A0723, then this substring example is equal to "SMITH9A0".

To define a substring, complete the following steps:

- a. Select **Substring** after clicking **Add After** or **Add Before**.
- b. Select the substring variable: **USERNAME**, **LASTNAME**, or **FIRSTNAME**.
- c. Highlight the range for the substring.
- d. Click **OK**. That substring appears within the e-mail address template.
- e. Click **File > Save** to save the e-mail address rule.

Scan/Probe Agent Administration

You can assign Storage Resource agents to run scan and probe jobs.

Assign the Storage Resource agents to perform scans against the following objects:

- File systems within NAS filers
- IBM Tivoli Storage SAN File Systems

The window associated with this node provides a complete listing of the NAS filers and SAN File System discovered by Data Manager.

Manual NAS Server Entry

Use the Manual NAS Server Entry node to enter, and view information about Network Attached Storage (NAS) servers.

In the manual NAS Server Entry node, you can complete the following tasks:

- Enter information about Network Attached Storage (NAS) servers that you want to monitor in your environment. After you enter information about the NAS servers, you can assign agents to the servers in the Scan/Probe Agent window.
- View a list of NAS filers. The NAS filers that were registered using the Data Manager are shown in the list.
- Delete NAS filers. The NAS filers that were registered using the Data Manager are deleted.

You can set up individual NAS servers for monitoring by Data Manager using this window, or you can use a discovery method to automatically add multiple servers simultaneously.

When you want to add multiple NAS Filers for monitoring, use the discovery method. When you want to add individual NAS Filers for monitoring, use the Manual NAS Server Entry window.

For more information about configuring NAS, see <http://www.redbooks.ibm.com/>. Search for **sg247490**.

Manually adding a NAS filer or gateway

To manually add a NAS filer or gateway:

1. Expand **Administrative Services > Configuration > Manual NAS Server Entry**.
2. Click the **Add NAS Server** button. The Add NAS Server window is displayed.
3. Enter the following information:

Network name

Enter the network name of the NAS server you want to add.

When manually adding a NAS Gateway or Filer that is monitored by a Storage Resource agent on a UNIX system, you must add the NAS using the same name that was used when mounting its file systems on that UNIX system. You can mount file systems using the short name, fully qualified name, or IP address of a NAS.

For example, if the file systems from a NAS Gateway were mounted to a UNIX computer (where the Storage Resource agent is installed) using the short name of the NAS Gateway, you must add the NAS Gateway in Tivoli Storage Productivity Center using the short name of the NAS Gateway. If the file systems from a NAS Gateway are mounted using an IP address, you must add the NAS Gateway in Tivoli Storage Productivity Center using the IP address as its name.

Consider the following example:

- a. The NAS filer named "oxide" was mounted on the UNIX system where a Storage Resource agent is located. The following commands used the short name and IP address to identify the NAS filer during a file system mount:

- oxide:/vol/john1 81920 22108 74% 123 5% /n3700_john1
 - 192.168.1.65:/vol/vol0 47566060 28715228 40% 1016085 52% /n3700_vol0
- b. To add this NAS filer to Tivoli Storage Productivity Center, enter the following values in the **Network Name** field:
- oxide
 - 192.168.1.65

If file systems from a NAS Gateway or Filer is mounted on a UNIX system in different ways, you must add that NAS to Tivoli Storage Productivity Center with the names used in both methods. For example, if one file system is mounted using the fully qualified name of a NAS Gateway and the other file system is mounted using the IP address, you must add that NAS Gateway to Tivoli Storage Productivity Center twice: once with the fully qualified name and once with the IP address. Only one row is displayed for this NAS on Manual NAS Server Entry panel, but both file systems are listed on the Scan/Probe Agent Administration panel.

Data Manager Agent OS Type

Select the operating system of the computer that contains the agent that gathers information about the NAS filer.

Accessible from

Select the agent that you want to use to "discover" the NAS filer. This list box displays agents that are:

- Running under the operating system selected in the **Data Manager Agent OS Type** field.
- Found on Windows or UNIX operating systems that are accessible to the NAS filers (Data Manager's agents are not found on the NAS filers themselves):
 - Windows: agents are found on Windows operating systems within the same domain as the NAS filers.
 - UNIX: agents are found on UNIX or Linux operating systems that have NFS imports for the file systems within the NAS filers.

SNMP Community

Enter the name of the SNMP communities that Data Manager uses when communicating with systems in your environment. If you do not enter the name of an SNMP community, the default community **public** is used. Data Manager uses the SNMP protocol to contact and identify NAS filers. This field is optional.

Login ID

(Windows operating systems only.) Enter the login ID that Data Manager uses when logging in to the NAS filer.

Password

(Windows operating systems only.) Enter the password that Data Manager uses when logging in to the NAS filer.

Add as Other NAS

Select this check box to add a NAS server as **Other NAS** filer. With this option, you can monitor and report on file system information about the NAS filer or gateway through Windows CIFS or UNIX NFS shares accessible to the scan or probe job for the agent. No controllers, disks, and logical volumes information are collected or reported.

NAS Server Vendor Name

Enter the vendor name (or manufacturer) hosting the file system of the NAS server. The default vendor is Network Appliance.

4. Click **OK** to have the Data Manager verify the filer for which you entered information. During this verification, Data Manager will:
 - Log in to the NAS filer.
 - Gather information about the file systems visible on those filers to the agent. For the UNIX or Linux operating system, it gathers information about the file systems that it can actually see (for example, file systems that are mounted to the UNIX operating systems). By default, file systems are discovered at the root. For the Windows operating system, it finds all the NAS filers that are visible through CIFS.
 - Determine which file systems are visible to which agents.
 - Enter the NAS server information into the repository.
5. Expand **Administrative Services > Configuration > Scan/Probe Agent Administration**. This window allows you to assign agents to each file system of the NAS server.

Note: At any time, you can change the login ID and password for a NAS filer on the **Administrative Services > Configuration > License Keys > Filer Logins** window.

Deleting a manually added NAS filer

To delete a NAS filer whose information was manually entered into Data Manager, follow these steps:

1. Expand **Administrative Services > Configuration > Manual NAS Server Entry**.
2. Highlight a row that represents the device you want to delete.
3. Click **Delete**. Remember that when you delete a device from this window, all information about that device is removed from the repository.

Editing Data Manager configuration files

Tivoli Storage Productivity Center provides you with the ability to edit Data Manager configuration files to further customize the settings for a component according to the standards at your site.

Data Manager has configuration files for customizing the operation of the server and agent components within your environment. These files are located in the *TPC_installation_directory/config/* directory, where *TPC_installation_directory* represents the directory where you installed the product.

The configuration file for the Tivoli Storage Productivity Center server is located in the following default installation directory:

Windows operating system

C:\Program files\IBM\TPC\Data\config\

UNIX or Linux operating system

/opt/IBM/TPC/Data/config

The agent configuration file for the Storage Resource agent is located in the following default installation directory:

Windows operating system

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

UNIX or Linux operating system

/opt/IBM/TPC/agent/config/

When you change the configuration files for the server component on the UNIX or Linux operating system, you must stop and start the server before those changes take effect.

Edit the `agent.config` file to configure the Data Manager agents in your environment. This file is located in the agent installation directory on every computer where an agent is installed. For information about the `agent.config` file, go to the information center. Search for *agent.config file*.

Edit the `nas.config` file to configure the Data Manager NAS feature for your environment.

The `nas.config` file contains the following information:

- On each line not beginning with #, the first blank-delimited field must contain the SNMP Enterprise code of a NAS filer that the agent discovers, probes, or scans.
- The second field contains identifying information about the filer.

Any remote host that cannot be reached by SNMP or whose enterprise code does not match one of these values is ignored.

For information about the server files to configure, go to the information center. Search for *server.config file*, *scheduler.config file*, and *TPCD.config file*.

For information about the `agent.config` file, go to the information center. Search for *agent.config file*.

Editing the NAS configuration file

This topic provides information about editing the `nas.config` file for the Data Manager NAS feature.

Edit the `nas.config` file to configure the Data Manager NAS feature for your environment.

The `nas.config` file contains the following information:

- On each line not beginning with #, the first blank-delimited field must contain the SNMP Enterprise code of a NAS filer that the agent discovers, probes, or scans.
- The second field contains identifying information about the filer.

Any remote host that cannot be reached by SNMP or whose enterprise code does not match one of these values are ignored.

History Aggregator

You can configure reports for data aggregation.

The History Aggregator in IBM Tivoli Storage Productivity Center defines and runs jobs to sum data in the database repository for historical reporting purposes. For example, you can view the sum of usage across multiple storage resources, by filesystem. Trending enables you to see patterns of your historical data across your entire network.

You have the option of turning aggregation off, although this action is not recommended. To turn off aggregation, access the History Aggregator window, clear the **Enabled** check box, and select **File > Save**.

To configure reports for data aggregation, complete the following steps:

1. Expand **Administrative Services > Configuration > History Aggregator**.
2. The Edit History Aggregator panel is displayed in the content pane. You can specify the following information:
 - How often to run the job
 - How to handle time zones
 - Triggering conditions
 - Triggered actions

Storage Resource agent Deployments

You can use this page to manage your Storage Resource agent deployments.

Deploying Storage Resource agents

Deploy Storage Resource agents through the user interface rather than a separate installation wizard. You can have only one agent per host that points to the same server.

You can enter host names from a Microsoft directory, enter host names manually, or import a host list file. You can define a schedule to add the Storage Resource agents at a time that is convenient for you. You can also specify the type of alerts to generate if a schedule fails. After you enter information for the Storage Resource agent, you can optionally validate the connection to the Data server. This step helps to eliminate most of the possible failure cases after you submit the deployment schedule.

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

To deploy Storage Resource agents, follow these steps:

1. See “Deployment guidelines and limitations for Storage Resource agents” on page 282 for a list of considerations before deploying Storage Resource agents. If you are deploying Storage Resource agents on UNIX or Linux operating systems, you must use root as the user ID.
2. Open the Tivoli Storage Productivity Center user interface.
3. In the navigation tree pane, expand **Administrative Services > Configuration > Storage Resource agent Deployments**. Right-click **Storage Resource agent Deployments** and click **Create Storage Resource agent Deployments**.
4. The Create Storage Resource agent Deployment window opens. Enter information for the Computers tab.
5. Enter information for the When to run tab.
6. Enter information for the Alert tab.
7. Click **File > Save**.

Related reference:

“Deployment guidelines and limitations for Storage Resource agents”

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

Deployment guidelines and limitations for Storage Resource agents

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

Use the following information when you deploy Storage Resource agents and migrate existing Data agents:

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

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

Platforms that support Storage Resource agents for the same

For a list of platforms on which you can deploy Storage Resource agents, go to the Tivoli Storage Productivity Center support site at [http://www.ibm.com/support/ctgibm.nsf/00000000000000000000000000000000](#), click the **Documentation** link, and enter **Platform Support: Agents, Servers and GUI** in the **Search** box. Click the link to the document for the appropriate release.

Product functions that are not available for storage devices monitored by Storage Resource agents

Before you migrate an existing Data agent to a Storage Resource agent or deploy a new Storage Resource agent, make sure that the product functions you want to use on the monitored devices are available for those agents. The following functions are not available for storage devices that are monitored by Storage Resource agents:

- Certain relational database monitoring. For list of relational databases that can be monitored by Storage Resource agents, go to the Tivoli Storage Productivity Center support site at [http://www.ibm.com/support/ctgibm.nsf/00000000000000000000000000000000](#), click the **Documentation** link, and enter **Platform Support: Agents, Servers and GUI** in the **Search** box. Click the link to the document for the appropriate release.
- The changing of zone configurations and the reporting of HBA, fabric topology, or zoning information for fabrics connected to hosts running Linux on IBM System z hardware. These limitations also apply to Storage Resource agents on all guest operating systems for VMware configurations.

Required authorities for deploying and running Storage Resource agents

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

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

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

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

Orphan zones

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

Firewalls and Storage Resource agent deployments

Before you can deploy a Storage Resource agent on a computer, you must turn off the firewall on that computer. If you do not turn off the firewall, the deployment fails. To turn off the firewall on a Windows 2008 computer, complete the following steps:

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

Multiple computers in a deployment

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

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

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

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

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

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

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

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

Daemon and non-daemon services

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

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

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

Tivoli Storage Productivity Center requires the correct authentication information (user ID, password, port, certificate location, or passphrase) for monitored computers each time it communicates with Storage Resource agents on those computers. If the authentication information changes for a host computer on which a Storage Resource agent is deployed, the authentication information for that agent must be updated using the **Change Authentication** or **Update Storage Resource agent List** functions on the panel located under the **Administrative Services > Data Sources** node in the navigation tree.

Storage Resource agents on the same computer

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

Time zones for computers monitored by Storage Resource agents

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

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

UNIX operating systems using Remote Shell protocol (RSH)

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

Deployments on Windows operating systems - NetBIOS setting

If you want to install a Storage Resource agent on Windows targets, the **Enable NetBIOS over TCP/IP** option must be selected in the Control Panel settings for the computer's network connections properties. To set this option, click **Start > Settings > Control Panel > Network and Dial-Up**

Connections > <some_connection> > Properties > Internet Protocol (TCP/IP) > Advanced > WINS > Enable NetBIOS over TCP/IP.

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

To determine if security policies are blocking the connection ports, click **Start > Settings > Control Panel > Administrative Tools.**

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

Locally stored policies

For locally stored policies, click **Start > Settings > Control Panel > Administrative Tools > Local Security Policy > IP Security Policies on Local Computer.**

Policies stored in Active Directory

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

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

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

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

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

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

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

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

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

Note: This task contains steps that tell you how to modify the registry. Serious problems might occur if you modify the registry incorrectly. Therefore, make sure that you follow these steps carefully. For added

protection, back up the registry before you modify it. Then you can restore the registry if problems occur. For information about how to back up and restore the registry, see .

To disable UAC remote restrictions, follow these steps:

1. Click **Start** > **Run**. Enter **regedit** and click **OK**.
2. Locate and click the following registry subkey:
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\
Policies\System
3. If the LocalAccountTokenFilterPolicy registry entry does not exist, follow these steps:
 - a. On the **Edit** menu, click **New** > **DWORD Value**.
 - b. Enter **LocalAccountTokenFilterPolicy** as the name for the DWORD value in the content pane. Click **Enter**.
 - c. Right-click **LocalAccountTokenFilterPolicy**, then click **Modify**.
 - d. In the Edit DWORD Value window, enter **1**, then click **OK**.
This value can be 0 or 1:
 - 0 This value builds a filtered token. This value is the default value. The administrator credentials are removed.
 - 1 This value builds an elevated token.
 - e. Exit the registry editor.

Additional information about Storage Resource agents

See “Planning for Storage Resource agents” on page 43 for more information about Storage Resource agents.

Creating a certificate for SSH protocol

Before installing the Storage Resource agents using the SSH protocol, you can optionally create a certificate.

Note: The Storage Resource agent only supports either DES-EDE3-CBC encryption or no encryption for the private key used in SSH protocol communication between the server and agent. The default encryption used in the **ssh-keygen** command on UNIX is always DES-EDE3-CBC. However, with Windows Cygwin, **ssh-keygen** generates a key with AES-128-CBC encryption if a passphrase is specified. If there is no passphrase, the private key is generated without encryption. For more information about encryption, see <http://www.openssl.org/docs/apps/enc.html>.

Creating a certificate for SSH protocol (non-Windows)

Note: The Storage Resource agent only supports either DES-EDE3-CBC encryption or no encryption for the private key used in SSH protocol communication between the server and agent. The default encryption used in the **ssh-keygen** command on UNIX is always DES-EDE3-CBC but with Windows Cygwin, it is using AES-128-CBC encryption if a passphrase is specified. If there is no passphrase, the private key is generated without encryption.

To create a certificate for SSH protocol, complete the following steps:

To create a certificate for SSH protocol, complete the following steps:

1. Telnet to the remote machine using the root user ID.
2. To create an SSH certificate on AIX, you must first install the following packages (if not already installed):

```
openssl.base.openssh.base.client
openssh.base.server
```

3. Go to the directory where you want to create the certificate:

```
cd to ~/.ssh
```

4. Enter **ssh-keygen**. Accept the default names (for example, **id_rsa**).

5. Enter the passphrase.

6. Two files are created:

id_rsa The private key.

id_rsa.pub

The public key.

7. Create an **authorized_key** file in the same location as **id_rsa.pub** by entering the following command:

```
cat id_rsa.pub >> authorized_keys
```

8. Copy the **id_rsa** (private key) to your server machine. For example, to copy the **id_rsa** file to **:\keys\id_rsa** on the IBM Tivoli Storage Productivity Center server (user responses are in boldface type):

```
# ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (//.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been save in //.ssh/id_rsa.
Your public key has been save in //.ssh/id_rsa.pub.
The key fingerprint is:
xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx root@server
# cat id_rsa >> authorized_keys
# ls -l
total 24
-rw-r--r-- 1 root system 1743 Oct 15 09:40 authorized_keys
-rw---- 1 root system 1743 Oct 15 09:39 id_rsa
-rw-r--r-- 1 root system 399 Oct 15 09:39 id_rsa.pub
#
```

Note: You must copy the file in binary mode.

9. To connect to the remote system using the private key, the following information needs to be entered on the Remote Agent Machines panel when you install the Storage Resource agent:

- User
- Certificate Location (c:\keys\id_rsa)
- Passphrase

Setting up an SSH daemon on Windows

On Windows Server 2003, Windows XP, Windows Server 2008, or Windows Vista, you must run the **ssh-host-config** command.

Note: Cygwin is not a prerequisite for the Storage Resource agent on Windows. To use the SSH protocol on Windows, an SSH software program must be used because Windows does not come with an SSH service. Cygwin is a free software program providing SSH access to a Windows server. Cygwin can be used if you want to run the Storage Resource agent using the SSH protocol.

You must be in a Cygwin window or be an X term user to create the **sshd** service. In most cases, you click the **cygwin.bat** to start the Bash shell.

Run the following steps:

1. Install Cygwin.
2. Set up your **sshd** service in Cygwin.
3. Create the certificate (this step is optional).

Installing Cygwin

To install Cygwin, go to <http://cygwin.com>. This page contains a link that displays help for the setup program and a link to download the setup program. Read the help before running the setup program. Then download the Cygwin program by clicking the **Install Cygwin now** link. Start the setup program on your computer by running the **setup.exe** program. Select the appropriate download option (**Install from Internet**, **Download from Internet**, or **Install from Local Directory**) as described in the help files.

If you are upgrading from an older version of Cygwin to a newer version, you need to remove the **sshd** service before installing the new version of Cygwin.

Accept the default installation options as they are presented to you (Root Directory, Install For, Default Text File Type, and so on). Select a download mirror that is geographically close to your location. Some sites require an FTP account before you can install Cygwin. You can either request an account or simply select another mirror.

During the installation process, a Select Packages list is displayed. Expand the plus sign (+) next to the Admin category and select **cygrunsrv** and the **Bin** check box. Expand the plus sign (+) next to the Net category and select **openssh**. Expand the plus sign (+) next to the Util category and select **diffutils**. Click **Next** to resume the setup program. The time required to download the packages depends on how busy the mirror is, and on the speed of your internet connection. With **openssh** and **cygrunsrv**, the downloaded files require approximately 70 MB of disk space. Allow 20 minutes to 30 minutes for the download and installation to complete.

Setting up your sshd service in Cygwin

Here is an example of the sequence of steps and responses. The responses to the prompts are in boldfaced type.

1. Run the **ssh-host-config** command.

Note: With Cygwin, you might experience permission problems when running the **ssh-host-config** command. If you have permission problems, run these commands:

```
chmod +r /etc/passwd
chmod +r /etc/group
chmod 777 /var
```



```

$ ssh-host-config
Generating /etc/ssh_host_key
Generating /etc/ssh_host_rsa_key
Generating /etc/ssh_host_dsa_key
Generating /etc/ssh_config file
Should privilege separation be used? (yes/no) no

Generating /etc/ssh_config file

Warning: The following functions require administrator privileges!

Do you want to install sshd as service? yes (If sshd is already
installed as a service, answer no.)

You appear to be running Windows 2003 Server or later. On 2003 and
later systems, it's not possible to use the LocalSystem account, if
sshd should allow passwordless logon (e. g. public key
authentication).

If you want to enable that functionality, it's required to create a
new account 'sshd_server' with special privileges, which is then used
to run the sshd service under.

Should this script create a new local account 'sshd_server' which
has the required privileges? (yes/no) yes

Please enter a password for new user 'sshd_server'. Please be sure
that this password matches the password rules given on your system.

Entering no password will exit the configuration.
PASSWORD=password (Specify a password for the sshd_server
account.)

User 'sshd_server' has been created with password 'password'.

If you change the password, please keep in mind to change the
password for the sshd service, too.

Also keep in mind that the user sshd_server needs read permissions
on all users' .ssh/authorized_keys file to allow public key
authentication for these users. (Re-)running ssh-user-config
for each user will set the required permissions correctly.

Which value should the environment variable CYGWIN have when sshd
starts? It's recommended to set at least "ntsec" to be able to
change user context without password.

Default is "ntsec".

CYGWIN=ntsec

The service has been installed under sshd_server account.

To start the service, call 'net start sshd' or 'cygrunsrv -S sshd'.

Note! If the service doesn't start because of a login failure

Host configuration finished. Have fun!

```

2. Start the **sshd** service:
 - a. Open a command prompt window.
 - b. Enter **net start sshd** or in a Bash prompt, enter **cygrunsrv -start sshd**.
 - c. Verify that the daemon is running.
 - d. Enter **ps -a**. Examine the output to see if `/usr/sbin/sshd` is contained in the list of running processes.

To stop the service from a Windows command prompt, enter **net stop sshd**. Alternatively, you can change to the C:\cygwin\bin directory (or open a Bash shell) and enter **cygrunsrv -stop sshd**.

3. When you have started the **sshd** service, test it by entering the following command from a Bash shell prompt:

```
ssh localhost -l user_ID
or
ssh host_name -l user_ID
```

If **localhost** does not work, use the short host name. If you receive a message indicating that the authenticity of localhost cannot be established, answer **Yes** to the question "Are you sure you want to continue connecting?" When prompted for your account password on **localhost**, enter the password you use when logging in to the computer.

4. Create the accounts that can log in to the computer:
 - a. Create the Windows accounts. Click **Start > Settings > Control Panel > User Accounts**. Make each user a member of the Administrators group. Perform this operation for each user you want to add before you create the corresponding Cygwin accounts.
 - b. Make a backup copy of the /etc/passwd file.
 - c. Create the Cygwin user accounts. Run the following command:

```
mkpasswd -l>/etc/passwd
```
 - d. Verify that a home directory has been created for each account that you have added. Change the ownership of the home directory to its owner (run the **chown** command). If a home directory for the user does not exist, create one. For example, enter the following command:

```
mkdir home/account_name;chown account_name/home/account_name
```
 - e. When you add users, you need to stop and start **sshd** before that account is recognized because **sshd** only reads the file /etc/passwd when the service starts.
 - f. If you need to create groups of accounts, create the Windows groups first, then create the Cygwin groups. After creating the Windows groups, run the following command:

```
mkgroup -l>/etc/passwd
```

When you add groups, you also need to stop and start **sshd** before the new group is recognized.

5. Set the TEMP environment variable. For information about setting the environment variable, see <http://www.cygwin.com/cygwin-ug-net/setup-env.html>.

Here is an example of setting the environment variable:

- a. Right-click **My Computer > Properties > Advanced > Environment Variables**.
- b. Under **System variables**, find out the value of TEMP. For example, "C:\WINNT\TEMP"
- c. Set the TEMP environment variable to point to the Cygwin format of TEMP in the ~/.bashrc file. For example run the following command:

```
export TEMP=/cygdrive/c/WINNT/temp
```

Uncomment and modify this line in the `~/.bashrc` file from the default:

```
# export TEMP=/tmp
to
export TEMP=/cygdrive/c/WINNT/temp
```

The Cygwin `sshd` service must be added as a service that starts automatically. To verify this step, click **Start > Settings > Control Panel > Administrative Tools > Services**. Look for **CYGIN sshd** in the name list. Verify that it is started and configured to start automatically.

Creating the certificate

To create a certificate for SSH protocol, complete the following steps:

1. Run this command:

```
cd to ~/.ssh
```

2. Generate the public and private keys with a passphrase. From the Bash shell prompt, here is an example of the input and output (user responses are in boldface type):

```
Administrator:
~/.ssh
$ openssl genrsa -des3 -out key 1024

Response:

Generating RSA private key, 1024 bit long modulus
...
...
e is 65537 (0x10001)
Enter pass phrase for key: (enter pass phrase for key)
Verifying - Enter pass phrase for key: (enter pass phrase for key again)

Administrator:
~/.ssh
$ chmod 600 ~/.ssh/key

~/.ssh
$ ssh-keygen -y -f ~/.ssh/key > key.pub

Response:

Enter passphrase: (enter passphrase)

Administrator:
~/.ssh
$ cat key.pub >> authorized_keys
```

3. Once the `key.pub` file has been added to the `authorized_keys` directory, copy the private key `key` to the IBM Tivoli Storage Productivity Center server.

Replacing a custom certificate for SSL protocol

IBM Tivoli Storage Productivity Center provides default SSL certificates for communication between the Data server and Storage Resource agent. You can replace your custom certificates with the default certificates.

Overview of replacing a custom certificate for SSL protocol

IBM Tivoli Storage Productivity Center uses SSL certificates for communication between the Data server and Storage Resource agent (daemon service). IBM Tivoli Storage Productivity Center provides default SSL certificates for this

communication. If you want to use your own certificates, you can replace the default certificates with your custom certificates.

Server certificate

The IBM Tivoli Storage Productivity Center Data server uses the `TPCDataServer.jks` and `server.pwd` files for communication with the Storage Resource agents. If you are using custom certificates, you need to replace these files.

There are two ways you can replace these certificates:

- Before installation of the Data server
- After installation of the Data server

Storage Resource agent certificate

The Storage Resource agent uses the certificate files `sra.pem` and `sra.pwd` for communication with the Data server. These two files are compressed into the `certs.zip` file on the server system for deployment purposes. If you are using custom certificates, you need to replace these files.

There are two ways you can replace the certificates:

- Before deployment of the agent
- After deployment of the agent

The general steps for replacing custom certificates are:

1. Generate the custom certificates.
2. Stop the Data server (and the Storage Resource agent, if the agent is already deployed) if applicable.
3. Replace the custom certificates on the Data server and Storage Resource agent or on the disk 1 or Storage Resource Agent installation image.
4. Start the Data server (and the Storage Resource agent, if the agent is already deployed) if applicable.

Note: When generating custom SSL certificates, the certificates have a start date, end date, and time when they are valid. These dates and times are related to the system where these custom certificates were generated (which is usually the server system). When installing a Storage Resource agent on a remote system, you must check the date and time on the Storage Resource agent system. If the server and agent systems are in the same time zone, they must have the same date and time. Otherwise, the time zone difference should be set.

For example, if the server system is 8:00 PM, the agent system should also be 8:00 PM. If the agent system is set at a different time (for example, 6:00 PM) at the time the SSL custom certificates are generated on the server system with a time of 8:00 PM, the deployment of the Storage Resource agent will fail.

How to generate custom certificates

The script file `createSRACerts.sh` (for Linux or UNIX) or `createSRACerts.bat` (for Windows) is located in the following directory:


`TPC_install_directory/data/sra/tools/certs`

`TPC_install_directory` is where the IBM Tivoli Storage Productivity Center servers are installed. The default directory is `/opt/IBM/TPC` for Linux or UNIX or `C:\Program Files\IBM\TPC` for Windows.

To generate custom certificates, follow these steps:

1. Create the custom certificates.

The createSRACerts script creates the custom certificates. The syntax is:



output_directory

Directory where the certificates are created. The default is sra_certs_out directory in the current working directory.

rootCAPassword

Root CA password (root common authority password). The default non-encrypted password is: s5umEvApR6cafruhustu.

server_key_password

Server key password. The default non-encrypted password is: drUtahaswefraf9uth.

server_store_password

Server store password. The default non-encrypted password is: wr4d5Xekaqafehet5u2a.

agent_password

Agent password. The default non-encrypted password is: jawUchezuthew6azEjef.

Note: You are prompted for the password when the script is run. Provide the value of the rootCAPassword on the command line (if specified on the command line). Otherwise, use the default values in the createSRACerts.sh or createSRACerts.bat script when prompted for the password (twice). The following example creates the SSL certificate using the default output directory and default passwords.

```
createSRACerts
```

The following examples create the SSL certificates using the directory /tmp on UNIX and C:\temp on Windows. These examples use the default passwords. The passwords are stored in the following files:

Windows

```
C:\temp\sra_certs_out
```

Here is an example:

```
createSRACerts C:\temp
```

UNIX or Linux

```
/tmp/sra_certs_out
```

Here is an example:

```
createSRACerts /tmp
```

The following examples create SSL certificates in the directory /tmp on UNIX and C:\temp on Windows. These examples use non-default passwords for root CA and server key. The passwords are stored in the following files:

Windows

```
C:\temp\sra_certs_out
```

Here is an example of creating a non-default password for the root CA password and server key password:

```
createSRACerts C:\temp newpasswordforrootCA newpasswordforserver
```

UNIX or Linux

```
/tmp/sra_certs_out
```

Here is an example of creating a non-default password for the root CA password and server key password:

```
createSRACerts /tmp newpasswordforrootCA newpasswordforserver
```

2. Regenerate the certificates again if you have a failure. Delete the files in the out directory before rerunning the createSRACerts script.
3. Stop the Data server and Storage Resource agent. If the Data server is installed and running, stop the Data server. Make sure that the JVM for the Data server is not running.

Stop the Storage Resource agent if the certificates for the agent are installed and running.

4. Replace the certificates. There are three scenarios:
 - Replacing the certificates after the server is installed.
 - Replacing the certificates before the server is installed.
 - Replacing the certificates after the agent is installed.

Replacing the certificates after the server is installed

The server certificates are in the following directory:

```
TPC_install_directory/data/sra/tools/certs/out
```

These are the server certificates:

```
TPCDataServer.jks  
server.pwd
```

For the server certificates, copy the server certificates to the following directory:

```
TPC_install_directory/data/sra/certs
```

The agent certificates are created in the certs.zip file in the following directory:

```
TPC_install_directory/data/sra/tools/certs/out/agent
```

Copy the certs.zip file into the following directory:

```
TPC_install_directory/data/sra/agent_operating_system
```

Extract the certs.zip file into the following directory:

```
TPC_install_directory/data/sra/server_operating_system
```

agent_operating_system is the operating system that the Storage Resource agent is running under. *server_operating_system* is the operating system on which the Data server is installed on.

Replacing the certificates before the server is installed

The server certificates are in the following directory:

```
disk1_or_SRA_image_install_directory/data/sra/tools/certs/out
```

These are the server certificates:

```
TPCDataServer.jks  
server.pwd
```

disk1_or_SRA_image_install_directory is the location of the disk1 or Storage Resource agent image.

For the server certificates, copy the server certificates into the following directory:

TPC_install_directory/data/sra/certs

The agent certificates are created in the *certs.zip* file in the following directory:

TPC_install_directory/data/sra/tools/certs/out/agent

Copy the *certs.zip* file into the disk1 or Storage Resource agent image directory:

*disk1_or_SRA_image_install_directory/data/sra/
agent_operating_system*

Disk1 or SRA is the location of the disk1 or Storage Resource agent image that was extracted for installation of the server and agents.

Extract the *certs.zip* file into the following directory:

disk1_or_SRA_image_install_directory/data/sra/server_operating_system

agent_operating_system is the operating system that the Storage Resource agent is running under. *server_operating_system* is the operating system on which the Data server is installed on.

Replacing the certificates after the agent is installed

The agent certificates are in the *certs.zip* file that gets created in the following directory:

TPC_install_directory/data/sra/tools/certs/out/agent

Extract the *certs.zip* file into the

storage_resource_agent_install_directory/certs directory.

storage_resource_agent_install_directory is the location where the Storage Resource agent has been installed.

5. Start the Data server and the Storage Resource agent. If the Data server has stopped for replacement of the certificates, then start the Data server after the replacement of the certificates.

If the Storage Resource agent has stopped for replacement of certificates, then start the Storage Resource agent after the replacement of the certificates.

Storage Resource Agent Upgrades

You can configure IBM Tivoli Storage Productivity Center for a Storage Resource agent upgrade.

When you apply maintenance to Tivoli Storage Productivity Center, you can upgrade the Storage Resource agent now or upgrade the agent at a later time. To ensure that all your agents are at the current release level and to manage your network load, schedule upgrades on a regular basis.

To configure Tivoli Storage Productivity Center for a Storage Resource agent upgrade, complete the following steps:

1. In the navigation tree, expand **Administrative Services > Configuration > Storage Resource Agent Upgrades**.
2. Right-click on **Storage Resource Agent Upgrades**. Click **Create Storage Resource Agent Upgrade**.

3. The Create Storage Resource Agent Upgrade page is displayed in the topic pane. You have the following tabs:

Computers

You can select Computer Groups or Computers.

When to Run

You can specify how often to run the job or how to handle time zones.

Options

You can specify to force an upgrade for an agent, maintenance upgrade for an agent, or language pack to apply to an agent.

Alerts You can specify the triggering conditions and triggering actions.

Enter your information in the appropriate fields.

4. Save and run the job.

Migrate Data/Fabric Agent

Define schedules to migrate existing Data agents or Inband Fabric agents to Storage Resource agents. The migration process installs a Storage Resource agent on a target host and then uninstalls the existing agent.

Tivoli Storage Productivity Center provides the following benefits when you migrate an existing Data agent or Inband Fabric agent to a Storage Resource agent:

- Storage Resource agents require less resources on a host computer than a Data agent or Inband Fabric agent that is based on Common Agent Services. Additionally, you can deploy Storage Resource agent as a daemon or non-daemon service.
- Reduced complexity when deploying agents. You can deploy Storage Resource agents directly from the Tivoli Storage Productivity Center user interface.
- Improved interface when deploying, upgrading, and administering agents. You can manage Storage Resource agents using the nodes in the **Administrative Services** section of the navigation tree. You can deploy and upgrade Storage Resource agents on schedules that you define.

See “Deployment guidelines and limitations for Storage Resource agents” on page 282 to view additional information about deploying and using Storage Resource agents.

To migrate a Data agent or Inband Fabric agent to a Storage Resource agent on a host, complete the following steps:

1. Expand **Administrative Services > Configuration** in the navigation tree.
2. Right-click **Migrate Data/Fabric Agent** and select **Create Data/Fabric Agent Migration** from the menu. The Create Data/Fabric Agent Migration panel is displayed.
3. Enter information about the schedule on the following pages:
 - **Computers:** Use this page to select the computers and groups of computers that contain the Data or Fabric agents that you want to migrate to Storage Resource agents.
 - **When to run:** Use this page to determine when the Data and Fabric agents in a schedule are migrated.
 - **Options:** Use this page to specify the runtime settings of a Storage Resource agent, the method for migrating an agent, and the login credentials for an agent.

- **Alert:** Use this page to define the actions that occur if a Data/Fabric Agent Migration schedule fails.
4. Click **File > Save** from the menu bar to save the migration schedule. The migration runs according to the date and time that you defined in the schedule definition. Use the **Administrative Services > Data Sources > Data/Storage Resource Agents** panel to view a list of the Storage Resource agents that are deployed in your environment.

Update Storage Subsystem Credentials

If you have a System Storage DS8000, XIV system, or SAN Volume Controller storage system, you must update your storage system credentials to use the Tivoli Storage Productivity Center native interfaces.

To enter the storage subsystem credentials, follow these steps:

1. Start the Tivoli Storage Productivity Center migration tool to enter storage subsystem credentials. You can start the migration tool in one of the following ways:
 - Use the migration tool from the Tivoli Storage Productivity Center DVD. Use the DVD or download the migration tool: `MigrateUserInfo.bat` (for Windows) or `MigrateUserInfo.sh` (for UNIX or Linux). Run the migration tool from the directory `UserMigrationTool`.
 - Run the Tivoli Storage Productivity Center upgrade installation program. When Tivoli Storage Productivity Center detects the storage system, use the panel that is displayed to run the migration tool. Run the migration tool.
 - After upgrading Tivoli Storage Productivity Center, open the GUI. In the navigation tree, expand **Administrative Services > Configuration > Update Storage Subsystem Credentials**. This node is visible only when Tivoli Storage Productivity Center detects that you have a System Storage DS8000, XIV system, or SAN Volume Controller storage system.
2. The Update Storage Subsystem Credentials panel is displayed.

The table lists the subsystems with credentials that can be updated automatically. Click the update button to update the selected subsystems.

Update Update All

Subsystem	Device Type	Name

The table lists the subsystems that can have their credentials manually updated. Click the update button to migrate the selected subsystem.

Update

Subsystem	Device Type	Name
xxx xxx xxx xxx	IBM DS8000	2107.1300241

Figure 13. Update Storage Subsystem Credentials panel

Top table

All System Storage DS8000 storage systems are displayed if Tivoli Storage Productivity Center can use the existing CIM agent user ID and password to connect to the HMC. If you want to use the existing credentials, click **Update** or **Update All** (if you have more than one System Storage DS8000). Tivoli Storage Productivity Center updates the credentials and removes the System Storage DS8000 from the table.

Bottom table

All XIV system and SAN Volume Controller storage systems are displayed in this table. All other System Storage DS8000 storage systems are displayed in this table if they are not listed in the other table. All XIV system and SAN Volume Controller storage systems must have the storage system credentials migrated manually.

If you want to manually update the credentials for the System Storage DS8000, select a storage system and click **Update**. The Update Subsystem Credentials window is displayed. Enter the following information:

Username

Enter the user ID for the System Storage DS8000 element manager.

Password

Enter the password for the user name.

HMC2 Address (Optional)

You can enter a second HMC address. This parameter is optional.

Click **Save**. The System Storage DS8000 is removed from the table after you update the credentials.

To update the credentials for the XIV system, select a storage system and click **Update**. The Update Subsystems Credentials window is displayed. Enter the following information:

Username

Enter the user ID for the XIV system.

Password

Enter the password for the user name.

Click **Save**. The XIV system is removed from the table after you update the credentials.

If you want to update the credentials for IBM SAN Volume Controller, select a storage system and click **Update**. The Update Subsystem Credentials window is displayed. Enter the following information:

Software Version

Enter the version of the SAN Volume Controller software that you want to configure.

IP address

Enter the IP address or host name of the SAN Volume Controller. You can enter an IPv4 or IPv6 address, for example:

- 9.47.97.128
- [2001:DB8::1234:0000:0000:5678:ABCD]

Select Key

Select one of the following actions for the Secure Shell (SSH) key:

- **Use existing uploaded key:** Select this option if you want to use a public key that has been uploaded using a method other than

through Tivoli Storage Productivity Center, such as the SAN Volume Controller web console. Enter the location of the private key file that corresponds to the public key in the **Private SSH Key** field. The **Admin Username** and **Admin Password** fields are not required if you select this option.

- **Upload new key:** Select this option to upload a new key. Public key content is extracted from the private key that is provided in the **Private SSH Key** field and is uploaded to the SAN Volume Controller cluster. If you select this action, the **Admin Username** and **Admin Password** fields are required.

Admin Username

If you selected **Upload new key** in the **Select Key** list, enter the name of a user that has the Security Administrator role for the SAN Volume Controller cluster that contains the device. Tivoli Storage Productivity Center uses this value during the setup process to contact the cluster and configure the SSH keys for the user. This user name must have privileges to modify other user accounts (Security Administrator role), otherwise Tivoli Storage Productivity Center cannot configure the SSH keys.

If you selected **Use existing uploaded key**, this field is optional.

Note: The **Admin Username** only applies to SVC5.x+; it does not apply to SVC 4.x.

Admin Password

Enter the password for the user that you entered in the **Admin Username** field.

Username

Enter an authorized user for the device. Tivoli Storage Productivity Center uses this value when monitoring and configuring the storage system. If the user name does not exist it is created. This user name must have the Administrator role.

This field is active if you selected **Upload new key** in the **Select Key** list and entered a user name and password in the **Admin Username** and **Admin Password** fields.

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

If you have this situation, create a user ID specifically for Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication on the SAN Volume Controller cluster. Add the SAN Volume Controller to Tivoli Storage Productivity Center for Replication using the new user ID and password (or change the user ID and password while in the disconnected state).

Select User

Click this button to select from a list of existing users that are defined on the SAN Volume Controller.

Private SSH Key

By default, the private key entry includes `${device.conf}\tpc_svc.pem`,

which represents the Tivoli Storage Productivity Center default key `tpc_svc.pem`. The `tpc_svc.pem` key is in the `conf` directory where the Device server is installed. If you want to use your own key, the key must exist on the system where you are running the Tivoli Storage Productivity Center user interface. The key is uploaded to the Tivoli Storage Productivity Center server system. To identify a new key, click **Browse**.

The private SSH key must be in OpenSSH format or in PuTTY (.ppk) format that is not password protected. If you have a password protected PuTTY .ppk key, use the PuTTYgen utility on a Windows operating system to convert that key to OpenSSH format by following these steps:

- a. Start the PuTTYgen utility.
- b. Load the .ppk file by clicking **File > Load private key**.
- c. From the menu bar, click **Conversion > Export OpenSSH key conversion**.

This action saves the key in OpenSSH format.

Passphrase (Optional)

Enter the passphrase for the SSH key pair. The passphrase protects your private SSH key file. If you do not have a passphrase, leave this field blank.

Click **Save**. The SAN Volume Controller is removed from the table after you updated the credentials.

After a successful migration of System Storage DS8000, XIV system, or SAN Volume Controller credentials, the CIMOMs that previously managed these devices must be removed from the Tivoli Storage Productivity Center configuration. The migration process tries to determine and remove these CIMOMs automatically, but if there are CIMOMs remaining, they must be deleted manually. For DS CIMOMs managing DS6000 or ESS devices and additional System Storage DS8000 devices, the System Storage DS8000 devices must be removed from the CIMOM if possible.

Resource History Retention

Specify how long to keep a history of the data that is collected by Tivoli Storage Productivity Center.

By specifying a number for days, weeks, or months for each element, you can control the amount of data that is retained and available for historical analysis and charting. The longer you keep the data, the more informative your analysis.

You can set retention periods for:

- Directories
- File systems
- Pings
- Computer Uptime
- Disks
- Storage subsystems
- Performance Monitors

Keep in mind the following considerations when specifying how long to retain historical data:

- If you do not select a check box, the data related to that check box is retained permanently. This action might cause a large amount of data to accumulate in the database repository over time.
- If you select a check box and enter a value of 0, Tivoli Storage Productivity Center stops collecting history data for the related element.

Note: If you enter 0 for the items under Performance Monitors, the data already collected by performance monitors is deleted and no more data is collected.

When you set the retention period for Performance Monitors, remember the following restrictions:

- The hourly and daily values for Performance Monitors can be set to 0 only if the per performance monitoring task is also set to 0.
- The hourly and daily values for Performance Monitors cannot be set lower than the per performance monitoring task value.

To set the retention periods, complete the following steps:

1. Expand **Administrative Services > Configuration > Resource History Retention**.
2. The Retain History panel is displayed in the content pane. Enter the information for the retention periods.

Note: Some default profiles have Accumulate history settings. If set, these settings dictate the number of days, weeks, or months to be retained. If a scan includes any of these profiles, then the history data is persisted and appears in the corresponding reports. The following default profiles have Accumulate history settings:

- TPCUser.By Access
- TPCUser.By Creation
- TPCUser.By Mod Not Backed Up
- TPCUser.Summary By File Type
- TPCUser.Summary By Filesystem/Directory
- TPCUser.Summary By Group
- TPCUser.Summary By Owner
- TPCUser.Temporary Files
- TPCUser.Wasted Space

Removed Resource Retention

You can specify the number of days to keep information about directories, file systems, disks, and so on, that have been removed from the system and can no longer be found.

Information is kept in the database repository. To clear the history record for a resource and activate a new period for resource retention, perform a discovery or probe job on the resource.

You can set the retention period for the following:

- Directories
- File systems

- Disks
- Filers
- Subsystems
- Fabric
- Tape
- Hypervisors
- Virtual Machines
- IBM Tivoli Storage Productivity Center Servers: When you remove a data source (agent) from a subordinate server or a storage entity that it monitors, the master server continues displaying information about that agent or storage entity in roll up reports for a specified amount of time. You can specify how quickly that information is removed by changing this setting.

If the check box for a storage resource is not selected, the data related to the resource is retained permanently. This action might cause a large amount of data to accumulate in the database repository over time. If you select a check box and enter a value of 0, the data related to that check box is removed immediately from the database repository.

The retention time for cluster resource groups that are monitored by Tivoli Storage Productivity Center is determined by the values that you enter in the **Filers** and **Filesystems** fields. If the values in the **Filesystems** and **Filers** fields are different, the field with the higher value (longer retention period) is applied to cluster resource groups.

The following list shows the fields that determine the retention period for Storwize V7000 File Module storage resources:

- **Filers:** Storage systems, storage pools, and nodes.
- **Disks:** Network Shared Disks (NSDs).
- **Filesystems:** File systems, exports, file sets, and quotas.

To set the retention period, complete the following steps:

1. Expand **Administrative Services > Configuration > Removed Resource Retention**.
2. The Retain Removed panel is displayed in the content pane. Select the resource and enter the number of days to keep.

Resource History Retention for Databases

You can specify how long to keep a history of the database-related statistical elements collected by the system.

By specifying a number of days, weeks, or months for each element, you can control the amount of data that is retained and is available for historical analysis and charting. The longer you keep the data, the more informative your analysis.

You can retain the histories for:

- Databases-Tablespaces
- Tables

Note: If you do not select a check box, the data related to that check box is retained permanently. This action might cause a large amount of data to

accumulate in the database repository over time. If you select a check box and enter a value of 0, the data related to that check box is removed immediately from the database repository.

To set the retention period, complete the following steps:

1. Expand **Administrative Services > Configuration > Resource History Retention for Databases**.
2. The Retain History panel is displayed in the content pane. Enter the retention period information.

Removed Resource Retention for Databases

You can specify how long to keep information in the enterprise repository that is related to a database entity that has been removed from the system and can no longer be found.

By specifying a number for days, you can indicate how long to keep information for table spaces and tables that have been removed from the system.

Note: If you do not select a check box, the data related to that check box is retained permanently. This action might cause a large amount of data to accumulate in the database repository over time. If you select a check box and enter a value of 0, the data related to that check box is removed immediately from the database repository.

To specify a retention period, complete the following steps:

1. Expand **Administrative Services > Configuration > Removed Resource Retention for Databases**.
2. The Retain Removed panel is displayed in the content pane. Enter the information for the retention period.

Configuration History Settings

This page allows you to specify how often the system captures snapshots of your configuration and when to delete them. It also gives the number of snapshots in the database, indicates when the last one was taken (and allows you to refresh this setting), and enables you to create and enter a title for a snapshot on demand.

You must configure and save the settings on this page before you can use the Configuration History function.

To specify the configuration history settings, complete the following steps:

1. In the Create snapshot every field, type how often (in hours) you want the system to take snapshot views of the configuration.
2. To automatically delete snapshots, click the check box to place a check mark before Delete snapshots older than. In the box that follows that field, type how long you want to keep the snapshots (in days) before they are automatically deleted.
3. To optionally refresh the date and time of when the latest snapshot was created, click **Update**.
4. To optionally create a snapshot on demand, click **Create Snapshot now**. If you choose, you might want to type a name for the snapshot in the box **Title this snapshot** (optional).
5. To change your settings to the default, click **Reset** to defaults. The defaults are create snapshots every 12 hours and delete snapshots older than 14 days.

6. Click **File > Save** to save the configuration history settings.

Configuration: Device server specific

This section provides configuration information that is specific to the Device server.

Configuring switches

This section provides configuration information for switches.

IBM Tivoli Storage Productivity Center is a storage area network (SAN) management application that discovers devices in the SAN and displays a topology of the SAN environment. Tivoli Storage Productivity Center is designed to operate using industry-based standards for communicating with Fibre Channel switches and other SAN devices. This communication can be done using the simple network management protocol (SNMP) interface for out-of-band agents, the FC-GS-3 interface for Storage Resource agents, Storage Management Initiative (SMI) agents, or a combination of these agent types. FC-GS-3 refers to the Fibre Channel Generic Services 3 standard. To gather and display the information as expected, the switches must be configured correctly. The configuration varies by vendor and the type or types of agents that are used. One of the common sources of customer problems is incorrect configuration of the switches that are being managed. This incorrect configuration leads to missing information and misconceptions that the Tivoli Storage Productivity Center product does not work with certain switches. The switch vendors supported are Brocade, Cisco, and QLogic. Other vendors such as IBM, often sell these switches under their own labels.

Determining the agent type or types to use with a switch

For Brocade fabrics, the preferred type of agent is the SMI agent. The SMI agent provides most fabric functions and the other agent types can be added for redundancy. However, Storage Resource agents are required to gather host bus adapter (HBA) information.

For QLogic and Cisco fabrics, a combination of different agent types is required to enable all functions.

For information about information that is gathered by each agent type, see the topic about discovering storage resources in the IBM Tivoli Storage Productivity Center User's Guide.

Using Storage Resource agents

With Tivoli Storage Productivity Center Storage Resource agent discovery, the Tivoli Storage Productivity Center for agent software is installed on SAN-attached hosts. The Storage Resource agents collect information about the fabric across the Fibre Channel network by querying the switch and the attached devices through the host bus adapter (HBA) in the system. For the switches to successfully receive and respond to the queries, there are some basic requirements:

- The switch must support the FC-GS-3 standard interface for discovery.
 - Name server
 - Configuration server
 - Unzoned name server

- For zone control functions, the fabric zone server must be supported, except in the case of Brocade switches.

For Storage Resource agent discovery, fabric events are automatically sent from the agent to Tivoli Storage Productivity Center. There is no need for configuration.

For Storage Resource agent discovery, you must configure SNMP traps to be sent from the switches in your fabric to the Tivoli Storage Productivity Center Device server.

Using out-of-band SNMP agents

Out-of-band SNMP agent discovery collects some of the same information that can be obtained by Storage Resource agents, but out-of-band agent discovery is performed differently. In out-of-band discovery, Tivoli Storage Productivity Center queries the switch directly rather than going through a Storage Resource agent and the Fibre Channel network. Tivoli Storage Productivity Center uses the SNMP protocol to send queries across the IP network to management information bases (MIBs) supported on the switch. Tivoli Storage Productivity Center uses the FC Management MIB (sometimes referred to as the FA MIB) and the FE MIB. The queries are sent only to switches that were added to Tivoli Storage Productivity Center for use as SNMP agents. SNMP information is collected for a single switch. The out-of-band discovery registers each switch. The following list shows the basic requirements for the switch to successfully receive and respond to the query:

- The FC Management MIB and FE MIB must be enabled on the switch.
- The switch must be configured to receive SNMPv1 queries and respond in SNMPv1. Some switches are configured to use SNMPv2 or SNMPv3 by default.
- The community string configured in Tivoli Storage Productivity Center must match one of the community strings configured on the switch with read access. Cisco switches must additionally have a community string match for write access. The default community strings in Tivoli Storage Productivity Center are "public" for read access and "private" for write access. Additional community strings can be defined on the switches, but are not used.
- SNMP access control lists must include the Tivoli Storage Productivity Center system. Some lists automatically include all hosts while others exclude all by default.

Another aspect of the SNMP configuration includes trap notification. SNMP traps are generated by the switch and directed to Tivoli Storage Productivity Center as an indication that something in the fabric changed and that a discovery must occur to identify the changes. The default configuration for handling switch traps is to send them from the switch to port 162 on the Tivoli Storage Productivity Center system. To successfully generate and receive traps, there are some configuration requirements:

- The trap destination parameter on the switch must be set. This parameter is the host that receives the trap and sends it to Tivoli Storage Productivity Center. The parameter is set on the switch.
- The destination port parameter on the switch must be set. Tivoli Storage Productivity Center listens on port 162 by default. The parameter is set on the host.
- The traps must be sent as SNMPv1. This parameter is set on the switch.
- The trap severity level must be set to generate traps for change conditions. This level typically means to send error level traps and anything more severe. This parameter is set in Tivoli Storage Productivity Center.

Using SMI agents

You must install or enable an SMI agent for to perform the following tasks:

- Gather fabric performance data.
- Collect and configuring Brocade zone aliases.

For information about installing or enabling a SMI agent for the switch, contact your switch vendor.

Configuring the inactive zone sets for a fabric data source

Some fabrics have multiple locations to retrieve a database for inactive zone sets. In some scenarios, and for some switch vendors, these definitions for inactive zone sets are not synchronized.

For fabrics composed of Brocade switches, IBM Tivoli Storage Productivity Center allows you to choose which CIM agent to use to retrieve inactive zone sets and when configuring zoning. To set the inactive zone sets for **Data Source**, go to **Fabric Manager > Fabrics**, select the fabric, and click **Select Inactive Zone Sets Data Source**.

Managing a SAN without agents

In the following situations, there might not be any agents on the SAN:

- The hosts do not currently have a Storage Resource agent or Fabric agent installed.
- The host operating system is not supported by the Storage Resource agent or Fabric agent.
- The customer requirements do not require the deployment of a Storage Resource agent or Fabric agent (topology map only).

In these cases, it is recommended that an agent is installed on the Device server itself. This action allows the Device server to use advanced features like Remote Node Identification (RNID) which requires an agent.

Normally the Device server does not have a Fibre Channel host bus adapter. In this configuration, the following steps are taken:

1. A Fibre Channel host bus adapter is added to the manager.
2. An agent is installed on the Device server(the Device server is installed first).
3. All storage devices are verified to ensure that they use LUN masking techniques. The LUN masking techniques prevent the Device server from accessing the disks used by the host systems.
4. The Fibre Channel host bus adapter is attached to the SAN to be managed. This host is added to each zone that is intended to be managed by the Device server.

Setting timeouts values for the Device server

If a probe or discovery of a storage subsystem times out before the operation completes, you can increase the timeout values for the Device server.

If a probe or discovery of a storage subsystem times out before the operation completes, you receive the following error message:

```
HWN021650E Encountered timeout while connecting to CIMOM IP:port.  
Check the CIMOM or increase timeout value.
```

where *IP* is the IP address, and *port* is the port number. If you determine that the Common Information Model Object Manager (CIMOM) is not the cause of the problem, you can use the command-line interface (CLI) to increase the timeout values for the Device server.

For those storage systems that use native interfaces to connect to Tivoli Storage Productivity Center you see this error message:

```
HWN020103E The external process exceeded the timeout limit and was cancelled.
```

The following storage systems use native interfaces to connect to Tivoli Storage Productivity Center:

- System Storage DS8000
- SAN Volume Controller
- The XIV system
- Storwize V7000
- Storwize V7000 Unified
- IBM SONAS

1. Run the **getdscfg** command to determine the current values of the timeout properties. From the command prompt, enter the following command:

```
cli>tpctool getdscfg -user user -pwd password -url host:port  
-property timeout_property timeout_value
```

where:

- *user* is an IBM Tivoli Storage Productivity Center user ID.
- *password* is the password for the Tivoli Storage Productivity Center user ID.
- *host* is the host name or IP address, and *port* is a valid port number for the HTTP service of the Device server. The default value for *port* is typically 9550.
- *timeout_property* is one of the following strings:
 - http.timeout
 - CIMClientWrapper.Timeout
 - Probe.Timeout.Array
 - Probe.Timeout.InBand
 - Probe.Timeout.LMM
 - Discovery.Timeout
 - CIMOMManager.TestConnectionTimeout

For the storage systems that use the native interface, the *timeout_property* strings are:

- NAPI.Timeout.TestConnection
 - NAPI.Timeout.Probe
 - NAPI.Timeout.EventPoll
- *timeout_value* is the value for the timeout.

2. Run the **setdscfg** command to increase the timeout value. Run the following command:

```
cli>tpctool setdscfg -user user -pwd password -url host:port  
-property timeout_property timeout_value
```

For more information about **tpctool**, see the *Tivoli Storage Productivity Center Command-Line Interface Reference*. You also can view help from the command line by issuing the command with the **-help** option.

Service Location Protocol (SLP) configuration considerations

This topic provides some configuration considerations for enabling Tivoli Storage Productivity Center to discover a larger set of storage devices through SLP. These recommendations cover some of the more common SLP configuration problems. This topic discusses router configuration, SLP directory agent configuration, and environment configuration.

For additional information about SLP, see the **Service Location Protocol Request for Comments** website at <http://www.ietf.org/rfc/rfc2165.txt>.

Note: The storage systems that use the native interfaces (DS8000, XIV system, SAN Volume Controller, and Storwize V7000) do not use SLP discovery.

Router configuration

Configure the routers in the network to enable general multicasting or to allow multicasting for the SLP multicast address and port, 239.255.255.253, port 427. The routers of interest are the ones associated with subnets that contain one or more storage devices that are to be discovered and managed by Tivoli Storage Productivity Center.

To configure your router hardware and software, refer to your router and configuration documentation.

SLP directory agent configuration

This topic provides information about suggestions for configuring the SLP directory agent.

Configure the SLP directory agents (DAs) to circumvent the multicast limitations. With statically configured DAs, all service requests are unicast by the user agent. Therefore, it is possible to configure one DA for each subnet that contains storage devices that are to be discovered by Tivoli Storage Productivity Center. One DA is sufficient for each of the subnets. Each of these DAs can discover all services within its own subnet, but no other services outside its own subnet. To allow Tivoli Storage Productivity Center to discover all the devices, it needs to be statically configured with the addresses of each of these DAs. This operation can be accomplished by using the Tivoli Storage Productivity Center Discovery Preference panel.

You can use this panel to enter a list of DA addresses. Tivoli Storage Productivity Center sends unicast service requests to each of these statically configured DAs, and sends multicast service requests on the local subnet on which Tivoli Storage Productivity Center is installed. Configure an SLP DA by changing the configuration of the SLP service agent (SA) that is included as part of an existing CIM Agent installation. This action causes the program that normally runs as an SLP SA to run as an SLP DA.

Note: The change from SA to DA does not affect the CIMOM service of the subject CIM Agent, which continues to function normally, sending registration and deregistration commands to the DA directly.

Environment configuration

This section provides information about the configuration of your environment.

It might be advantageous to configure SLP DAs in the following environments:

- In environments where there are other non-Disk Manager SLP UAs that frequently perform discovery on the available services, an SLP DA must be configured. This action ensures that the existing SAs are not overwhelmed by too many service requests.
- In environments where there are many SLP SAs, a DA helps decrease network traffic that is generated by the multitude of service replies. It also ensures that all registered services can be discovered by a given UA. The configuration of an SLP DA is recommended when there are more than 60 SAs that need to respond to any given multicast service request.

SLP registration and slptool

IBM Tivoli Storage Productivity Center uses Service Location Protocol (SLP) discovery, which requires that all the CIMOMs that Disk Manager discovers are registered using the SLP.

In a non-multicast network environment, SLP can only discover CIMOMs that are registered in its IP subnet. For CIMOMs outside of the IP subnet, you need to use an SLP DA and register the CIMOM using **slptool**. Ensure that the **CIM_InteropSchemaNamespace** and **Namespace** attributes are specified.

For example, enter the following command:

```
slptool register service:wbem:https://myhost.com:port
```

Where *myhost.com* is the name of the server hosting the CIMOM, and *port* is the port number of the service, for example 5989.

Note: **slptool** is installed with a CIMOM. Run the command from the machine hosting the CIMOM.

SLP discovery

A common problem with SLP discovery is due to IP multicasting being disabled on the network router. Communication between the SLP SA and UA is done with IP multicasting. Follow these recovery procedures when there are SLP discovery problems and IP multicasting is disabled on the network router.

Note: The storage systems that use native interfaces (DS8000, XIV system, SAN Volume Controller, and Storwize V7000) do not use SLP discovery.

There are two recovery procedures when there are SLP discovery problems and IP multicasting is disabled on the network router:

1. Configure one DA for each subnet within the environment.
2. Enable IP multicasting on the router which is disabled by default. Here is a list of common router configurations for multicasting:
 - Internet Group Management Protocol (IGMP) is used to register individual hosts in particular multicast groups and to query group membership on particular subnets.

- Distance Vector Multicast Routing Protocol (DVMRP) is a set of routing algorithms that use a technique called reverse path forwarding. These algorithms provide the best solution for how multicast packets are to be routed in the network.
- Protocol-Independent Multicast (PIM) comes in two varieties: dense mode (PIM-DM) and sparse mode (PIM-SM). The dense mode and sparse mode routines are optimized for networks where either a large percentage of nodes requires multicast traffic (dense) or a small percentage of nodes requires the sparse traffic.
- Multicast Open Shortest Path First (MOSPF) is an extension of OSPF. It is a link-state unicast routing protocol that attempts to find the shortest path between any two networks or subnets to provide the most optimal packet routing.

To properly configure the routers for multicasting, see the reference and configuration documentation from the router manufacturer.

Configuring Tivoli Storage Productivity Center for a Windows domain

There are two options available if you want the Tivoli Storage Productivity Center system to authenticate users against an instance of Microsoft Active Directory.

To authenticate user names on Tivoli Storage Productivity Center by using an instance of Active Directory, choose one of the following methods:

Authenticate user names: Method one:

1. Install Tivoli Storage Productivity Center on an AIX, Linux, or Windows operating system.
2. Configure Tivoli Storage Productivity Center for LDAP authentication.

The target LDAP repository is an instance of Active Directory running on Windows 2003 or later. This method provides the following benefits:

- The Tivoli Storage Productivity Center target system does not have to be a Windows system.
- The Tivoli Storage Productivity Center target system does not have to be a member of a Windows domain.
- There are fewer configuration steps for Tivoli Storage Productivity Center on Windows when you use this method.
- The Tivoli Storage Productivity Center user names and group names do not have to include the Windows domain name.
- You can use the single-sign on feature for Tivoli Storage Productivity Center and storage system element managers.

Authenticate user names: Method two:

1. Install Tivoli Storage Productivity Center on a Windows system that is a member of a Windows domain or on the Windows Domain Controller.
2. Configure Tivoli Storage Productivity Center for operating system authentication.

Configuring the Data server for a Windows domain

By default, the IBM Tivoli Storage Productivity Center Data server is installed to operate under a local account created at installation time. However, if you are

functioning in a Windows domain environment, and want to use domain accounts for authentication and authorization with Tivoli Storage Productivity Center, then you must reconfigure the service.

Note: These configuration steps are for a Windows system that is a member of a Windows domain and not for the Windows Domain Primary Domain Controller.

To configure the Data server, complete the following steps:

1. Open the Windows Services panel and verify that the Windows service **IBM Tivoli Storage Productivity Center - Data Server** is configured to log on as a domain user in order to perform Network Management Functions.
2. If the Data server is not configured as indicated in step 1, perform the following steps:
 - a. In the Windows Services panel, highlight **IBM Tivoli Storage Productivity Center - Data Server**.
 - b. Right-click the service menu.
 - c. Select **Properties**.
 - d. Select the **Log on** tab.
 - e. Specify a domain account under **This account** using the *domain_name\domain_user* format.
 - f. Enter and confirm the password for the domain account.
 - g. Click **Apply** for the setting, and then click **OK**.
 - h. Restart the Data server.
3. Log on to the Tivoli Storage Productivity Center GUI as a Tivoli Storage Productivity Center superuser.
4. Click **Administrative Services > Configuration > Role-to-Group mappings**.
5. Delete any instances of previous domain group mappings.
6. Save the changes.
7. Create new Role-to-Group mappings by specifying only the domain group name.
8. Save the changes.
9. You can now log on to the Tivoli Storage Productivity Center GUI with a domain account belonging to the domain group by using the *domain_name\domain_user* format.

Granting local administrative privileges to a domain account

For Windows users, the user account which the Storage Resource agent runs under requires local administrative rights. Because these rights are not necessarily guaranteed for domain users in a Windows domain environment, this topic provides information about how to grant local administrative rights for domain users. Using this procedure, you do not have to manually process each machine in the domain.

Note: These steps are for a Windows system that is a member of a Windows domain and not for the Windows Domain Primary Domain Controller.

To use Group Policy to grant local administrative privileges to a domain account, complete the following steps:

1. On the domain controller, go to **Administrative Tools > Active Directory Users and Computers** (you must be running with Domain Administrator privileges).

2. Right-click on the Organizational Unit (OU) upon which you want to apply the Group Policy. Click **Properties**.
3. The Group Policy Properties panel is displayed. Select the Group Policy tab and click **New** to create a Group Policy.
4. Designate a name for the new Group Policy. Select the new Group Policy and click **Edit**.
5. The Group Policy Object Editor panel is displayed. Go to **New Group Policy Object *your_policy* > Computer Configuration > Windows Settings > Security Settings > Restricted Groups**. Right-click on **Restricted Groups**. Click **Add Group**.
6. For example, name the new group "Administrators." Under "Properties", add the user "Administrator", and the domain accounts or groups upon which you want the Group Policy in effect for. For example, you can add "TPC\tapeadmin", "TPC\tapegroup", and "TPC\TestGroup". Click **OK**.
7. Add these user rights to the domain account:
 - Act as part of the operating system
 - Log on as a service

In the Group Policy Object Editor, go to **New Group Policy Object *your_policy* > Computer Configuration > Windows Settings > Security Settings > Local Policies > User Rights Assignments**. In the content pane, select "Log on as a service" and double-click. Add the domain user for whom you are granting user rights and click **OK**. Repeat this step for "Act as part of the operating system."

8. The group policy is now enforced for the Organizational Unit to include the domain accounts and groups specified under the local Administrators group on each computer in the Organizational Unit. In addition, the domain user has been granted the necessary rights. To verify the user rights, log in to a domain computer and open the Computer Management console. Select **Groups**, double-click the Administrators group, and verify the membership of the domain users.

Adding Tivoli Storage Productivity Center to the Windows domain

Use this information to add Tivoli Storage Productivity Center to the Windows domain for OS authentication. This procedure requires two steps, one for configuring Windows and one for configuring Tivoli Storage Productivity Center.

Configuring Windows

Note: These steps are for a Windows system that is a member of a Windows domain and not for the Windows Domain Primary Domain Controller.

Add the DNS IP as an IP of the domain controller, follow these steps:

1. Click **Start > Settings > Control Panel > Network Connections**.
2. Right-click **Local Area Connection**. Click **Properties**.
3. The Local Area Connection Properties window is displayed. Select **Internet Protocol Version 4 (TCP/IPv4)**. Click **Properties**.

Note: This procedure uses Windows 2008. If you use Windows 2003, select **Internet Protocol** .

4. The Internet Protocol Version 4 window is displayed. On the General tab, enter the IP address of the domain controller in the **Preferred DNS server** field. Click **Advanced**.
5. The Advanced TCP/IP Settings window is displayed. On the DNS tab, enter the domain controller information in the **DNS Suffix for this connection** field. Select the check box for **Register this connection's addresses in DNS**. Click **OK**.
6. Right-click **My Computer** > **Properties**. Click **Change settings**.
7. The System Properties window is displayed. On the Computer Name tab, click **Change**.
8. The Computer Name/Domain Changes window is displayed. Click **More**.
9. The DNS Suffix and NetBIOS Computer Name window is displayed. Enter the domain controller information in the **Primary DNS suffix of this computer** field. Click **OK**.
10. The Computer Name/Domain Changes window is displayed. Click **Domain**.
11. Provide the authentication credentials to the domain controller.
12. Restart the computer for the changes to take effect.

Configure Tivoli Storage Productivity Center with domain controller credentials

Follow these steps:

1. On the desktop, click **My Computer** > **Manage**.
2. Click **Configuration** > **Local Users and Groups** > **Groups**.
3. In the Groups pane, double-click **Administrators**.
4. The Administrators window is displayed. Click **Add**.
5. The Select Users, Computers, or Groups window is displayed. Enter the domain controller user in the object names to select field. Click **Check Names**. Click **OK** twice to exit each window.
6. Click **Configuration** > **Services**.
7. Perform these steps for each of the services in the content pane:
 - IBM Tivoli Storage Productivity Center - Data Server
 - IBM WebSphere Application Server V6.1 - CSM
 - IBM WebSphere Application Server V6.1 - DeviceServer
 - Tivoli Integrated Portal - TIPProfile_Port_xxxx.

To view information about Tivoli Integrated Portal port numbers, see "Tivoli Integrated Portal port numbers" on page 35.

 - a. Double-click the service.
 - b. Click **Properties**.
 - c. On the Properties window, click the **Log On** tab.
 - d. On the Log On tab, select **This account**. Enter the domain controller password in the Password and Confirm password fields. Click **OK**.
 - e. In response to the following message, click **OK**:


```
The account domain_controller\user_ID has been granted the Log On
As a Service right.
```
 - f. Stop and restart each service.
8. You can now log in to Tivoli Storage Productivity Center (as well as Tivoli Storage Productivity Center for Replication) using the domain controller and user ID.

Configuring IP addressing

This section provides information about configuring IP addressing.

Configuring Tivoli Storage Productivity Center with multiple IP addresses

If the system where IBM Tivoli Storage Productivity Center is to be installed has multiple IP addresses, then a configuration value must be set manually as a post-installation task using the **tpctool setdscfg** command. The value to be set is for the local IP address which must be used for subscription for CIM Indications for CIM agents.

Note: This task does not apply to storage systems that use the native interfaces (DS8000, XIV system, SAN Volume Controller, and Storwize V7000).

If you are using IPv6 machines, see “Planning for Internet Protocol Version 6” on page 46.

For multiple IPv6 addresses, the IPv6 address to use for CIM indication subscription by Tivoli Storage Productivity Center can be specified by setting the property `System.LocalIPv6Address` as described.

With dual stack IPv4 and IPv6 Tivoli Storage Productivity Center servers, two IP addresses are required to subscribe to IPv4 CIMOMs and IPv6 CIMOMs. The configuration property `System.LocalIPv6Address` is used for IPv6 CIMOMs and the property `System.LocalIPAddress` is used for IPv4 CIMOMs.

To change the IP address, follow these steps:

1. Open a command prompt window on the server system.
2. Change to the following directory:

```
cd TPC_installation_directory\cli
```

3. Enter the following command:

```
tpctool setdscfg -user user_ID -pwd password -url  
host:port -property System.LocalIPv6Address value
```

Where:

user_ID

is the user ID.

password

is the password for the user.

host is either the host name or IP address of the system running Tivoli Storage Productivity Center.

port is a valid port number for the HTTP service of the Device server (the default is 9550).

value is the local IP address which must be used for subscription for CIM Indications for CIM agents.

4. Verify that the command was successful by entering this command:

```
tpctool getdscfg -user user_ID -pwd password  
-url host:port -property System.LocalIPv6Address
```

Changing the HOSTS file

When you install IBM Tivoli Storage Productivity Center on your Windows computers, you must follow these steps to avoid addressing problems with the systems you want to manage. These problems are caused by the address resolution protocol returning the host short name rather than the fully qualified host name. You can avoid this by modifying the entries in the corresponding host tables on the DNS server and on the local computer system. The fully qualified host name must be listed before the short name in each entry associated with systems managed by IBM Tivoli Storage Productivity Center.

The **HOSTS** file is located in the %SystemRoot%\system32\drivers\etc\ directory. To change the HOSTS file, follow these steps:

1. Open the HOSTS file in a text editor.
2. Add, remove, or modify the host entries. In the following example of a HOSTS file, the short name is incorrectly listed before the fully qualified host name. This can cause address resolution problems in IBM Tivoli Storage Productivity Center.

```
# Copyright (c) 1993-1995 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows NT.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name.
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be
# inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#       102.54.94.97       rhino.acme.com   # source server
#       38.25.63.10      x.acme.com       # x client host
#
192.168.123.146          jason           jason.groupa.mycompany.com
```

3. In the following example, the order of the host names has been changed so that the fully qualified host name is placed before the short name. The host names must be entered in the order shown so IBM Tivoli Storage Productivity Center can locate the host. Use this format for any hosts associated with IBM Tivoli Storage Productivity Center.

```
# For example:
#
#       102.54.94.97       rhino.acme.com   # source server
#       38.25.63.10      x.acme.com       # x client host
#
192.168.123.146          jason.groupa.mycompany.com          jason
```

Note: Host names are case sensitive. This is a WebSphere requirement. For example, if your computer shows the name as JASON (uppercase), then you must enter JASON in the HOSTS file.

Configuring tracing

Configure tracing for the Data server, agent, scheduler, Device server, and GUI by setting the debug level and log file characteristics. The level of information saved in a log file is determined by the tracing level. Tracing settings can only be changed by users with administrator authority.

To configure tracing for IBM Tivoli Storage Productivity Center, complete the following steps:

1. In the navigation tree pane, expand **Administrative Services > Services > Data Server**.
2. Right-click on one of the following components:
 - Server
 - Agent
 - Scheduler
 - Device
 - GUI

Click **Configure Tracing**.

3. In the Server Tracing Configuration window, check the **Enable Trace** check box. Specify additional tracing options:

Level Select the level of debugging you want to perform. Select one of the following debugging level:

DEBUG_MIN

Provides minimal information about errors. This option is the default value.

DEBUG_MID

Provides more information about errors.

DEBUG_MAX

Lists all debugging trace information.

Maximum number of files

Select the maximum number of trace files that is created. When the maximum number of files has been reached, tracing rolls over and start writing to the first file. (default is five)

The value in this field also applies to the other services for the Data server for which tracing is enabled.

Maximum File Size (kilobytes)

Select the maximum size of each log file. When the current trace file has reached its maximum capacity, tracing continues to the next available trace file as defined in the Maximum Number of Files field. (default is 20 MB).

4. Click **OK**.

Tracing can be turned off by clearing the **Enable Trace** option in the Server Tracing Configuration window.

Configuration guidelines for 500 or more agents

This topic provides information if you have 500 or more agents in IBM Tivoli Storage Productivity Center.

If you have 500 or more agents for the Data Server, follow these steps.

Note: Probes can overwhelm Data Server resources, because by default, they can return information all at once.

1. Probe the machines at least once a day or more, depending on when you want to test for alert conditions. This action applies to alert conditions other than directory alerts, quotas, or constraints.
2. If you use anything but the "ALL" groups (ALL file systems, ALL computers), you need to manually populate the groups.
3. Always run a probe before a scan. Set the following parameters in the `server.config` file:

MaxConnections=1200

The default is 500. Agents can have multiple connections to the server.

routerThreads=3 (max)

Incoming connections need to be routed to the correct Data Manager "service" queue and can stack up behind this thread. You can watch this behavior by watching the connections in the "server service" and the "agent service" in the GUI. The server service runs the router and the agent service is where the connections are queued once routed and saved by any of three threads here to the repository.

4. Set the following parameter in the `Scheduler.config` file:

MaxSubmitthreads=8

Tells how many threads are used to tell the agents to start a job. Agent connections can queue up the scheduler service. After a job is run, the agent makes a connection to communicate with this thread to give it the job status.

Configuring LUN provisioning for Oracle Solaris

Tivoli Storage Productivity Center for Data provides a file system extension feature that can be used to automatically increase file system capacity for managed hosts when utilization reaches a specified level. This function allows for the automatic provisioning of (TotalStorage Enterprise Storage Server, DS6000, DS8000) LUNs when there is not enough space available in a volume group to extend a file system. This section provides information about LUN provisioning for Solaris.

LUNs can be provisioned for file system hosts that run Solaris, but you must configure the hosts must to avoid a restart after provisioning. Before you install the Storage Resource agent, complete the following steps:

1. Assign TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUNs to Solaris Host Bus Adapters (HBAs).
2. Modify the HBA configuration file to include persistent name binding.
3. Modify the SCSI Disk configuration file to allow the maximum number of LUNs.
4. If you are using multipathing, ensure that TotalStorage Enterprise Storage Server, DS6000, or DS8000 multipaths are detected by the Veritas Dynamic Multipathing (VxDMP) utility.

This section provides basic instructions for performing these configuration steps. For detailed information, see the HBA and VxDMP documentation.

Assigning TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUNs to Oracle Solaris HBAs

This section provides information about assigning TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUNs to Solaris HBAs.

You must assign at least one TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUN to each HBA on the Solaris host.

If you are using multipathing, there are different ways to configure either the host and TotalStorage Enterprise Storage Server, DS6000, or DS8000. For example:

- For an TotalStorage Enterprise Storage Server, DS6000, or DS8000 without internal multipath configuration, assign the same LUNs to the World Wide Port Node (WWPN) of each HBA.
- For an TotalStorage Enterprise Storage Server, DS6000, or DS8000 with internal multipath configuration, assign the LUNs to the WWPN of one HBA or assign the same LUNs to the WWPNs of two or more HBAs.

Modifying the HBA configuration file

The HBA configuration file must be modified to include Persistent Name Binding on HBAs and targets so that both the controller and target numbers remain the same across system reboots. This section provides information about what to modify in the configuration file.

The HBA configuration file (for example, `qla2200.conf`) must be modified to include Persistent Name Binding on HBAs and targets so that both the controller and target numbers remain the same across system reboots. You must reboot the system with the new configuration for the changes to take effect.

QLogic QLA2200 and QLA2300 HBAs have been tested for use with IBM Tivoli Storage Productivity Center. You can use the QLogic SANblade Control FX (`scfx`) application to modify the configuration file for these HBAs. The `scfx` application is included as part of the device driver installation package. The `scfx` application is installed in the `/opt/QLogic_Corporation/SANblade_Control_FX` directory.

Setting Persistent Name Binding for QLA2200 and QLA2300 HBAs by using the `scfx` command

This section describes how to set Persistent Name Binding in the HBAs by using the `scfx` command for LUN provisioning under Oracle Solaris.

Follow these steps:

1. Install the QLogic HBA Driver, Common API Library, and QLogic SANblade Control FX (`scfx`) application if you have not already done so. For installation instructions, see the *SANblade 2200 Series User's Guide* or *SANblade 2300 Series User's Guide*. After these packages are installed successfully, restart and reconfigure the system by using the `reboot -- -rv` command.
2. After the system is rebooted, use `scfx` to configure Persistent Bind on HBAs and Targets in the `/kernel/drv/qla2xxx.conf` file.
 - a. Start the `scfx` application. For example:

```
# /opt/QLogic_Corporation/SANblade_Control_FX/scfx
```

The main window of the **scfx** application consists of three sections:

Menu Bar

The menu bar provides three options: **File**, **Tools**, and **Help**.

HBA Tree

The HBA Tree displays the host with its connected adapters (HBAs), devices and LUNs. The HBAs are displayed with a model name and instance number. For example, Adapter 2200 (Instance #0). If a device is connected to an HBA, it has a plus sign (+) by the HBA, which can be expanded to view the list of attached devices. The devices are listed with their World Wide Port Names (WWPN).

Click the plus sign next to a device to expand the tree to show all the LUNs in that device. For a RAID device, such as an TotalStorage Enterprise Storage Server, DS6000, or DS8000, there are multiple LUNs per device.

Note: Expand all the devices to search the TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUNs assigned to the system and note the WWPN of the target device. This information is required to identify the SCSI Target ID assigned or specified for the Persistent Bind Targets Setting.

Tabbed Pages

The contents of the Tabbed Pages changes depending on what is currently selected in the HBA Tree.

- b. Select an HBA.

Select an adapter in the HBA Tree. The Tabbed Pages show the **HBA Information**, **HBA Options**, **Target Settings**, **Boot Device**, **Diagnostics**, and **Utilities** tabs.

- c. Select the Persistent Bind HBA Setting.

Click the **HBA Options** tab. In the **Select Parameter Section** drop-down list, select **Advanced Host Parameters**. Select the check box for **Persistent Bind HBA**. Click **Save**.

- d. Select the Persistent Bind Target Setting.

Click the **Target Settings** tab. Select the check box for each target in the **Bind** column. If the check boxes are already checked and disabled, proceed to the next step. In the **Target ID** column, you can either accept the pre-selected SCSI Target ID or change to a different value. Each SCSI target ID must be unique and range from 0 to 255.

Note: Write down the selected Target ID for each TotalStorage Enterprise Storage Server subsystem device.
Click **Save**.

- e. Repeat Steps b through d for the next HBA.
- f. Exit the **scfx** application.

From the **Menu Bar**, select **File | Exit**. A Reboot Reminder dialog is displayed. Click **OK** to exit.

3. Restart and reconfigure the system by using the **reboot -- -rv** command.

Modifying the SCSI disk configuration file

You must configure the SCSI disk configuration file for the maximum number of LUNs per target for LUN provisioning for Oracle Solaris.

You must configure the SCSI disk (sd.conf) configuration file for the maximum number of LUNs (256) per target. The system must then be rebooted with the new configuration for the changes to take effect. Follow these steps:

1. Identify the SCSI Target ID assigned to the TotalStorage Enterprise Storage Server.
2. Edit the /kernel/drv/sd.conf file to include all the possible target and LUN mappings for the RAID device. For example, assume the SCSI Target ID assigned for an TotalStorage Enterprise Storage Server is 2. You can allow up to 256 LUNs (0 - 255) for this target:

```
name="sd" class="scsi" target=0 lun=0;
name="sd" class="scsi" target=1 lun=0;
name="sd" class="scsi" target=2 lun=0;
name="sd" class="scsi" target=2 lun=1;
name="sd" class="scsi" target=2 lun=2;
:
:
name="sd" class="scsi" target=2 lun=253;
name="sd" class="scsi" target=2 lun=254;
name="sd" class="scsi" target=2 lun=255;
name="sd" class="scsi" target=3 lun=0;
name="sd" class="scsi" target=4 lun=0;
:
:
name="sd" class="scsi" target=253 lun=0;
name="sd" class="scsi" target=254 lun=0;
name="sd" class="scsi" target=255 lun=0;
```

In this example, the system can detect up to 256 targets with 1 LUN (for example, multiple RAID devices with a total of 256 LUNs) and up to 256 LUNs for target 2 (for example, a RAID device with a total of 256 LUNs).

3. Restart and reconfigure the system by using the **reboot -- -rv** command.

Checking for TotalStorage Enterprise Storage Server, DS6000, or DS8000 multipaths in VxDMP

If you are using IBM TotalStorage Enterprise Storage Server, DS6000, or DS8000 LUNs with multipaths, you must ensure that all the paths are detected by Veritas Dynamic Multipathing (VxDMP) utility. This section provides information about how to check for multipathing in the VxDMP utility.

The VxDMP utility is an administrative interface to the Veritas Volume Manager (VxVM) Dynamic Multipathing (DMP) facility. It lists the paths under a DMP device, gets the DMP device corresponding to a path, lists all the disk controllers on the system, lists all the paths through a host disk controller, lists all the DMP nodes through a disk array, and enables or disables a host disk controller on the system. For more information, and detailed instructions, see the VxDMP documentation.

To list all disk controllers on the system, enter the following command:

```
# vxdmpadm listctlr all
```

The following sample output shows that controllers c3 and c4 are connected to the IBM TotalStorage Enterprise Storage Server with an Enclosure Type of IBM_SHARK and an Enclosure Name of IBM_SHARK0.

CTLR-NAME	ENCLR-TYPE	STATE	ENCLR-NAME
c1	Disk	ENABLED	Disk
c3	IBM_SHARK	ENABLED	IBM_SHARK0
c4	IBM_SHARK	ENABLED	IBM_SHARK0

To list all subpaths for controller c3, enter the following command:

```
# vxdkmpadm getsubpaths ctlr=c3
```

The following sample output shows that the **DMPNODENAME** is the same as the device name for each TotalStorage Enterprise Storage Server LUN:

NAME	STATE	PATH-TYPE	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME
c3t4d0s2	ENABLED	-	c3t4d0s2	IBM_SHARK	IBM_SHARK0
c3t4d1s2	ENABLED	-	c3t4d1s2	IBM_SHARK	IBM_SHARK0

To list all subpaths for controller c4, enter the following command:

```
# vxdkmpadm getsubpaths ctlr=c4
```

The following sample output shows that the **DMPNODENAME** for each TotalStorage Enterprise Storage Server LUN is from controller c3. This means that VxDMP refers to the TotalStorage Enterprise Storage Server, DS6000 or DS8000 LUNs as devices from controller c3 and mask devices on controller c4 from VxVM:

NAME	STATE	PATH-TYPE	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME
c4t4d0s2	ENABLED	-	c3t4d0s2	IBM_SHARK	IBM_SHARK0
c4t4d1s2	ENABLED	-	c3t4d1s2	IBM_SHARK	IBM_SHARK0

How to import authentication information for a Storage Resource agent

The Storage Resource agent is installed as a non-daemon or daemon process. IBM Tivoli Storage Productivity Center stores the authentication information to connect to the host on which the Storage Resource agent has installed for the non-daemon agent. This authentication information can be changed depending on the environment.

To change a Storage Resource agent's authentication information (for non-daemon service), follow these steps:

1. Export the authentication information for a Storage Resource agent. For more information, go to the information center. Search for **Exporting a Storage Resource agent list**.
2. The data file exported contains information such as the host name, user ID, password, certificate location, and passphrase for every agent selected. The information is separated by the pipe character (`|`). For example,


```
agent_host|user|password|certificate|passphrase
```

You can update the password or passphrase in encrypted format or plain text format. If you want to update the password or passphrase in encrypted format, then you can use the **tpctool**. For example go to this directory and run the **tpctool**:

```
cd TPC_install_directory/cli
tpctool encrypt string_to_be_encrypted
```

This generates an encrypted string. Place this string in the data file to be imported and add `@ENC@` to the end of the encrypted string. For example,

```
agent_host|usera|encrypted_password@ENC@|certificate|
encrypted_passphrase@ENC@
```

encrypted_password is the encrypted string for the password and
encrypted_passphrase is the encrypted string for the passphrase.

3. Import the data file. For more information, go to the information center. Search for **Updating a Storage Resource agent list**.

Configuring Tivoli Storage Productivity Center for DS8000 LDAP authentication

You must configure IBM Tivoli Storage Productivity Center to use LDAP for single sign-on support for the DS8000 R4.2.

Overview

Configuring Tivoli Storage Productivity Center and DS8000 for single sign-on involves this process:

For Tivoli Storage Productivity Center

1. Extract the certificate. This certificate is used for securing communication between the Authentication Client on the HMC and the Authentication Service (server component) on Tivoli Storage Productivity Center.
2. Create a truststore which includes the certificate from step 1.
3. Know the URL for the Authentication Service.

For DS8000

1. Create a Storage Authentication Service (SAS) policy with information collected from Tivoli Storage Productivity Center and the LDAP server.
2. Test the Storage Authentication Service policy by using a valid LDAP user mapped to a DS8000 user role in the policy.
3. Activate the Storage Authentication Service policy by using a valid LDAP user mapped to the DS8000 administrative user role in the policy.

Configuring Tivoli Storage Productivity Center

This procedure assumes that Tivoli Storage Productivity Center is configured with an LDAP repository for authentication. To configure Tivoli Storage Productivity Center:

1. You need to know the URL for the Authentication Service.

Tip: An example of the Authentication Service URL is:

```
https://TIP_server_host:9569/TokenService/services/Trust
```

Here is an example:

```
https://tpcserver1.storage.mycompany.com:9569/TokenService/services/Trust
```

The port for the Authentication Service is one plus the default IBM Tivoli Integrated Portal port. The port numbers for IBM Tivoli Integrated Portal can vary. To view information about Tivoli Integrated Portal port numbers, see

“Tivoli Integrated Portal port numbers” on page 35. Contact your Tivoli Storage Productivity Center administrator if you need to verify the Tivoli Integrated Portal port.

2. Access the Tivoli Integrated Portal administrative console. Start a web browser and enter the following in the **Address** field:

`http://hostname:port/ibm/console`

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. To view information about Tivoli Integrated Portal port numbers, see “Tivoli Integrated Portal port numbers” on page 35. Contact your Tivoli Storage Productivity Center administrator if you need to verify the host name and port number. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have Tivoli Integrated Portal administrator authorization.

3. Create the truststore in Tivoli Integrated Portal. In the Tivoli Integrated Portal navigation tree, click **Settings > WebSphere Administrative Console**. In the Tivoli Integrated Portal content pane, click **Launch WebSphere administrative console**. In the WebSphere Administrative Console, click **Security > SSL certificate and key management > Key stores and certificates > NodeDefaultKeyStore > Personal Certificates**. On this page, select the default certificate and click **Extract**. On the next page, enter the following information:

Certificate file name

Enter a file name for the extracted certificate. This file automatically is created in `TIP_installation_directory/profiles/TIPProfile/etc/`.

For Windows, the default directory is `C:\TIP_installation_directory\profiles\TIPProfile\etc\`.

Accept and select the default data type. Click **OK**.

4. Create the truststore file and import the certificate into the truststore file by using the **ikeyMan** tool.

- a. Start the **ikeyMan** tool.

For example on Windows: `C:\Program Files\IBM\tipv2\bin\ikeyman.bat`

- b. Click **Key Database File > New**. On the New panel, enter the following information and click **OK**:

Key database type

In the drop-down list, select **JKS**.

File Name

Enter a file name. For example, enter `tpc_ess.jks`.

Location

Enter a location. For example, enter `C:\tpc\`. Click **OK**.

Tip: The default location is: `C:\Program Files\IBM\tipv2\bin`

- c. The next panel prompts you to specify a password for this truststore. Specify a password that you can remember. Click **OK**.
- d. On the main Key database information panel, in the Key database content section, select **Signer Certificates** and click **Add**. In the Open panel, click **Browse** and select the certificate file you created in step 3. Click **Open** and **OK**.

Note: Look for the certificate file, change the Files of Type to **All files**. Click **Open**.

- e. You see a prompt to specify a label. Provide any label. An example of a label is: ESS_Cert. Click **OK**.
 - f. The **ESS_Cert** is now one of the listed Signer Certificates.
 - g. Exit the **ikkeyman** tool and locate the truststore file you created in Step 4b (for this example, tpc_ess.jks). You need this truststore file and the password for configuring the LDAP-based policy on DS8000.
 - h. You are now finished with IBM Tivoli Integrated Portal and the truststore setup.
5. Find the user ID and password that is used in LDAP to use for the DS8000 Storage Authentication Service policy configuration page.
This user ID is used for authenticating with the Authentication Service. It can be any user ID in LDAP, or a user ID that is also used by Tivoli Storage Productivity Center. This user ID is used as the *Application Client User ID* for a Storage Authentication Service policy on the DS8000.
 6. Find the name of a group in LDAP with which you can log in to Tivoli Storage Productivity Center and the DS8000. You would use this LDAP group on the DS8000 also, for mapping to DS8000 roles.
You can go to the Tivoli Storage Productivity Center **Role-to-Group Mapping** node to find out which LDAP group is mapped to the role in Tivoli Storage Productivity Center.
To find the LDAP group name, open the Tivoli Storage Productivity Center GUI and click **Tivoli Storage Productivity Center > Configuration > Role-to-Group Mapping**.
The information gathered in steps 1, 3, 4, 5, and 6 is used on the DS8000 Storage Authentication Service policy creation page.
 7. Configure DS8000 R 4.2.

Configuring DS8000 for LDAP authentication

To configure DS8000 for LDAP authentication, complete the following steps:

1. Add the IP address of the DS8000 Hardware Management Console to the Internet Explorer list of trusted sites by using the following steps:
 - a. Open the Internet Explorer by clicking the Internet Explorer icon on the Quick Launch toolbar.
 - b. From the Internet Explorer toolbar, click **Tools > Internet options**.
 - c. Click the **Security** tab, click the **Trusted sites** icon and then click **Sites**.
 - d. In the **Add this web site to the zone** field, type the IP address of the DS8000 Hardware Management Console (HMC). Click **Add** and the IP address is added to the **Websites** field.
 - e. Click **Close** and then click **OK** to exit the Internet Options window, and then close the Internet Explorer.
2. To access the DS8000 GUI, complete the following steps:
 - a. In the Tivoli Storage Productivity Center GUI, expand the **Element Manager** tree.
 - b. Click DS8000.
 - c. Complete one of the following steps:
 - In the content pane, select a DS8000 system.
 - Click **DS8000 Element Manager > Select Action > Add Element Manager**.

- d. In **Select Action**, select **Launch Default Element Manager**. This action opens the DS8000 Storage Manager GUI so that you can administer DS8000. Enter the user name and password and click **OK**.
3. On the DS8000 Storage Manager Welcome page, click **Real-time manager > Monitor System > User Administration**.
4. On the User and Authentication Policy Administration Summary page, select a Complex Name. Under the Select action menu, select **Create Storage Authentication Service Policy**.
5. The Authentication Service Configuration page is displayed. Enter the following information:
 - Policy Name
 - Authentication Service URL (primary)
 - Authentication Service Client User ID
 - Authentication Service Client Password
 - Confirm Authentication Service Client Password

Click **Next**.

Tip: An example of the authentication URL is:

`https://TIP_server_host:9569/TokenService/services/Trust`

Here is an example:

`https://tpcserver1.storage.mycompany.com:9569/TokenService/services/Trust`

The port for the Authentication Service is one plus the default Tivoli Integrated Portal port. The Tivoli Integrated Portal port number can vary. To view information about Tivoli Integrated Portal port numbers, see “Tivoli Integrated Portal port numbers” on page 35. Contact your Tivoli Storage Productivity Center administrator if you need to verify the Tivoli Integrated Portal port.

6. In the Truststore file Information page, enter the following information:
 - Truststore File Location
 - Truststore File Password
 - Confirm Truststore File Password
- Click **Next**.
7. In the Map External Users and User Groups to DS8000 User Roles page, enter the following information:
 - External Entity Name
 - External Entity Type
 - DS8000 User Role

Click **Add**. The entry is entered in the table on this page. Select the entry you created and click **Next**.

8. The Verification page is displayed. Verify the information and click **Next**.
9. The Summary page is displayed. Click **Activate the Policy** if you want to activate the policy immediately. If you want to test the policy before activating it, do not select **Activate the Policy** and click **Finish** to create the policy. This scenario assumes that you want to test the policy before activating it. You see a message dialog that indicates whether the policy was successfully created. If the policy was successfully created, close the message dialog.

10. The Manage Authorization Policy page is displayed. Select a policy. Under the Select action menu, click **Test Authentication Policy**.
11. The Test Storage Authentication Service Policy page is displayed. Enter the following information:
 - External User Name
 - External User Password

Provide an LDAP user ID and password for External User Name and External User Password. The user ID must already be mapped to a valid DS8000 user role in the Storage Authentication Service policy. This user ID does not have to be in the Administrator group. Click **OK**.
12. The Manage Authentication Policy page is displayed. Select the policy you want. Under the Select action menu, click **Activate Authentication Policy**.
13. The Activate Storage Authentication Service Policy page is displayed. Enter the following information:
 - External User Name
 - External User Password

Provide an LDAP user ID and password for External User Name and External User Password. The user ID must already be mapped to a valid DS8000 user role in the Storage Authentication Service policy. This user ID must be in the Administrator group. Click **OK**. The policy is now activated. Close the page.
14. To access the DS8000 GUI again from Tivoli Storage Productivity Center, update the username and password:
 - a. In the stand-alone GUI, expand the **Element Manager** tree.
 - b. Click DS8000.
 - c. In the content pane, select the DS8000 system you just configured.
 - d. Click **DS8000 Element Manager > Select Action > Modify Element Manager**.
 - e. Change the user name and password to the External User Name and External User Password that were previously used to configure the DS8000 system.

Configuring multiple Tivoli Storage Productivity Center servers with one DS8000 R4.2

You can configure multiple Tivoli Storage Productivity Center servers to use LDAP for single sign-on support for the DS8000 R4.2.

Configuring multiple Tivoli Storage Productivity Center servers

To configure multiple Tivoli Storage Productivity Center servers:

1. Configure one server as described in “Configuring Tivoli Storage Productivity Center for DS8000 LDAP authentication” on page 322. This server is called TPC_server1.
2. Install a second Tivoli Storage Productivity Center server and configure it with the same LDAP repository information as the first Tivoli Storage Productivity Centerserver. The second server is called TPC_server2.
3. Open a command prompt window on TPC_server2:


```
TIP_installation_directory/bin
```
4. Run the **wsadmin** command to export LTPA keys from TPC_server1 into a file on TPC_server2.

```
wsadmin -user TIP_admin_ID -password TIP_admin_password -lang jython
-port TIP_SOAP_port -host TPC_server1_hostname_or_IP_address
-f "TPC_install_dir_on_TPC_server2/tip/scripts/exportLTPAKeys.py"
"LTPA_keys_file_name" LTPA_keys_password
```

An example is:

```
C:\Program Files\IBM\tipv2\bin>wsadmin -user tpcsuperuser
-password tpcsuperuser -lang jython
-port 9571 -host 9.56.98.41
-f "c:/program files/ibm/tpc/tip/scripts/exportLTPAKeys.py"
"c:/share/ltpaKeys_serv1" ltpa123
```

This creates a file on TPC_server2 named ltpaKeys_serv1 that contains the LTPA keys of TPC_server1.

Note: Use forward slashes.

5. In the same command window, run the following command to import the LTPA keys into Tivoli Integrated Portal on TPC_server2.

```
wsadmin -user TIP_admin_ID -password TIP_admin_password -lang jython
-f "TPC_install_dir_on_TPC_server2/tip/scripts/importLTPAKeys.py"
"LTPA_keys_file_name" LTPA_keys_password
```

An example is:

```
C:\Program Files\IBM\tipv2\bin>wsadmin -user tpcsuperuser
-password tpcsuperuser -lang jython
-f "c:/program files/ibm/tpc/tip/scripts/importLTPAKeys.py"
"c:/share/ltpaKeys_serv1" ltpa123
```

Note: Use forward slashes.

6. The LTPA keys in TPC_server1 and TPC_server2 are now synchronized. You can perform a successful single sign-on launch from TPC_server2 to the DS8000 R4.2. The DS8000 uses the same policy that was set up when you set up TPC_server1.

The same steps can be used to launch the same DS8000 from any number of Tivoli Storage Productivity Center servers. This is not a high-availability setup because the policy in DS8000 is still pointing to only one Embedded Security Service, which is that of TPC_server1

Note: Use forward slashes.

Setting up dual Tivoli Storage Productivity Center servers for high availability

This section describes how to set up dual Tivoli Storage Productivity Center servers for high availability of LDAP Authentication and single sign-on.

Procedure

Follow these steps:

1. Configure one Tivoli Storage Productivity Center server as described in “Configuring Tivoli Storage Productivity Center for DS8000 LDAP authentication” on page 322. In this example, this server is called TPC_server1.
2. Install a second Tivoli Storage Productivity Center server with the same LDAP information as the first server. In this example, this server is called TPC_server2.
3. Open a command prompt window and navigate to the following directory:

TIP_install_directory/bin

- On Windows systems, the default installation path for Tivoli Integrated Portal is C:\Program Files\IBM\Tivoli\TIP. The default installation path for Tivoli Storage Productivity Center is C:\Program Files\IBM\TPC.
 - On Linux or AIX systems, the default installation path for Tivoli Integrated Portal is /opt/IBM/Tivoli/tip. The default installation path for Tivoli Storage Productivity Center is /opt/IBM/TPC.
4. On TPC_Server2, run the following WebSphere command to export the LTPA keys from TPC_server1 into a file on TPC_server2:

```
wsadmin -user TIP_admin_ID -password TIP_admin_password
-lang jython
-port TIP_SOAP_port
-host TPC_server1_hostname_or_IP_address
-f "TIP_install_directory_on_TPC_server2/tip/scripts/exportLTPAKeys.py"
"LTPA_keys_file_name" LTPA_keys_password
```

Here is an example on a Windows system using the default Tivoli Storage Productivity Center installation file path:

```
C:\Program Files\IBM\tipv2\bin>wsadmin -user tpcsuperuser
-password tpcsuperuserpassword
-lang jython
-port 9571
-host 9.54.91.40
-f "c:/Program Files/IBM/TPC/tip/scripts/exportLTPAKeys.py"
"c:/share/ltpakeys_serv1" ltpa123
```

This step creates a file named ltpakeys_serv1 which contains the LTPA keys of TPC_server1. The LTPA keys are imported into TPC_server2.

Note: Use forward slashes with the **-f** parameter.

5. In the same command window, run the following WebSphere command to import the LTPA keys into Tivoli Integrated Portal:

```
wsadmin -user TIP_admin_ID -password TIP_admin_password
-lang jython
-f "TIP_install_directory_on_TPC_server2/tip/scripts/importTPAKeys.py"
"LTPA_keys_file_name" LTPA_keys_password
```

Here is an example on a Windows system using the default Tivoli Storage Productivity Center installation file path:

```
C:\Program Files\IBM\tipv2\bin>wsadmin -user tpcsuperuser
-password tpcsuperuserpassword
-lang jython
-f "c:/Program Files/IBM/TPC/tip/scripts/importLTPAKeys.py"
"c:/share/ltpakeys_serv1" ltpa123
```

Note: Use forward slashes with the **-f** parameter.

6. Synchronize the second authentication service (on TPC_server2) with the correct LTPA keys:

```
C:\Program Files\IBM\tipv2\bin>wsadmin -user tpcsuperuser
-password tpcsuperuserpassword
-lang jython
-c "AdminTask.importESSLTPAKeys
('[-pathname c:/share/ltpakeys_serv1 -password ltpa123]')
```

Restart the second Tivoli Integrated Portal server (on TPC_server2).

The LTPA keys in TPC_server1 and TPC_server2 are now synchronized.

7. If you are using the Java client, add the SSL certificates for all servers to the truststore file:
 - a. Log on to Tivoli Integrated Portal on TPC_server1 and extract the certificate. In this example, the certificate is named cert1.cer.
 - b. Log on to Tivoli Integrated Portal on TPC_server2 and extract the certificate. In this example, the certificate is named cert2.cer.
 - c. On TPC_server1, take the cert1.cer and cert2.cer certificates at one location and use the Java keytool command to create a truststore.
 - d. Go to C:\Program Files\IBM\tipv2\java\bin and add these two certificates:


```
keytool -import -alias TPCServer1
-file c:\cert1.cer
-keystore c:\ess.truststore.jks
-storetype jks
-storepass password

keytool -import -alias TPCServer2
-file c:\cert2.cer
-keystore c:\ess.truststore.jks
-storetype jks
-storepass password
```
 - e. Verify that the two certificates exist in the keystore by running this command:


```
keytool -list
-keystore c:\ess.truststore.jks
-storepass password
```

This command must list the two aliases (TPCServer1 and TPCServer2) in the keystore.
 - f. Copy the ess.truststore.jks keystore to c:\Program Files\IBM\TPC\device\conf on TPC_server1 and TPC_server2.

Configuring Tivoli Storage Productivity Center and SAN Volume Controller or Storwize V7000 for single sign-on

To use one user ID and password to access multiple applications, configure Tivoli Storage Productivity Center and SAN Volume Controller or Storwize V7000 for single sign-on.

Overview

Configuring Tivoli Storage Productivity Center and SAN Volume Controller or Storwize V7000 for single sign-on involves these general steps:

Tivoli Storage Productivity Center

Know the web address for the Authentication Service and the HTTP basic authentication user name and password.

SAN Volume Controller or Storwize V7000

Configure the remote authentication service of the cluster.

Configuring Tivoli Storage Productivity Center for single sign-on

Before you begin, ensure that Tivoli Storage Productivity Center is configured with an LDAP repository.

To configure Tivoli Storage Productivity Center for single-sign on:

1. Extract the Secure Sockets Layer (SSL) certificate.

Remember: The certificate must be in the Privacy Enhanced Mail (PEM) format.

- On the Windows operating system, complete the following steps:
 - a. Go to `https://TIP_server_host:9569/ibm/console/logon.jsp` where *TIP_server_host* is the IP address or host name of the Tivoli Integrated Portal server.

On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have Tivoli Integrated Portal administrator authorization.
 - b. In the Tivoli Integrated Portal tree, click **Settings > WebSphere Administrative Console**.
 - c. In the Tivoli Integrated Portal topic pane, click **Launch WebSphere administrative console**.
 - d. In the WebSphere Administrative Console, click **Security > SSL certificate and key management > Key stores and certificates > NodeDefaultKeyStore > Personal Certificates**.
 - e. Select the default certificate and click **Extract**.
 - f. Enter the following information:
 - **Certificate file name**

Enter a file name for the extracted certificate and click **OK**. This file is automatically created in *TIP_install_directory/profiles/TIPProfile/etc/*. In Microsoft Windows, the default directory is `C:\TIP_install_directory\profiles\TIPProfile\etc\`.
 - **Data Type**

Select Base64-encoded ASCII data.
- Alternatively, on the Linux or UNIX operating systems, you can run the following Shell command to extract the SSL certificate from Tivoli Integrated Portal:

```
echo | openssl s_client -connect TIP_server_ip:9569 2>&1 \  
| sed -ne '/-BEGIN CERTIFICATE-/,/-END CERTIFICATE-/p'
```

where *TIP_server_ip* is the IP address of the Tivoli Integrated Portal server.

2. Determine the web address for the authentication service.

Remember:

- SAN Volume Controller or Storwize V7000 does not support Domain Name System (DNS) host names. IP addresses must be used.
- The default port for the unencrypted (HTTP) Tivoli Integrated Portal Authentication Service is 9568.
- The default port for the encrypted (HTTPS) Tivoli Integrated Portal Authentication Service is 9569.
- Web address for HTTP:

```
http://TIP_server_ip:port/TokenService/services/Trust
```

where *TIP_server_ip* is the IP address of the Tivoli Integrated Portal server.
- Web address for HTTPS:

```
https://TIP_server_ip:port/TokenService/services/Trust
```

where *TIP_server_ip* is the IP address of the Tivoli Integrated Portal server.

3. Determine the authentication service user ID and authentication service password to be used on the cluster for remote authentication. It can be any user ID in the LDAP repository that is also used by Tivoli Storage Productivity Center. Use this user ID and password when you configure SAN Volume Controller or Storwize V7000.

Configuring SAN Volume Controller or Storwize V7000 for remote authentication service

You can use either the SAN Volume Controller or Storwize V7000 Management GUI or the command-line interface (CLI) to configure the remote authentication of a cluster. To configure remote authentication service on SAN Volume Controller or Storwize V7000, you must have the Security Administrator role.

To use the SAN Volume Controller Console Version 5.x to configure remote authentication of a cluster:

1. Log on to the SAN Volume Controller Management GUI and select **Launch the SAN Volume Controller Console** for the particular cluster.
2. From the main menu, select **Manage Authentication > Remote Authentication**.
3. Check **Enabled** to enable remote authentication.
4. Enter the following information:
 - **Remote service user name**
 - **Remote service password**
 - **Protocol:**
 - HTTP
 - HTTPS
 - **Remote service web address**
5. Click **OK**.
6. From the main menu, click **Manage Authentication > User Groups**.
7. Select **Create a Group** and then click **Go**.
8. Enter a **Name**.

The name of the group you create in the SAN Volume Controller Management GUI must be identical to the name of a group in the LDAP repository, which is also used by Tivoli Storage Productivity Center. The members of this LDAP group must be the LDAP users who use single sign-on between Tivoli Storage Productivity Center and the SAN Volume Controller cluster.

9. Select a **Role**.
10. Select the **Enable this user group to be visible to the remote authentication service** check box.
11. Click **OK**.
12. Repeat steps 6 - 11 to create another group in the SAN Volume Controller Management GUI that corresponds to another group in the LDAP repository that is also used by Tivoli Storage Productivity Center.

To use the SAN Volume Controller Management GUI Version 6.1 or later or Storwize V7000 Management GUI to configure remote authentication of a cluster, follow these steps:

1. In a browser, enter the IP address for the Management GUI:
`https://cluster_domain_name or IP_address`
2. Log on to the Management GUI with your user name and password.

3. In the navigation portion of the Management GUI, click **Settings > Directory Services**.
4. On the Directory Services panel click **Global Actions > Enable Remote Authentication**.
5. After successfully enabling remote authentication, click **Global Actions > Configure Remote Authentication**.
6. Select **IBM Tivoli Integrated Portal** and click **Next**.
7. Enter the following information:
 - **Remote Service Credentials Name**
 - **Remote Service Credentials Password**
 - **Web address for the remote authentication service**
8. Click **OK**.
9. If you are using the encrypted authentication service (HTTPS), you see a message that the SSL certificate is automatically retrieved from the web address.
10. Click **OK**.
11. In the navigation portion of the Management GUI, click **Access > Users**.
12. Click **New User Group**.
13. Enter a **Group Name**.

The name of the group you create in the Management GUI must be identical to the name of a group in the LDAP repository that is also used by Tivoli Storage Productivity Center. The members of this LDAP group must be the LDAP users that use single sign-on between Tivoli Storage Productivity Center and the cluster.
14. Select a **Role**.
15. Under **Remote Authentication (IBM Tivoli Integrated Portal)**, select the **Enable for this group** check box.
16. Click **Create**.
17. Repeat steps 11-16 to create another group in the Management GUI that corresponds to another group in the LDAP repository that is also used by Tivoli Storage Productivity Center.

Configuring and controlling the Tivoli Storage Productivity Center Monitoring Agent

You might want to change the parameters or start and stop the Tivoli Storage Productivity Center Monitoring Agent. This topic describes how to configure the Tivoli Storage Productivity Center Monitoring Agent to change the Tivoli Enterprise Monitoring Services connection, log path, or authentication parameters. This topic also describes how to start and stop the Tivoli Storage Productivity Center Monitoring Agent.

Configuring the Tivoli Storage Productivity Center Monitoring Agent on Windows

This topic describes how to configure the Tivoli Storage Productivity Center Monitoring agent on Windows.

To configure the Tivoli Storage Productivity Center Monitoring agent on Windows, complete the following steps:

1. Log on to the system with administrator authority on Windows.

2. Open the Tivoli Enterprise Monitoring Services Console. Click **Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services**.

Note: Depending on how you have your Tivoli Monitoring Services components distributed across your enterprise, you will see different types of Tivoli Enterprise Monitoring Services components displayed.

3. On the Manage Tivoli Enterprise Monitoring Services window, select **Monitoring Agent for TPC**. Right-click **Monitoring Agent for TPC** and click **Reconfigure**.
4. For information about the different parameters to change, see “Installing and configuring the Tivoli Storage Productivity Center Monitoring Agent on Windows” on page 173.

Configuring the Tivoli Storage Productivity Center Monitoring Agent on AIX or Linux

This topic describes how to configure the Tivoli Storage Productivity Center Monitoring Agent on AIX or Linux.

To configure the Tivoli Storage Productivity Center Monitoring agent on AIX or Linux, complete the following steps:

1. Log on to the system with the root user ID or a user ID that has root access rights.
2. Open the Tivoli Enterprise Monitoring Services Console. Open a terminal session window and go to the Tivoli Monitoring Services bin installation directory. The default directory is:

```
/opt/IBM/ITM/bin
```

3. Run the following command:

```
./itmcmd manage
```

Note: Depending on how you have your Tivoli Monitoring Services components distributed across your enterprise, you will see different types of Tivoli Enterprise Monitoring Services components displayed.

4. On the Manage Tivoli Enterprise Monitoring Services window, select **Monitoring Agent for TPC**. Right-click **Monitoring Agent for TPC** and click **Configure**.
5. For information about the different parameters to change, see “Installing and configuring the Tivoli Storage Productivity Center Monitoring Agent on AIX or Linux” on page 176.

Starting and stopping the Tivoli Storage Productivity Center Monitoring Agent

This topic describes how to start and stop the Tivoli Storage Productivity Center Monitoring Agent.

To start or stop the Tivoli Storage Productivity Center Monitoring agent on Windows, complete the following steps:

1. Open the Tivoli Enterprise Monitoring Services Console. Click **Start > All Programs > IBM Tivoli Monitoring > Manage Tivoli Monitoring Services**.
2. On the Manage Tivoli Enterprise Monitoring Services window, select **Monitoring Agent for TPC**. Right-click **Monitoring Agent for TPC**. Click **Start** to start the agent, **Stop** to stop the agent, or **Recycle** to recycle the agent.

3. Wait for the agent to start, stop, or recycle.

To start or stop the Tivoli Storage Productivity Center Monitoring agent on AIX or Linux, complete the following steps:

1. Open the Tivoli Enterprise Monitoring Services Console. Open a terminal session window and go to the Tivoli Monitoring Services bin installation directory. The default directory is:
`/opt/IBM/ITM/bin`
2. Run the following command:
`./itmcmd manage`
3. On the Manage Tivoli Enterprise Monitoring Services window, select **Monitoring Agent for TPC**. Right-click **Monitoring Agent for TPC**. Click **Start Service** to start the agent, **Stop Service** to stop the agent, or **Recycle Service** to recycle the agent.
4. Wait for the agent to start, stop, or recycle.

Installing and configuring the Tivoli Storage Productivity Center server with multiple NIC cards

If your Tivoli Storage Productivity Center server has multiple network interface cards (NIC), install the Tivoli Storage Productivity Center server using a fully qualified hostname that resolves to the IP address of NIC card you want to use. After you install the Tivoli Storage Productivity Center server, all incoming and outgoing communication are successfully handled.

Installing Tivoli Storage Productivity Center for a multiple network configuration

If the Tivoli Storage Productivity Center server you are installing has multiple NIC, and is configured to use multiple network addresses, ensure that you use the fully qualified hostname that resolves to the appropriate IP address during installation. You can either setup the HOSTS file or the DNS to resolve the fully qualified host names to appropriate IP addresses.

Outgoing communication initiated by the Tivoli Storage Productivity Center server

All the outgoing communication that is initiated by the Tivoli Storage Productivity Center server is not affected if the server is configured for a multiple network environment.

For example, if you have a Tivoli Storage Productivity Center server with two IP addresses: 10.10.10.11 and 9.9.9.10, and 10.10.10.11 is used during installation, all outgoing transmissions can be sent to the devices and agents in both networks.

The following list includes examples of outgoing communication that is initiated by the Tivoli Storage Productivity Center server:

Storage systems using native interfaces

Run probe, performance management, and provisioning jobs, and collect data events (SAN Volume Controller, Storwize V7000 Unified, Storwize V7000, and XIV Systems)

Switches (SNMP)

Run SNMP discovery and probe jobs

CIM agents

Run discovery, probe, and performance management jobs, and provisioning jobs

VMware vSphere or vCenter

Run discovery and probe jobs

Agents (Storage Resource agents)

Deploy agents, run probe, discovery, scan, and batch report jobs, and run scripts

Tivoli Storage Productivity Center servers

Run probe jobs

Incoming communication that is initiated by the devices, agents, and GUI

Incoming communication that is initiated by the devices or agents can work with only the IP address that is specified during the installation with the exception of DS8000 events.

For DS8000 events, the Tivoli Storage Productivity Center server must initiate and establish a socket connection directly with the Hardware Management Console (HMC) to receive events. The DS8000 HMC uses that socket connection to send events. As long as the Tivoli Storage Productivity Center server can initiate the communication to the HMC, DS8000 events can be received.

Tivoli Storage Productivity Center informs devices and agents to initiate communication to the IP address provided during the installation. This example uses the IP address 10.10.10.11. However, depending on the communication, you might be able to change the IP address. For example, Tivoli Storage Productivity Center does not configure SAN switches to send SNMP traps to Tivoli Storage Productivity Center, so you can use either 9.9.9.10 or 10.10.10.11.

The following list includes examples of incoming communication that are initiated by the devices, agents, and the GUI:

DS8000 events

Events sent by the HMC to the Tivoli Storage Productivity Center server

SNMP trap notifications

SNMP traps sent from the switches and other devices

CIM indications

Indications sent by the CIM agents

Servers (agents)

Job results and registration

Tivoli Storage Productivity Center GUI

Any request.

CIM indications

A CIM indication is an event that occurs on a managed object, for example, the completion or failure of an operation. The CIM indications are managed by the CIM object manager. Tivoli Storage Productivity Center uses the CIM agents for the managed objects to gather information about the device.

Manually customize CIM indications on a Tivoli Storage Productivity Center system that has multiple IP addresses. To configure Tivoli Storage Productivity Center to receive CIM indications in an IPv4, IPv6, and dual stack (IPv4 and IPv6) environment, see “Configuring Tivoli Storage Productivity Center with multiple IP addresses” on page 314.

The manual customization task does not apply to storage devices that use the native interfaces.

Creating an SSH certificate for the root user ID

You can create a Secure Shell (SSH) certificate for authentication for the Virtual I/O Server. Follow the certificate-generation instructions. However, if you want to use Telnet to connect to the Virtual I/O Server using the padmin user ID, you must follow this procedure.

Procedure

To create an SSH certificate using the padmin user ID, follow these steps:

1. Telnet to the remote system using the padmin user ID.
2. Set up the AIX environment. Run the following command:
oem_setup_env
3. Change to the following /.ssh directory.
4. Enter **ssh-keygen**. Accept the default names (for example, id_rsa).
5. Enter the passphrase. Two new files are created:

id_rsa This is the private key.

id_rsa.pub

This is the public key.

6. Create an authorized_key file in the same location as the id_rsa.pub file. Enter the following command:

```
cat >> id_rsa.pub >> authorized_keys
```

The following example shows the command input and output (the commands are in bold):

```
# ssh-keygen

Generating public/private rsa key pair.
Enter file in which to save the key (//.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
You identification has been save in //.ssh/id_rsa.
Your public key has been save in //.ssh/id_rsa.pub.
The key fingerprint is:
xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx root@<server>

# cat id_rsa >> authorized_keys
# ls -l

-rw-r--r- 1 root system 1743 Oct 15 09:40 authorized_keys
-rw---- 1 root system 1743 Oct 15 09:39 id_rsa
-rw-r--r- 1 root system 399 Oct 15 09:39 id_rsa.pub
```

7. Copy the id_rsa (private key) to your server machine.

Note: You must copy the file in binary mode.

8. To connect to the remote system using the private key, enter the following information in the Remote Agent Machines window when you install the Storage Resource agent:
 - User
 - Certificate Location (c:\keys\id_rsa)
 - Passphrase

Configuring DB2, AIX, and Linux for IPv6-only environment

Use this information to configure DB2, AIX, and Linux for an IPv6-only environment.

Configuring the AIX system for IPv6 only

Note: For IPv6 support, the AIX operating system must have level TL 5300–06 installed.

Follow these steps:

1. Obtain the most recent versions of **openssh** and **openssl** packages for AIX and install them. Some older version of **openssh** does not work in an IPv6-only environment.
2. Change **sshd** (Secure Shell Daemon) on AIX system to accept IPv6 connections.
 - a. In the `/etc/ssh/sshd_config` file, uncomment the line "ListenAddress:".
 - b. Restart **sshd** with the following commands:


```
stopsrc -g ssh
startsrc -g ssh
```
 - c. From another IPv6 system, verify that you contact AIX over IPv6 (using `ssh`).
3. In SMIT, set the IPv4 address to 0.0.0.0 for all interfaces. Save the file.
4. Edit the `/etc/resolv.conf` file to use IPv6 DNS server or servers.

Configuring DB2 on AIX for IPv6 systems

To get DB2 on AIX to work on IPv6 systems, follow these steps:

1. Identify the host name used by DB2 in the `db2nodes.cfg` file:


```
# cat ~db2inst1/sqllib/db2nodes.cfg
0 myhost 0
#
```
2. Edit the `/etc/hosts` file and make sure that the host name found in the `db2nodes.cfg` file resolves to an IPv6 address. Use the `vi` editor to verify that the host name is not on any line with an IPv4 address. In particular, ensure that the host name is not listed as an alias for the IPv4 loopback address 127.0.0.1.


```
# vi /etc/hosts
127.0.0.1 loopback localhost
::1 localhost
2001:db8:0:0:209:6bff:fe09:63fa myhost.mydomain myhost
```
3. Stop DB2 and set DB2 to use IPv6 addressing. Restart DB2.
 - a. Source the DB2 profile:


```
. ~db2inst1/sqllib/db2profile
```
 - b. Stop DB2:


```
db2stop
```
 - c. Configure DB2 to use IPv6.

```
db2set
```

An example of the output is: DB2FCMCOMM=TCPIP6.

- d. Start DB2.

```
db2start
```

In some installations, the AIX server does not have a graphical console attached to the server. In this situation, you can select another system with an X11 server to display the IBM Tivoli Storage Productivity Center installation and Tivoli Storage Productivity Center application. The X11 server must have IPv6 configured and an SSH client installed. Open an SSH connection from a shell on the X11 server desktop with the `-X` option to permit forwarding of X11 applications from the remote AIX server. Start the Tivoli Storage Productivity Center installation program or application from the SSH shell.

```
ssh -X my_IPv6_host  
/opt/IBM/TPC/gui/TPCD.sh
```

Configuring DB2 on Linux for IPv6-only systems

To get DB2 on Red Hat Enterprise Linux Version 5 and 6 systems to work in an IPv6-only environment, follow these steps:

1. Install DB2 in dual-stack configuration.
2. Stop DB2 and set DB2 to use IPv6 addressing:
 - a. As the root user from the UNIX command line, run this command:

```
su - db2inst1
```
 - b. Stop DB2 by running this command:

```
db2stop
```
 - c. Configure DB2 to use IPv6 by running this command:

```
db2set
```

An example of the output is: DB2FCMCOMM=TCPIP6.

Note: The host name in the `db2nodes.cfg` file resolves to an IPv6 address. This action can require you to change the domain or search directive in the `/etc/resolv.conf` file to specify a domain in which the host name can resolve to IPv6. You can also edit the `/etc/hosts` file so that the host name resolves to an IPv6 address.

- d. Start DB2 by running this command:

```
db2start
```

Changing the timeout value for the web-based GUI session

You can change the timeout value for the web-based GUI session through Tivoli Integrated Portal so that you are not logged off the session before you complete your tasks.

1. Open an Internet Explorer or Firefox web browser.
2. Enter the following information in the address bar:

```
https://hostname:port/ibm/console/login.jsp
```

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. For more information about Tivoli Integrated Portal port numbers, see “Tivoli Integrated Portal port numbers” on page 35.

3. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password.
4. In the Tivoli Integrated Portal navigation tree, click **Settings > WebSphere Administrative Console**.
5. In the Tivoli Integrated Portal content pane, click **Launch WebSphere administrative console**.
6. In the WebSphere Administrative Console navigation tree, click **Security > Global Security**.
7. Click the link labeled **LTPA** under the Authentication section to change the timeout value.
8. Restart the Tivoli Integrated Portal server.

Chapter 4. Upgrading and migrating Tivoli Storage Productivity Center

You can upgrade any Tivoli Storage Productivity Center Version 4 release to Tivoli Storage Productivity Center Version 5.

The following section provides information for a successful upgrade:

- Tivoli Storage Productivity Center Version 5 does not use the Data agent or Fabric agent, so you must migrate the Data agents and Fabric agents to Storage Resource agents.
- You must decide whether you want to migrate performance and configuration history data. Depending on how much data you have, it can take from a one to several hours or more to migrate the data. If you do not migrate the data, this information is deleted from the database repository.
- If the upgrade process fails and a Tivoli Storage Productivity Center component cannot be upgraded, the upgrade process ends and the status of the component is shown as failed. When you restart the upgrade process, the installation resumes upgrading the failed component.
- The Tivoli Storage Productivity Center Version 4 ports are reused in Tivoli Storage Productivity Center Version 5.

Complete the following tasks before you upgrade Tivoli Storage Productivity Center:

- Ensure that your system meets the hardware and software requirements for upgrading. For more information, see “Hardware requirements” on page 1 and “Software requirements” on page 4.
- Back up your Tivoli Storage Productivity Center environment, including the database repository. For more information, see “Backups” on page 349
- Migrate any Data agents or Fabric agents to a Storage Resource agent. For more information, see “Migrating Data agents and Fabric agents to Storage Resource agents” on page 365.

Restriction: To migrate Data agents or Fabric agents before the upgrade, use the Tivoli Storage Productivity Center Version 4 GUI or use the command-line interface. To migrate Data agents or Fabric agents after the upgrade, use the Tivoli Storage Productivity Center 5 GUI or the migration command in the command-line interface.

- You must manually migrate BIRT reports. These reports must be exported from Tivoli Integrated Portal 1.1 before you migrate Tivoli Storage Productivity Center. For more information about migrating BIRT reports, see “Migrating BIRT reports” on page 369.

Upgrading and migrating Tivoli Integrated Portal

If the Tivoli Storage Productivity Center installation wizard does not detect Tivoli Integrated Portal Version 2.2, a new instance of Tivoli Integrated Portal Version 2.2 is installed. Tivoli Common Reporting Version 2.1.1, Extended Security Services, and the web-based GUI components are also installed. You must enter a Tivoli Common Reporting port.

Note: If you are upgrading on the Linux operating system, you do not need to enter a Tivoli Common Reporting port.

If the installation wizard detects Tivoli Common Reporting Version 2.1.1, and Tivoli Integrated Portal Version 2.2 on the system, those instances are used by Tivoli Storage Productivity Center installation. When you use an existing instance of Tivoli Integrated Portal, you must enter the Tivoli Integrated Portal administrative user name and password that was used when you installed that instance. You cannot enter a new user name and password for that instance.

The installation program migrates the user authentication settings including the LDAP configuration. The installation program does not migrate any data from Tivoli Integrated Portal 1.1 to Tivoli Integrated Portal 2.2.

Uninstalling Tivoli Integrated Portal

If you uninstall Tivoli Integrated Portal Version 1.1, and an instance of Tivoli Integrated Portal Version 2.2 does not exist, the ports, administrative user name, and password that were used in Version 1.1 are reused for the Version 2.2 instance.

If you uninstall Tivoli Integrated Portal Version 1.1, and an instance of Tivoli Integrated Portal Version 2.2 exists, you must change the values to the Tivoli Integrated Portal administrative user name and password associated with the Version 2.2 instance on the Tivoli Integrated Portal and Tivoli Common Reporting: Specify Server Information page. The installation location and port fields must be changed to contain the values associated with the Version 2.2 instance.

If you do not uninstall Tivoli Integrated Portal Version 1.1, the Tivoli Integrated Portal and Tivoli Common Reporting: Specify Server Information page is displayed. You can enter information for Tivoli Integrated Portal Version 2.2. The Tivoli Integrated Portal Version 1.1 server must be running before Tivoli Storage Productivity Center can uninstall the product.

Restriction: You can uninstall Tivoli Integrated Portal Version 1.1 only if it is not shared with another product. If it is shared with another product, do not uninstall it.

Upgrading Tivoli Storage Productivity Center components

In a multiple-servers environment, you must upgrade the Tivoli Storage Productivity Center components in the order listed in this table:

Table 40. Upgrading Tivoli Storage Productivity Center

Server A includes...	Server B includes...	Upgrade in this order...
Data server, Device server, replication server, database repository, web-based GUI, stand-alone GUI, CLI, Storage Resource agent, Tivoli Integrated Portal, Tivoli Common Reporting (single-server environment)	Not applicable	Not applicable

Table 40. Upgrading Tivoli Storage Productivity Center (continued)

Server A includes...	Server B includes...	Upgrade in this order...
Data server, Device server, replication server, web-based GUI, stand-alone GUI, CLI, Storage Resource agent, Tivoli Integrated Portal, Tivoli Common Reporting (multiple-server environment)	Database repository	<ol style="list-style-type: none"> 1. Run the Tivoli Storage Productivity Center upgrade program on Server B to upgrade the database repository. 2. Run the Tivoli Storage Productivity Center upgrade program on Server A to upgrade the remainder of the Tivoli Storage Productivity Center components.
Data server, Device server, replication server, database repository, stand-alone GUI, CLI, Storage Resource agent, (multiple-server environment)	Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI	<ol style="list-style-type: none"> 1. Run the Tivoli Storage Productivity Center upgrade program on Server A to upgrade all the components except for Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI. 2. Run the Tivoli Storage Productivity Center upgrade program on Server B to upgrade Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI.

Choosing the upgrade method

You can upgrade Tivoli Storage Productivity Center by using the upgrade wizard or the command line in silent mode. In silent mode, a command is provided with the values in a response file.

Important: On UNIX and Linux operating systems, you must have an X Window system program to display the installation wizard.

A silent-mode upgrade is useful if your system is running from a terminal that cannot display graphics.

Related tasks:

“Upgrading Tivoli Storage Productivity Center by using the wizard” on page 351
The upgrade wizard is one method of upgrading Tivoli Storage Productivity Center.

“Upgrading Tivoli Storage Productivity Center by using silent mode” on page 352
Silent-mode upgrade is one method of upgrading Tivoli Storage Productivity Center. This installation method is useful if your system is running from a terminal that cannot display graphics.

Upgrading Tivoli Storage Productivity Center for Replication

With Tivoli Storage Productivity Center for Replication version 4.2 or later, an embedded database is used for the database repository. DB2 is no longer supported as the database repository for Tivoli Storage Productivity Center for Replication.

As a part of the upgrade process, an embedded repository is created and any data from an existing DB2 Tivoli Storage Productivity Center for Replication database is copied to the new database repository. This process is automatic and does not require any input or action from the user.

For a typical upgrade procedure, follow this general procedure:

1. Issue the takeover command to the standby server and then upgrade the standby server.

Note: This action makes both the Replication servers active.

2. Wait for the standby server to complete installation and start the standby server. .
3. While upgrading the active server, avoid making any configuration changes to the sessions.
4. Upgrade the active server.
5. If no changes have been made to the configuration while the active server is being upgraded, reestablish the High Availability function from the active server to the standby server. If configuration changes were made to the standby server, synchronize the High Availability function from the standby server to the active server. Then perform a takeover operation and reestablish the High Availability function from the active server to the standby server.

During the initial synchronization, the current information in the database is saved and held until the synchronization is complete. If an error occurs during this process, the server database is restored to its original state before the synchronization process began. If an error occurs during the synchronization process which causes the status to be in the disconnected or inconsistent state, you can reconnect to a synchronized state.

Fix packs and patches

All fix packs and patches use the upgrade procedure to upgrade Tivoli Storage Productivity Center.

You must have a valid Tivoli Storage Productivity Center license to use the upgrade procedure. For information about fix packs and patches, go to <http://www.ibm.com/support/docview.wss?&uid=swg21320822>.

Storage Subsystem Credential Migration tool

Use this tool if you have DS8000, XIV, or SAN Volume Controller storage systems and are upgrading Tivoli Storage Productivity Center Version 4.1x to Version 5.1. These storage systems use the native storage system interface instead of the CIM agents.

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

- Use the migration tool from the Tivoli Storage Productivity Center DVD. Use the DVD or download the migration tool: `MigrateUserInfo.bat` (for Windows) or `MigrateUserInfo.sh` (for UNIX or Linux). Run the migration tool from the directory `UserMigrationTool`.
- Use the Tivoli Storage Productivity Center GUI to migrate the credentials (see “Update Storage Subsystem Credentials” on page 297)
- Run the Tivoli Storage Productivity Center upgrade installation program and specify that you want to use the migration tool. Run the migration tool during the Tivoli Storage Productivity Center upgrade procedure.

Note: For Windows, the credential migration tool has a dependency on `MSVCR90.dll` if you run the tool from a command window. If you get an error exception when running this tool, check to see that your system has this DLL file. If you use the tool from the installation wizard, the wizard automatically loads the required DLL files.

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

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

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

Note: After a successful migration of DS8000, XIV, or SAN Volume Controller credentials, the CIMOMs that manage these devices previously must be removed from the Tivoli Storage Productivity Center configuration. The migration process tries to determine and remove these CIMOMs during migration automatically, but if there are CIMOMs remaining, they must be deleted manually. For DS CIMOMs managing DS6000 or ESS devices and additional DS8000 devices, the DS8000 devices must be removed from the CIMOM if possible.

Upgrading DB2 from DB2 9.1 or DB2 9.5 to DB2 9.7

If you are upgrading DB2 from DB2 9.1 or DB2 9.5 to DB2 9.7, follow this general procedure.

For information about upgrading to DB2 9.7, go to this website:
<http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp>.

For the general procedure to upgrade the DB2 database repository to DB2 9.7, follow these steps:

1. Stop the Tivoli Storage Productivity Center services and Agent Manager. For information about stopping the services, go to the Tivoli Storage Productivity Center information center. Search for **Starting and stopping the Tivoli Storage Productivity Center services**.
2. Check to see if your databases are ready for migration:
 - a. On DB2 9.1 or DB2 9.5, disconnect all users, stop the back-end process, and stop your existing instances by running the following commands:


```
db2stop force (disconnects all users and stops the instance)
db2 terminate (terminates back-end process)
```
 - b. Verify that the local databases that are owned by the current instance are ready to be upgraded and generate a log file. Run the **db2ckupgrade** command.

Note: For Windows, click **Start > Programs > IBM DB2 > Command Line Tools > Command Window**.

Go to the directory for the **db2ckupgrade** command:

```
cd DB2_9.7_installer_kit\ESE\image\db2\Windows\utilities
```

Run this command:

```
db2ckupgrade database_name -l db2ckupgrade.log -u user_name
-p password
```

where *database_name* is the name of the database and *db2ckupgrade.log* is the name of the log file created in the current directory. The log file includes error and warning messages. Each time you run this command, it overwrites the existing log file. You can rename the log file to avoid losing the information from a previous run. You must correct these errors before you upgrade.

- c. Check the log file. For a successful run, the output provides this message:


```
db2ckupgrade was successful. Database(s) can be upgraded.
```
3. Back up the database. Run the following DB2 command from a command prompt window:


```
DB2 BACKUP DATABASE TPCDB USER user_name USING password
TO path_to_backup_folder
```

where *user_name* is the user ID, *password* is the password, and *path_to_backup_folder* is the directory path to the backup folder. Test the image integrity by running this command:

```
DB2CKBKP *.001
```

4. Run the DB2 extended security command:


```
db2set DB2_EXTSECURITY=NO
```
5. Install DB2 9.7. For information about installing DB2 9.7, see “Installing DB2” on page 180. Select the **Install New** option on the Install a Product panel. Install all DB2 add-on products that you installed on your previous DB2 version.
6. Upgrade the DB2 instances.
 - a. On DB2 9.1 or DB2 9.5, disconnect all users, stop the back-end process, and stop your existing instances by running the following commands:


```
db2stop force (disconnects all users and stops the instance)
db2 terminate (terminates back-end process)
```
 - b. Log on to the DB2 9.1 or DB2 9.5 server with root authority on UNIX or local administrator authority on Windows.

- c. Upgrade your existing instances by running the **db2iupgrade** command from the target DB2 version 9.7 copy location.

For Windows, run this command:

```
"%DB2PATH%" \bin\db2iupgrade InstName /u:user,password
```

Where DB2PATH is set to the location you specified during DB2 Version 9.7 installation, *user,password* are the user name and password under which the DB2 service runs, and *InstName* is the name of the instance.

For UNIX or Linux, run this command:

```
$DB2DIR/instance/db2iupgrade [ -u fencedID ] InstName
```

Where DB2DIR is set to the location you specified during DB2 Version 9.7 installation, *fencedID* is the user name under which the fenced user-defined functions (UDFs) and stored procedures run, and *InstName* is the login name of the instance owner. This example upgrades the instance to the highest level for DB2 database product that you installed.

- d. Check the following configuration file for Windows (2008):

```
C:\ProgramData\IBM\DB2\DB2COPY1\DB2\db2node.cfg
```

Remove the extra line. The line should look like this:

```
0 beesspc BEESPC 0
```

Where beesspc and BEESPC is the name of the Tivoli Storage Productivity Center server.

- e. Log on to the DB2 9.7 server as a user with sufficient authority to start your instance.
- f. Restart your instance by running the **db2start** command.
- g. Verify that your instance is running on DB2 9.7 by running the **db2level** command. The output provides a string like:

```
DB2 v9.7.x.x
```

Where x is a digit number.

7. Upgrade the DB2 Administration Server.

- a. Log on to the DB2 9.7 server with root authority on UNIX or Linux or local administrator authority on Windows.
- b. Upgrade your existing DB2 Administration Server by running the **dasmigr** command.

For Windows, run this command:

```
%DB2PATH%\bin\dasmigr
```

Where DB2PATH indicates the location that you specified during DB2 Version 9.7 installation.

For UNIX or Linux, run this command:

```
$DB2DIR/instance/dasmigr
```

Where DB2DIR indicates the location that you specified during DB2 Version 9.7 installation.

8. Upgrade the databases.

- a. Log on to the DB2 9.7 server as the instance owner or a user with SYSADM authority.

- b. Migrate the TPCDB database. Run this command from the DB2 command prompt window:

```
db2 migrate db TPCDB
```

- c. From a DB2 command prompt window, run this command:

```
db2 update dbm cfg using JDK_PATH 'C:\Program Files\IBM\SQLLIB\java\jdk'
```

- d. Verify that your database upgrade is successful. Connect to the upgraded databases and run a small query:

```
db2 connect to database_name
```

Database Connection Information

```
Database server      = DB2/INT 9.7.0
SQL authorization ID = ADMINISTRATOR
Local database alias = database_name
```

```
db2 "select * from syscat.dbauth"
```

- e. For SSPC only: Open the DB2 - DB2COPY1 - DB2-0 Properties window. Click the **Log On** tab. Check to see if the "Log On As" service for DB2 - DB2COPY1 - DB2-0 is .\administrator.

Set this DB2 service to Automatic. Open the Services window and find the DB2 - DB2COPY1 - DB2-0 service. Right-click this service and click **Automatic**.

9. Start the Tivoli Storage Productivity Center services and Agent Manager (if the Agent Manager is installed). For information about starting the services, go to the Tivoli Storage Productivity Center. Search for **Starting and stopping the Tivoli Storage Productivity Center services**.
10. Upgrade Tivoli Storage Productivity Center.

Related reference:

Chapter 4, "Upgrading and migrating Tivoli Storage Productivity Center," on page 341

You can upgrade any Tivoli Storage Productivity Center Version 4 release to Tivoli Storage Productivity Center Version 5.

Preparing for an upgrade

To upgrade Tivoli Storage Productivity Center, use the same starting procedure as the installation program. Depending on the components you install, you can see different pages.

Before you upgrade, you must review the following information:

- "Hardware requirements" on page 1
- "Software requirements" on page 4

For a complete list of installation images, see "Tivoli Storage Productivity Center installation" on page 135.

Note:

- If you are deploying Storage Resource agent on UNIX or Linux systems, you must use root as the user ID.
- You can no longer install the Storage Resource agent from the disk 1 image on computers that do not have a Data server installed. To install a Storage Resource agent on a computer where the Data server is not installed, you must use the separate Storage Resource agent electronic image or DVD.

- After you upgrade Tivoli Storage Productivity Center, all trace settings are reset to the default values. If you changed the trace settings, you must reset the trace settings after an upgrade.

Important: This procedure is disruptive to the Tivoli Storage Productivity Center for Replication services. For Tivoli Storage Productivity Center for Replication, you can complete a takeover to a standby server to ensure that consistency can be maintained for any active servers.

Backups

Before you can migrate, you must back up the entire Tivoli Storage Productivity Center system.

The process involves backing up the following parts:

- Tivoli Storage Productivity Center database by using DB2 backup utilities.
For more information about backing up your database, see http://pic.dhe.ibm.com/infocenter/tivihelp/v59r1/index.jsp?topic=%2Fcom.ibm.tpc_V51.doc%2Ffqz0_t_backing_up_tpc_database.html
- The entire Tivoli Storage Productivity Center server by using software such as IBM Tivoli Storage Manager.

Remember: The frequency of backups depends on the data protection policies of your company, the Tivoli Storage Productivity Center work load, and the size of the Tivoli Storage Productivity Center environment. Your backup plan must include documentation with the user names and passwords for DB2, Tivoli Storage Productivity Center, and Tivoli Integrated Portal.

If you experience problems during the migration, use these backups to restore your original Tivoli Storage Productivity Center system and try to migrate again. For more information about restoring the Tivoli Storage Productivity Center database, see <http://pic.dhe.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=%2Fcom.ibm.db2.luw.admin.ha.doc%2Fdoc%2Fc0006237.html>

Upgrade response file

You do not need an upgrade response file to upgrade Tivoli Storage Productivity Center in silent mode from Version 5 to another Version 5 fix pack. To accept the license agreement, you must specify the value of the `LICENSE_ACCEPTED` variable to `true`. To upgrade from Version 5 to a Version 5 fix pack, run the command `setup.bat | bin -i silent -DLICENSE_ACCEPTED=true`.

You must use the `silent_Upgrade4x.properties` file to upgrade Tivoli Storage Productivity Center in silent mode from Tivoli Storage Productivity Center Version 4 to Version 5. When you maintain a current version of this response file, you minimize upgrade errors.

Upgrade response file parameters

The response file contains the following parameters:

CHOSEN_INSTALL_TYPE="4.x Upgrade"

Specifies the upgrade type. You can upgrade any Tivoli Storage Productivity Center version 4 release to version 5.

LICENSE_ACCEPTED=false

Specify whether the user accepts the Tivoli Storage Productivity Center license agreement. The options are:

true

The user accepts all the terms and conditions of the Tivoli Storage Productivity Center license agreement.

false or any other value

The user does not accept the Tivoli Storage Productivity Center license agreement. The installation program will exit.

varMigrateCASAgentsLater=option

Specify when to migrate the Data agents and Fabric agents. The options are:

false

You have migrated all the Data agents and Fabric agents to Storage Resource agents before this upgrade. This option is the default.

true

You plan to migrate the Data agents and Fabric agents to Storage Resource agents after this upgrade.

varMigratePerfAndConfigHistoryData=option

Specify whether you want to migrate performance and configuration history data.

If you generated performance and configuration history data in Tivoli Storage Productivity Center Version 4 and want to save this information, migrate the data to Version 5. This migration can take from a few minutes up to a few hours or more depending on the amount of data to migrate.

If you do not migrate this data, the performance and configuration history data is deleted from the database repository. The options are:

true

Migrate the performance and configuration history data. This option is the default.

false

Delete the performance and configuration history data from the database repository.

varPreserveTIP1x=value

Specify whether to uninstall Tivoli Integrated Portal 1.1 by selecting one of the following values:

0 Uninstall Tivoli Integrated Portal 1.1. This option is the default.

1 Do not uninstall Tivoli Integrated Portal 1.1.

varTIPAdmUsr=user_name

If you use an existing Tivoli Integrated Portal 2.2 instance, you must provide a user name that has administrative privileges and enter the Tivoli Integrated Portal administrative user name and password that you used when you installed that instance. You cannot enter a new user name and password for that instance.

varTIPAdmPW=password

Specifies the password for the Tivoli Integrated Portal user name.

varReuseExistingTIP=value

Specifies whether to install a new instance of Tivoli Integrated Portal or reuse an existing instance of Tivoli Integrated Portal. Select one of these values:

0 Install a new instance of Tivoli Integrated Portal 2.2. This value is the default.

- 1 Reuse an existing instance of Tivoli Integrated Portal 2.2.

varTIPPath=*path*

Specifies the directory where Tivoli Integrated Portal is installed. If you reuse an existing instance of Tivoli Integrated Portal, provide the directory where Tivoli Integrated Portal is installed.

For Windows directories, use double backslashes for the path. For example:

```
varTIPPath=c:\\TIP\\version2.2\\
```

varTIPPortRangeSP=*path*

Specifies the first port in a range of ports for Tivoli Integrated Portal. The default port range is 9568 - 9581.

This parameter overrides **varTPCPortRangeSP**.

varTCRPort=*port*

Specifies the Tivoli Common Reporting port. The default is 9581.

This parameter overrides **varTPCServersPortRangeSP** and **varTIPPortRangeSP**.

Upgrading Tivoli Storage Productivity Center in a single-server environment

You can upgrade Tivoli Storage Productivity Center by using the upgrade wizard or silent mode.

In a single-server environment, to upgrade Tivoli Storage Productivity Center, see the following topics:

- “Upgrading Tivoli Storage Productivity Center by using the wizard”
- “Upgrading Tivoli Storage Productivity Center by using silent mode” on page 352

Migrating from an older version of Tivoli Storage Productivity Center to Tivoli Storage Productivity Center 5.1

To migrate from Version 3.x to Tivoli Storage Productivity Center 5.1, you must first migrate from Version 3.x to Version 4.x. For more information about migrating from Version 3.x to Version 4.x, see http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=%2Fcom.ibm.tpc_V4221.doc%2Ffqz0_r_migrating_tpc.html. You can use your Version 3.x license to migrate from Tivoli Storage Productivity Center Version 3.x to Tivoli Storage Productivity Center 4.x.

Initially, Tivoli Storage Productivity Center 5.1 will be available through IBM Passport Advantage (download). If your license maintenance is current, you can download Tivoli Storage Productivity Center 5.1 from IBM Passport Advantage or access the physical media.

Upgrading Tivoli Storage Productivity Center by using the wizard

The upgrade wizard is one method of upgrading Tivoli Storage Productivity Center.

Before you start the Tivoli Storage Productivity Center upgrade program on the AIX or Linux operating systems, you must source the DB2 profile, `db2profile`, for the instance owner of the DB2 database. The following example sources the DB2 profile:

```
. /home/db2inst1/sqllib/db2profile
```

To upgrade Tivoli Storage Productivity Center in a single-server environment:

1. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
2. Back up Tivoli Storage Productivity Center.
3. Start the Tivoli Storage Productivity Center upgrade program by running `setup.bin ./setup/bin`.
4. Select the language for your installation and click **OK**. The language that you select is used to install the product.
5. On the Choose Installation Location and Type page, select **Single server**.

Note: Tivoli Storage Productivity Center detects the currently installed version and components and determines whether this step is an upgrade operation.

6. Follow the wizard instructions to complete the upgrade process.
7. After the upgrade is completed, review the message that is displayed on the summary page. Review the installation log files, and correct any errors

Related concepts:

“Backups” on page 349

Before you can migrate, you must back up the entire Tivoli Storage Productivity Center system.

Related tasks:

“Preparing for an upgrade” on page 348

To upgrade Tivoli Storage Productivity Center, use the same starting procedure as the installation program. Depending on the components you install, you can see different pages.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Reviewing the log files to resolve installation issues” on page 165

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

Upgrading Tivoli Storage Productivity Center by using silent mode

Silent-mode upgrade is one method of upgrading Tivoli Storage Productivity Center. This installation method is useful if your system is running from a terminal that cannot display graphics.

Before you start the Tivoli Storage Productivity Center upgrade program on the AIX or Linux operating system, you must source the DB2 profile, `db2profile`, for the instance owner of the DB2 database. An example of sourcing the DB2 profile is:

```
. /home/db2inst1/sqllib/db2profile
```


To upgrade Tivoli Storage Productivity Center by using silent mode:

1. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
2. Back up Tivoli Storage Productivity Center.
3. Edit and save the appropriate response file.
4. Run the upgrade program. This is a list of the languages that can be used with the upgrade command:

- Czech - cs
- English - en
- French - fr
- German - de
- Hungarian - hu
- Italian - it
- Japanese - ja
- Korean - ko
- Polish - pl
- Brazilian Portuguese - pt_BR
- Russian - ru
- Spanish - es
- Chinese (Simplified) - zh_CN
- Chinese (Traditional) - zh_TW

- For Windows operating systems, run the following command:

```
setup.bat -l language -i silent -f absolute_path_to_response_file
```

absolute_path_to_response_file

The absolute path to the response file. This example specifies the language and the path:

```
setup.bat -l fr -i silent -f c:\TPC51\silent_Upgrade4x.properties.
```

- For the AIX or Linux operating system, run the following command:

```
./setup.bin -l language -i silent -f ./path_to_response_file
```

This example specifies the language and the path:

```
./setup.bin -l de -i silent -f ./TPC51/silent_Upgrade4x.properties
```

If you do not specify a language or specify the language incorrectly, the installation program uses the English language as the default.

5. Optional: Monitor the progress of the upgrade.

To view the progress of the silent mode installation on Windows operating systems, review the trace log in the following path:

```
TPC_installation_directory\logs\traceTPCInstall.log
```

To view the progress of the silent mode installation on AIX or Linux operating systems, review the trace log in the following path:

```
TPC_installation_directory/logs/traceTPCInstall.log
```

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Upgrade response file” on page 349

You do not need an upgrade response file to upgrade Tivoli Storage Productivity Center in silent mode from Version 5 to another Version 5 fix pack. To accept the license agreement, you must specify the value of the `LICENSE_ACCEPTED` variable to `true`. To upgrade from Version 5 to a Version 5 fix pack, run the command `setup.bat | bin -i silent -DLICENSE_ACCEPTED=true`.

“Reviewing the log files to resolve installation issues” on page 165

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

Upgrading Tivoli Storage Productivity Center in a multiple-server environment

You can upgrade Tivoli Storage Productivity Center by using the upgrade wizard or silent mode.

In a multiple-server environment, to upgrade Tivoli Storage Productivity Center, see the following topics:

- Using the upgrade wizard:
 - “Upgrading Tivoli Storage Productivity Center with a remote database by using the wizard” on page 355
 - “Upgrading Tivoli Integrated Portal and Tivoli Common Reporting by using the wizard” on page 356
- Using silent mode:
 - “Upgrading Tivoli Storage Productivity Center with a remote database by using silent mode” on page 357
 - “Upgrading Tivoli Integrated Portal and Tivoli Common Reporting by using silent mode” on page 360

Migrating from an older version of Tivoli Storage Productivity Center to Tivoli Storage Productivity Center 5.1

To migrate from Version 3.x to Tivoli Storage Productivity Center 5.1, you must first migrate from Version 3.x to Version 4.x. For more information about migrating from Version 3.x to Version 4.x, see http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=%2Fcom.ibm.tpc_V4221.doc%2Ffqz0_r_migrating_tpc.html. You can use your Version 3.x license to migrate from Tivoli Storage Productivity Center Version 3.x to Tivoli Storage Productivity Center 4.x.

Initially, Tivoli Storage Productivity Center 5.1 will be available through IBM Passport Advantage (download). If your license maintenance is current, you can download Tivoli Storage Productivity Center 5.1 from IBM Passport Advantage or access the physical media.

Upgrading Tivoli Storage Productivity Center with a remote database by using the wizard

The upgrade wizard is one method of upgrading Tivoli Storage Productivity Center.

Before you start the Tivoli Storage Productivity Center upgrade program on the AIX or Linux operating systems, you must source the DB2 profile, `db2profile`, for the instance owner of the DB2 database. The following example sources the DB2 profile:

```
. /home/db2inst1/sqllib/db2profile
```

You can upgrade Tivoli Storage Productivity Center in a multiple-server environment. With this configuration, you upgrade Tivoli Storage Productivity Center components on two servers. You upgrade the database server first on one server, and the other Tivoli Storage Productivity Center components on the second server.

In this procedure, the terms *Server A* and *Server B* denote the two servers. Server A has the Tivoli Storage Productivity Center database repository installed. Server B has the other Tivoli Storage Productivity Center components installed.

To upgrade Tivoli Storage Productivity Center by using the wizard:

1. On Server B, stop the all Tivoli Storage Productivity Center components.
2. Complete the following steps on Server A:
 - a. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
 - b. Back up the Tivoli Storage Productivity Center database repository.
 - c. Start the Tivoli Storage Productivity Center upgrade wizard.
 - d. Select the language for your installation and click **OK**.
 - e. On the Choose Installation Location and Type page, select **Multiple servers**.
 - f. Follow the wizard instructions to complete the upgrade process.
 - g. After the upgrade is completed, review the message that is displayed on the summary page.
3. Complete the following steps on Server B:
 - a. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
 - b. Back up Tivoli Storage Productivity Center (except for the database repository).
 - c. Select the language for your installation and click **OK**.
 - d. Follow the wizard instructions and complete the upgrade process.
 - e. After the installation is finished, review the message log to ensure that no errors occurred.

Related concepts:

“Backups” on page 349

Before you can migrate, you must back up the entire Tivoli Storage Productivity Center system.

Related tasks:

“Preparing for an upgrade” on page 348

To upgrade Tivoli Storage Productivity Center, use the same starting procedure as the installation program. Depending on the components you install, you can see different pages.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Reviewing the log files to resolve installation issues” on page 165

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

Upgrading Tivoli Integrated Portal and Tivoli Common Reporting by using the wizard

The upgrade wizard is one method of upgrading Tivoli Integrated Portal and Tivoli Common Reporting.

Before you start the Tivoli Storage Productivity Center upgrade program on the AIX or Linux operating systems, you must source the DB2 profile, `db2profile`, for the instance owner of the DB2 database. The following example sources the DB2 profile:

```
. /home/db2inst1/sqllib/db2profile
```

You can upgrade Tivoli Integrated Portal and Tivoli Common Reporting in a multiple-server environment. With this configuration, you upgrade the Tivoli Storage Productivity Center components first on one server, and Tivoli Integrated Portal and Tivoli Common Reporting on the second server.

The terms *Server A* and *Server B* denote the two servers. Server A has the Tivoli Storage Productivity Center components installed, and Server B has Tivoli Integrated Portal and Tivoli Common Reporting installed.

To upgrade Tivoli Integrated Portal and Tivoli Common Reporting by using the wizard:

1. On Server B, stop all Tivoli Storage Productivity Center components.
2. Complete the following steps on Server A:
 - a. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
 - b. Back up Tivoli Storage Productivity Center.
 - c. Start the Tivoli Storage Productivity Center upgrade wizard.
 - d. Select the language for your installation and click **OK**.
 - e. On the Choose Installation Location and Type page, select **Multiple servers**.
 - f. Follow the wizard instructions to complete the upgrade process.

- g. After the installation is finished, review the message log to ensure that no errors occurred.
3. Complete the following steps on Server B:
 - a. Follow the steps to create the instance owner ID. The administrative ID, `dasusr1`, and the fenced ID, `db2fenc1`, are not required.
 - b. If you are upgrading on a Windows operating system, log out and then log on again before upgrading Tivoli Storage Productivity Center.
 - c. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
 - d. Select the language for your installation and click **OK**.
 - e. On the Choose Installation Location and Type page, select **Multiple servers**.
 - f. Follow the wizard instructions to complete the upgrade process.
 - g. After the installation is finished, review the message log to ensure that no errors occurred.

Related concepts:

“Backups” on page 349

Before you can migrate, you must back up the entire Tivoli Storage Productivity Center system.

Related tasks:

“Installing DB2” on page 180

This topic describes how to install DB2 Version 9.7 for operating systems such as Linux, AIX, or Windows.

“Preparing for an upgrade” on page 348

To upgrade Tivoli Storage Productivity Center, use the same starting procedure as the installation program. Depending on the components you install, you can see different pages.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Reviewing the log files to resolve installation issues” on page 165

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

Upgrading Tivoli Storage Productivity Center with a remote database by using silent mode

Silent-mode upgrade is one method of upgrading Tivoli Storage Productivity Center.

Before you start the Tivoli Storage Productivity Center upgrade program on the AIX or Linux operating systems, you must source the DB2 profile `db2profile` for the instance owner of the DB2 database. An example of sourcing the DB2 profile is:

```
. /home/db2inst1/sqllib/db2profile
```

You can upgrade Tivoli Storage Productivity Center in a multiple-server environment. With this configuration, you upgrade the database server first on one server, and the other Tivoli Storage Productivity Center components on the second server.

The terms *Server A* and *Server B* denote the two servers. Server A has the Tivoli Storage Productivity Center database repository installed. Server B has the other Tivoli Storage Productivity Center components installed.

To upgrade Tivoli Storage Productivity Center by using silent mode:

1. On Server B, stop the all Tivoli Storage Productivity Center components.
2. Complete the following steps on Server A:
 - a. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
 - b. Back up the Tivoli Storage Productivity Center database repository.
 - c. Edit and save the appropriate response file.
 - d. Run the upgrade program. The following languages apply to commands you run:

language

- Czech – cs
- English – en
- French – fr
- German – de
- Hungarian – hu
- Italian – it
- Japanese – ja
- Korean – ko
- Polish – pl
- Brazilian Portuguese – pt_BR
- Russian – ru
- Spanish – es
- Chinese (Simplified) – zh_CN
- Chinese (Traditional) – zh_TW

absolute_path_to_response_file

The absolute path to the response file. This example specifies the language and the path:

```
setup.bat -l fr -i silent -f c:\TPC51\silent_Upgrade4x.properties.
```

- For Windows operating systems, run the following command:

```
setup.bat -l language -i silent -f absolute_path_to_response_file
```
- For the AIX or Linux operating system, run the following command:

```
./setup.bin -l language -i silent -f /absolute_path_to_response_file
```

This example specifies the language and the path:

```
./setup.bin -l de -i silent -f /TPC51/silent_Upgrade4x.properties
```

If you do not specify a language or specify the language incorrectly, the installation program uses the English language as the default.

- e. Optional: Monitor the progress of the upgrade.

To view the progress of the silent mode installation on Windows operating systems, check the trace log in the following path:

```
TPC_installation_directory\logs\traceTPCInstall.log
```

To view the progress of the silent mode installation on AIX or Linux operating systems, check the trace log in the following path:

TPC_installation_directory/logs/traceTPCInstall.log

3. Complete the following on Server B:
 - a. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
 - b. Back up Tivoli Storage Productivity Center (except for the database repository).
 - c. Before you start the Tivoli Storage Productivity Center upgrade program on the AIX or Linux operating system, you must source the DB2 profile `db2profile` for the instance owner of the DB2 database. An example of sourcing the DB2 profile is:

```
. /home/db2inst1/sqllib/db2profile
```
 - d. Edit and save the appropriate response file.
 - e. Run the upgrade program. Use one of the following languages when you run the command:

Language

- Czech – cs
- English – en
- French – fr
- German – de
- Hungarian – hu
- Italian – it
- Japanese – ja
- Korean – ko
- Polish – pl
- Brazilian Portuguese – pt_BR
- Russian – ru
- Spanish – es
- Chinese (Simplified) – zh_CN
- Chinese (Traditional) – zh_TW

absolute_path_to_response_file

The absolute path to the response file. The following example specifies the language and the path:

```
setup.bat -l fr -i silent -f c:\TPC51\silent_Upgrade4x.properties.
```

- For Windows operating systems, run the following command:

```
setup.bat -l language -i silent -f absolute_path_to_response_file
```
- For the AIX or Linux operating system, run the following command:

```
./setup.bin -l language -i silent -f /absolute_path_to_response_file
```

The following example specifies the language and the path:

```
./setup.bin -l de -i silent -f /TPC51/silent_Upgrade4x.properties
```

If you do not specify a language or specify the language incorrectly, the installation program uses the English language as the default.

- f. Optional: Monitor the progress of the upgrade.
 - To view the progress of the silent mode installation on Windows operating systems, check the trace log in the following path:

```
TPC_installation_directory\logs\traceTPCInstall.log
```
 - To view the progress of the silent mode installation on AIX or Linux operating systems, check the trace log in the following path:

Related tasks:

“Installing DB2” on page 180

This topic describes how to install DB2 Version 9.7 for operating systems such as Linux, AIX, or Windows.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Upgrade response file” on page 349

You do not need an upgrade response file to upgrade Tivoli Storage Productivity Center in silent mode from Version 5 to another Version 5 fix pack. To accept the license agreement, you must specify the value of the `LICENSE_ACCEPTED` variable to `true`. To upgrade from Version 5 to a Version 5 fix pack, run the command `setup.bat | bin -i silent -DLICENSE_ACCEPTED=true`.

“Reviewing the log files to resolve installation issues” on page 165

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

Upgrading Tivoli Integrated Portal and Tivoli Common Reporting by using silent mode

Silent-mode upgrade is one method of upgrading Tivoli Integrated Portal and Tivoli Common Reporting.

Before you start the Tivoli Storage Productivity Center upgrade program on the AIX or Linux operating systems, you must source the DB2 profile `db2profile` for the instance owner of the DB2 database. An example of sourcing the DB2 profile is:

```
. /home/db2inst1/sqllib/db2profile
```

You can upgrade Tivoli Storage Productivity Center in a multiple-server environment. With this configuration, you upgrade the Tivoli Storage Productivity Center components first on one server, and Tivoli Integrated Portal and Tivoli Common Reporting on the second server.

The terms *Server A* and *Server B* denote the two servers. Server A has the Tivoli Storage Productivity Center components installed. Server B has the Tivoli Integrated Portal and Tivoli Common Reporting installed.

To upgrade Tivoli Integrated Portal and Tivoli Common Reporting by using silent mode:

1. On Server B, stop the all Tivoli Storage Productivity Center components.
2. Complete the following steps on Server A:
 - a. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
 - b. Back up Tivoli Storage Productivity Center.
 - c. Edit and save the appropriate response file.
 - d. Run the upgrade program. Use one of the following languages when you run this command:

language

- Czech – cs
- English – en
- French – fr
- German – de
- Hungarian – hu
- Italian – it
- Japanese – ja
- Korean – ko
- Polish – pl
- Brazilian Portuguese – pt_BR
- Russian – ru
- Spanish – es
- Chinese (Simplified) – zh_CN
- Chinese (Traditional) – zh_TW

absolute_path_to_response_file

The absolute path to the response file. This example specifies the path and the language:

```
setup.bat -l fr -i silent -f c:\TPC51\silent_Upgrade4x.properties.
```

- For Windows operating systems, run the following command:

```
setup.bat -l language -i silent -f absolute_path_to_response_file
```
- For the AIX or Linux operating system, run the following command:

```
./setup.bin -l language -i silent -f /absolute_path_to_response_file
```

This example specifies the path and the language:

```
./setup.bin -l de -i silent -f /TPC51/silent_Upgrade4x.properties
```

If you do not specify a language or specify the language incorrectly, the installation program uses the English language as the default.

- Optional: Monitor the progress of the upgrade one of the following operating systems:
 - On Windows operating systems:

```
TPC_installation_directory\logs\traceTPCInstall.log
```
 - On AIX or Linux operating systems, check the trace log in the following path:

```
TPC_installation_directory/logs/traceTPCInstall.log
```

3. Complete the following steps on Server B:

- Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
- Before you start the Tivoli Storage Productivity Center upgrade program on the AIX or Linux operating system, you must source the DB2 profile `db2profile` for the instance owner of the DB2 database. An example of sourcing the DB2 profile is:

```
./home/db2inst1/sqllib/db2profile
```
- Edit and save the appropriate response file.
- Run the upgrade program. Use one of the following languages when you run the program:

language

- Czech – cs

- English – en
- French – fr
- German – de
- Hungarian – hu
- Italian – it
- Japanese – ja
- Korean – ko
- Polish – pl
- Brazilian Portuguese – pt_BR
- Russian – ru
- Spanish – es
- Chinese (Simplified) – zh_CN
- Chinese (Traditional) – zh_TW

absolute_path_to_response_file

The absolute path to the response file. This example specifies the language and the path:

```
setup.bat -l fr -i silent -f c:\TPC51\silent_Upgrade4x.properties.
```

- For Windows operating systems, run the following command:

```
setup.bat -l language -i silent -f absolute_path_to_response_file
```
- For the AIX or Linux operating system, run the following command:

```
./setup.bin -l language -i silent -f /absolute_path_to_response_file
```

This example specifies the language and the path:

```
./setup.bin -l de -i silent -f /TPC51/silent_Upgrade4x.properties
```

If you do not specify a language or specify the language incorrectly, the installation program uses the English language as the default.

- e. Optional: Monitor the progress of the upgrade on one of the following operating systems:
- On Windows operating systems, check the trace log in the following path:

```
TPC_installation_directory\logs\traceTPCInstall.log
```
 - On AIX or Linux operating systems, check the trace log in the following path:

```
TPC_installation_directory/logs/traceTPCInstall.log
```

Related tasks:

“Installing DB2” on page 180

This topic describes how to install DB2 Version 9.7 for operating systems such as Linux, AIX, or Windows.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Upgrade response file” on page 349

You do not need an upgrade response file to upgrade Tivoli Storage Productivity Center in silent mode from Version 5 to another Version 5 fix pack. To accept the license agreement, you must specify the value of the `LICENSE_ACCEPTED` variable to `true`. To upgrade from Version 5 to a Version 5 fix pack, run the command `setup.bat | bin -i silent -DLICENSE_ACCEPTED=true`.

“Reviewing the log files to resolve installation issues” on page 165

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

Deleting duplicate entries in a DB2 table when an upgrade to Tivoli Storage Productivity Center Version 5.1 fails

When you upgrade your system from Tivoli Storage Productivity Center 4 to Version 5.1, the DB2 repository upgrade might fail if the `t_stat_fs_hist` table contains duplicate entries.

To remove these duplicate entries, complete the following steps:

1. Start the DB2 command editor and connect to the Tivoli Storage Productivity Center DB2 repository. For more information about the DB2 command editor, see the DB2 information center and search for *Command Editor overview*.
2. Run the following query:

```
"select t1.filesystem_id, t1.statistic_type, t1.statistic_time,
t1.product_id from tpc.t_stat_fs_hist t1, (SELECT FILESYSTEM_ID as fs_id,
STATISTIC_TYPE as stype, STATISTIC_TIME as stime,
PRODUCT_ID as id from tpc.t_stat_fs_hist group by FILESYSTEM_ID,
STATISTIC_TYPE, STATISTIC_TIME, PRODUCT_ID having count(*) > 1)
where t1.filesystem_id=fs_id and t1.statistic_type = stype and
t1.statistic_time = stime and t1.product_id = id"
```
3. Click **Query Results** to see all the resulting entries.
4. Select a duplicate entry and click **Delete Row**. Deleting only the duplicate entries and keeping the original entries ensure that no DB2 repository data is lost.
5. Click **Commit**.

The table is updated, and you are ready to upgrade Tivoli Storage Productivity Center.

Upgrading Storage Resource agents

This section provides general steps for upgrading an existing or older version of Storage Resource agents to the current release.

If you are upgrading the Storage Resource agents, you can upgrade them using one of the following methods:

- Use the Tivoli Storage Productivity Center user interface. You can use the **Configuration** node or **Data Sources** node to upgrade the agents.
- Use a Storage Resource agent command.
- Use the installation wizard when upgrading the server. The installation wizard also upgrades the Storage Resource agent on the server.

Upgrading the Storage Resource agent by using the Data Sources node

Use this information to upgrade a previous Storage Resource agent to the current release.

To upgrade 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 topic pane, select one or more agents for which you want to upgrade. Click **Upgrade Agents**.
3. The Create Storage Resource Agent Upgrade page is displayed. Use this page to schedule an upgrade of the Storage Resource agent. You have the following tabs:

Computers

You can select Computer Groups or Computers.

When to Run

You can specify how often to run the job or how to handle time zones.

Options

You can specify to force an upgrade for an agent, maintenance upgrade for an agent, or language pack to apply to an agent.

Alerts You can specify the triggering conditions and triggering actions.

Enter your information in the appropriate fields.

4. Save and run the job.

Note: The Storage Resource agent upgrade job delivers the software upgrade packages to the agent computer. The job log displays the status of the delivery. The actual status of the upgrade is found in the agent log on the agent computer. If the agent log indicates that the upgrade failed and the state of the Storage Resource agent remains in the "Upgrading agent software status", try restarting the agent and running the upgrade job again.

Upgrading the Storage Resource agent by using the Configuration node

Use this information to upgrade a previous Storage Resource agent to the current release.

To upgrade a Storage Resource agent, see "Storage Resource Agent Upgrades" on page 295.

Upgrading Storage Resource agents by using a command

This section describes how to upgrade the Storage Resource agents manually.

To upgrade the Storage Resource agent agent, follow these steps:

1. Go to the DVD location of the installation program (using the Storage Resource Agent image) and go to the bin directory:

```
cd DVD_image_location/data/sra/operating_system_name
```

Where *DVD_image_location* is the Storage Resource agent installation image location.

2. Run the upgrade command:

```
► bin/Agent -upgrade -installLoc agent_install_directory
(1)
► -commType Daemon
```

Notes:

- 1 Parameter when the agent is run as a daemon service.
The parameters are:

-installLoc "*agent_install_directory*"

Location where the agent is installed. Enclose the directory name in quotation marks, for example, "C:\Program Files\IBM\TPC_SRA\".

-commType Daemon

If the agent is run as a daemon service, then this parameter must be specified.

Here is an example for a daemon-based service using the default location:

```
bin/Agent -upgrade
-installLoc "/opt/IBM/TPC/"
-commType Daemon
```

Here is an example for a non-daemon service using a non-default location:

```
bin/Agent
-upgrade -installLoc "C:\Program Files\IBM\TPC_SRA\"
```

Note:

If you run the upgrade program outside of the *DVD_image_location* installation directory, then you must specify the full path.

If the upgrade fails, see the return codes in the Information Center. Search for **Return codes used by Storage Resource agent**.

Migrating Data agents and Fabric agents to Storage Resource agents

The Storage Resource agents now perform the functions of the Data agents and Fabric agents. You can migrate the Data agents and Fabric agents to Storage Resource agents by using the Tivoli Storage Productivity Center user interface, upgrade installation wizard (for server system only), or using a command.

Migrating the Data agent or Fabric agent using the user interface

This topic describes how to migrate the Data agent or Fabric agent to a Storage Resource agent using the Tivoli Storage Productivity Center user interface.

To schedule a migration of the Data agent or Fabric agent to a Storage Resource agent through the user interface, complete the following steps:

1. Open the Tivoli Storage Productivity Center user interface.
2. In the navigation tree, expand **Administrative Services > Configuration > Migrate Data/Fabric Agent**.
3. Right-click the **Migrate Data/Fabric Agent** node.
4. Click **Create Data/Fabric Agent Migration**.
5. Specify the information for the migration job in the appropriate fields.
6. Click **File > Save** in the menu bar to save the migration definition. The agents you specify are migrated at the scheduled time.

Migrating Data agents and Fabric agents using a command

This section describes how to migrate the Data agents and Fabric agents to Storage Resource agents using a command.

To migrate the Data agent or Fabric agent to the Storage Resource agent, follow these steps:

1. If the Data agent or Fabric agent is on a Virtual I/O system, you must follow this log in procedure:
 - a. Log in with the **padmin** user ID.
 - b. Run the following command to set up the AIX environment:

```
eom_setup_env
```

2. Go to the DVD location of the installation program and go to the bin directory. The directory for disk 1:

```
DVD_image_location/data/sra/operating_system/bin
```

Where *DVD_image_location* is the Storage Resource agent installation image location and *operating_system* is the operating system type of the agent system.

3. Run the **Migration** command:

```
bin/Migration--MIGRATE [-COMMTYPE DAEMON] [-FORCE] [-SERVERPORT--port_number]
                        [-USERID--user_ID--password] [-PASSWORD--password]
                        [-CERTFILE--certificate_file] [-PASSPHRASE--passphrase]
                        [-DUSER--daemon_user_name] [-DPASSWORD--daemon_password] [-DEBUG min mid max]
```

The parameters are:

-COMMTYPE DAEMON

Specify this parameter if you want the Storage Resource agent to run as a daemon service.

-FORCE

This forces the migration to continue even if the Storage Resource agent cannot perform all the functions of the Data agent. The **-FORCE** option applies to the following conditions:

- The Data agent is monitoring an unsupported RDBMS type.
- The Data agent is responsible for a file system extension job.

-SERVERPORT *port_number*

Specify the port number of the Data server. This value is only required when the Data agent is not installed on the local system. The default port for the Data server is 9549.

-USERID *user_name*

For non-daemon service. The user ID defined on the agent system. This user ID is used by the server to connect to the agent system.

-PASSWORD *password*

For non-daemon service. Password for the user ID.

-CERTFILE *certificate_file*

For non-daemon service. The certificate used for SSH communication between the server and agent. This certificate needs to be stored on the server system.

-PASSPHRASE *passphrase*

For non-daemon service. The passphrase defined for the certificate used in SSH communication.

-DUSER *daemon_user_name* **(for Windows only)**

For daemon service. Login account used to start the Storage Resource agent service on Windows.

-DPASSWORD *daemon_password* **(for Windows only)**

For daemon service. Password for the Windows login account specified by the **-duser** parameter.

-DEBUG [**min** | **mid** | **max**]

Optional parameter for debugging purposes. You can specify the following:

min The minimum amount of debugging information.

mid A moderate amount of debugging information.

max The maximum amount of debugging information.

Here is an example for a daemon-based service:

```
bin/Migration -MIGRATE
-COMMTYPE DAEMON
-SERVERPORT 9549
-DEBUG max
```

Here is an example for a non-daemon service:

```
bin/Migration -MIGRATE
-SERVERPORT 9549
-USERID myuserid
-CERTFILE mycertfile
-PASSPHRASE mypassphrase
-DEBUG max
```

Note:

If you run the migration program outside of the *DVD_image_location* installation directory, then you must specify the full path.

Migrating Data agents and Fabric agents by using the upgrade wizard

This section describes how to migrate the Data agents and Fabric agents on the Tivoli Storage Productivity Center server to Storage Resource agent by using the upgrade wizard.

You can migrate the Data agent or Fabric agent to the Storage Resource agent when you upgrade Tivoli Storage Productivity Center. If you upgrade the Tivoli Storage Productivity Center servers, any Data agent or Fabric agent detected on the server is automatically migrated to a Storage Resource agent. For information about migrating the Data agent or Fabric agent, see “Upgrading Tivoli Storage Productivity Center by using the wizard” on page 351 and follow the upgrade wizard directions.

Upgrading CIM agents for storage systems

Before you upgrade a Common Information Model (CIM) agent, make sure that IBM Tivoli Storage Productivity Center supports the CIM agent that you want and that the agent is compatible with the firmware versions of your storage systems.

To ensure that a CIM agent is supported, review the supported products list for the current release of Tivoli Storage Productivity Center. Go to http://www.ibm.com/support/docview.wss?rs=40&context=SWJ50&q1=matrix&uid=swg21386446&loc=en_US&cs=utf-8&lang=en and click the applicable Tivoli Storage Productivity Center release in the Storage column. If you are uncertain about the CIM agent support, contact IBM customer support to help you with your upgrade plans.

Upgrading CIM agents

To upgrade a Common Information Model (CIM) agent, you must have the upgrade instructions that are supplied by the provider of the agent.

To upgrade the CIM agent, follow these steps:

1. Stop all Tivoli Storage Productivity Center server activity that might be dependent on the CIM agent (for example, performance monitor jobs, and discovery or probe jobs).
2. Stop the Tivoli Storage Productivity Center server services.
3. If any Tivoli Storage Productivity Center server processes remain, shut down the Tivoli Storage Productivity Center server for the duration of the CIM agent upgrade. If it is not possible to shut down the Tivoli Storage Productivity Center server, stop or kill the running processes.
4. Upgrade the CIM agent.
 - a. Refer to the instructions supplied by the CIM agent to perform the upgrade.
 - b. Verify that the CIM agent configuration is intact after the upgrade. Make sure that user accounts and passwords used with Tivoli Storage Productivity Center are still in place, and that all devices managed by the CIM agent are still listed.
5. Restart the Tivoli Storage Productivity Center server and Tivoli Storage Productivity Center server services.
6. Open the Tivoli Storage Productivity Center graphical user interface (GUI).

7. Run a CIM object manager (CIMOM) discovery job.
8. Run a probe job for each storage system that is managed by the CIM agent.

Migrating BIRT reports

You must migrate Business Intelligence and Reporting Tools (BIRT) reports before you upgrade Tivoli Storage Productivity Center.

Because Tivoli Storage Productivity Center does not migrate BIRT reports during the upgrade process, you must migrate the reports before the upgrade.

Tivoli Storage Productivity Center Version 4 can include Tivoli Integrated Portal Version 1.1 and Tivoli Common Reporting Version 1.2. Tivoli Storage Productivity Center Version 5 includes Tivoli Integrated Portal Version 2.2 and Tivoli Common Reporting Version 2.1.1.

The BIRT reports are in the following Tivoli Common Reporting installation directory:

- For Tivoli Storage Productivity Center Version 4:
 - For Windows operating systems: `c:/Program Files/IBM/Tivoli/tip/products/tcr`
 - For Linux and AIX operating systems: `/opt/IBM/Tivoli/tip/products/tcr`
- For Tivoli Storage Productivity Center Version 5:
 - For Windows operating systems: `c:/Program Files/IBM/tipv2Components/TCRComponent`
 - For Linux and AIX operating systems: `/opt/IBM/tipv2Components/TCRComponent`

To migrate BIRT reports from Tivoli Storage Productivity Center Version 4 to Version 5:

1. Export BIRT reports from Tivoli Storage Productivity Center Version 4. Run the following command:

```
TCR_installation_directory/bin/tcrmd -export
```

where *TCR_installation_directory* is the location of the BIRT reports for Tivoli Storage Productivity Center Version 4.


2. Copy the custom jar files to a temporary location. These files are in the *TCR_installation_directory/lib* directory.
3. Upgrade Tivoli Storage Productivity Center Version 4 to Tivoli Storage Productivity Center Version 5.
4. Import the BIRT reports that you exported in step 1 into Tivoli Storage Productivity Center Version 5. Run the following command:

```
TCR_installation_directory/bin/tcrmd -import
```

where *TCR_installation_directory* is the location of the BIRT reports for Tivoli Storage Productivity Center Version 5. For information about the **-import** command, see the Tivoli Common Reporting information center.

5. Copy the custom jar files that you saved in step 2. Copy these files to the *TCR_installation_directory/lib* directory.

Related information:

 http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/index.jsp?topic=%2Fcom.ibm.tivoli.tcr.doc_211%2Fic-home.html

Upgrading System Storage N Series Gateway servers

Follow this procedure to upgrade Tivoli Storage Productivity Center before adding N Series Gateway servers by using the **Other NAS** node.

Deleting and unlicensing existing N Series Gateway servers

If you are upgrading Tivoli Storage Productivity Center from TotalStorage Productivity Center 3.3.x, there are specific steps to follow. You must first delete and unlicense existing N Series Gateway servers before adding the N Series Gateway servers.

Follow these steps:

1. Stop all Tivoli Storage Productivity Center services.
2. Make a backup copy of the Tivoli Storage Productivity Center database.
3. Open the Tivoli Storage Productivity Center GUI.
4. Manually delete all N Series Gateway servers.
 - a. Expand **Administrative Services > Configuration > Manual NAS Server Entry**.
 - b. Select a NAS gateway server (for example, ratbert1) from the Manual NAS Server table.
 - c. Click **Delete**.
 - d. Click **Yes** to **Delete Manual NAS Server Entry** prompt.
 - e. Repeat steps b through d for other manually added NAS gateway servers.
5. Unlicense all NAS gateway servers that were automatically discovered
 - a. Click **Administrative Services > Configuration > License Keys**.
 - b. Expand **IBM TPC for Data** under the Licensing tab.
 - c. Clear the Licensed check box for all automatically discovered NAS gateway servers (for example, zinc).
 - d. Click the disk icon to save the changes.
 - e. Click **Yes** to the Save Confirmation prompt.

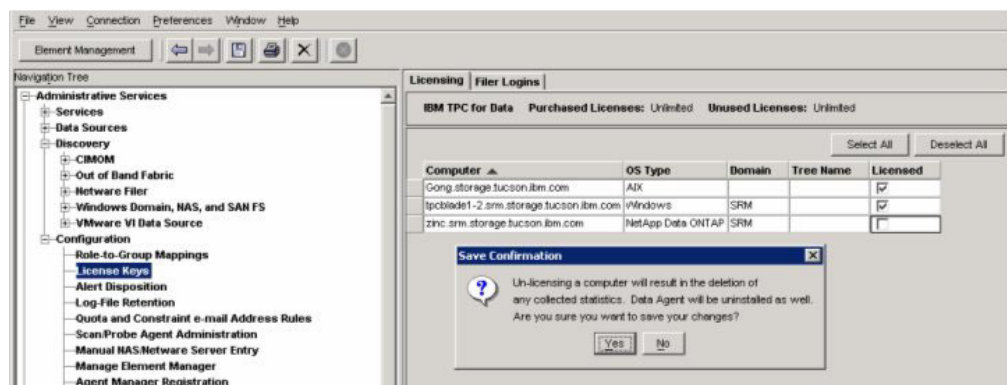


Figure 14. Save Confirmation prompt

Note: The OS Type and Licensed fields of existing NAS gateway servers are updated after it is manually added as **Other NAS**. See “Manually adding an N Series Gateway server .”

6. Exit the Tivoli Storage Productivity Center GUI.
7. Make another backup of the Tivoli Storage Productivity Center database after the N Series Gateway servers are deleted and unlicensed (make sure all Tivoli Storage Productivity Center services are stopped).
8. Restart all the Tivoli Storage Productivity Center services.
9. Go to “Manually adding an N Series Gateway server .”

Manually adding an N Series Gateway server

1. Start the Tivoli Storage Productivity Center GUI.
2. Manually add the N Series Gateway server as **Other NAS**.
 - a. Click **Administrative Services > Configuration > Manual NAS Server Entry**.
 - b. Click **Add NAS Server**.
 - c. Enter information for the Network Name, Accessible from, SNMP Community, Login ID, and Password fields for an N Series Gateway server (for example, **ratbert1**).
 - d. Check the **Add as Other NAS** check box.
 - e. Enter the NAS Server Vendor Name or accept the Network Appliance as the default.
 - f. Click **OK** to save the input. This action enables IBM Tivoli Storage Productivity Center to verify the NAS gateway server information and discover CIFS or NFS shares accessible to the selected Data agent or Storage Resource agent.

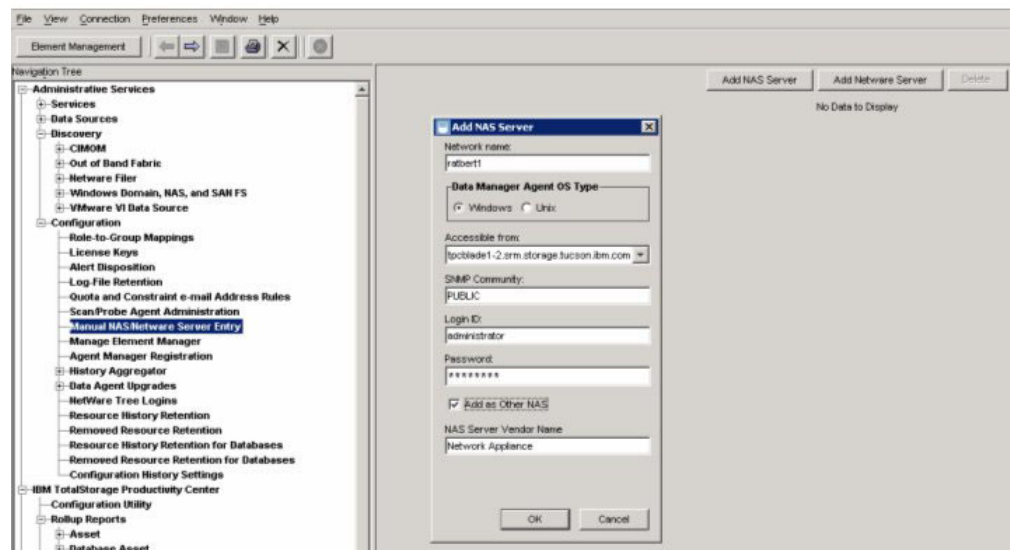


Figure 15. Add NAS Server panel

- g. Verify that the N Series Gateway server is added and listed with **Other NAS** as OS Type.

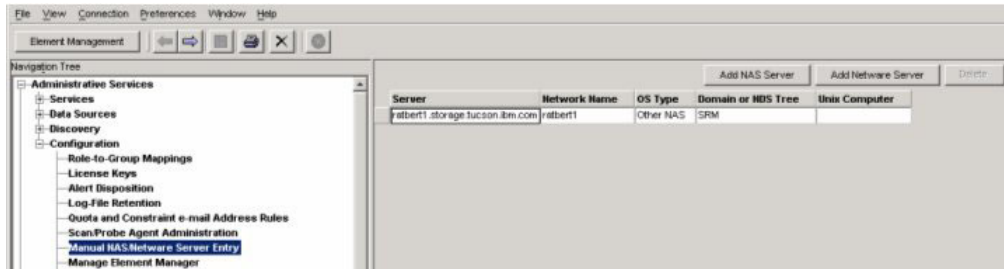


Figure 16. Verify that the N Series Gateway server is added

- h. Repeat steps b to g for other N Series Gateway servers.
3. Assign a Scan/Probe agent or agents to the exported CIFS or NFS shares accessible to the selected Data agent or Storage Resource agent.
 - a. Click **Administrative Services > Configuration > Scan/Probe Agent Administration**.
 - b. Select one or more file systems (for example, /vol/tpc1, /vol/tpc2, and /vol/vol0 for ratbert1) and assign the Data agent or Storage Resource agent by clicking **Set agent per row** or **Set agent for all selected rows**.

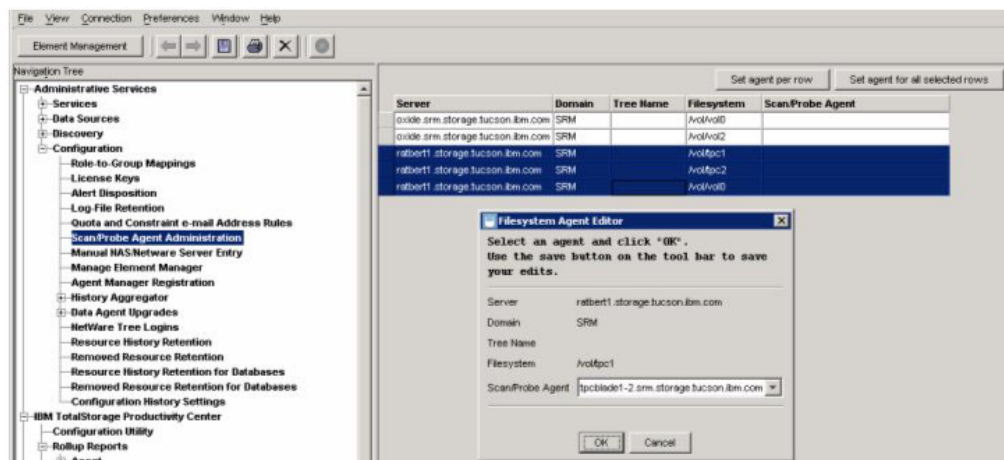


Figure 17. File system Agent Editor panel

- c. Click **OK** to save the assignment.




Figure 18. Scan/Probe Agent Administration panel

- d. Repeat steps b and c for other file systems (for example, /vol/vol0 and /vol/vol2 for oxide).

- e. Click the disk icon to save the **Scan/Probe Agent assignment** (so that the N Series Gateway servers and their file systems can be displayed when a probe or scan job is run).
4. Create a probe job to probe one or more N Series Gateway servers.
5. Create a scan job to scan one or more file systems from the N Series Gateway servers.
6. View asset report on a NAS gateway server and its file systems as **Other NAS**.
 - a. Click **Administrative Services > Data Manager > Reporting > Asset > By OS Type > Other NAS**.
 - b. Click an N Series Gateway server (for example, ratbert1) to view information for the server.

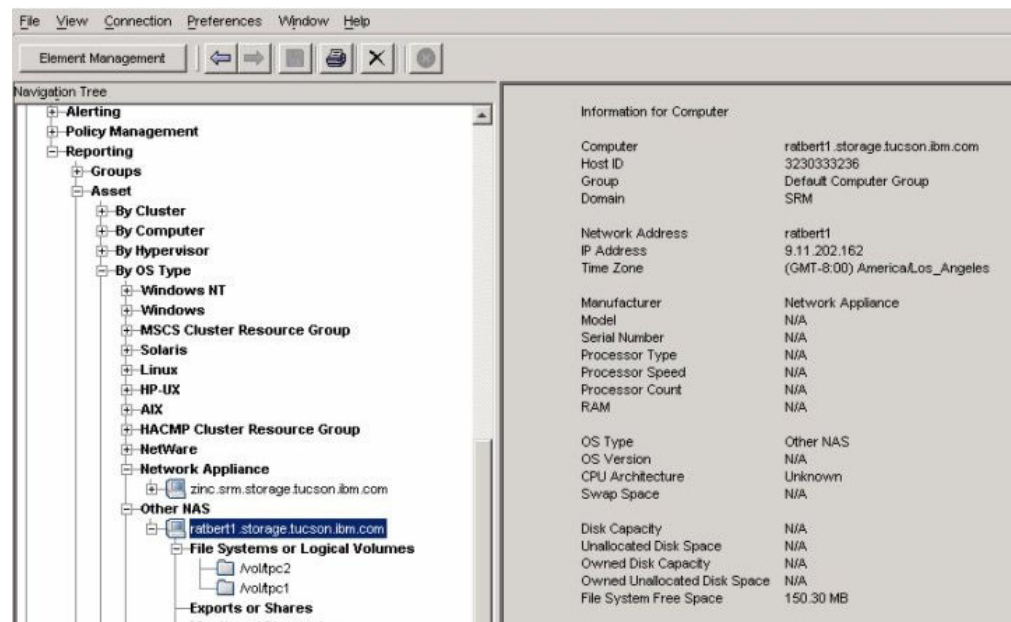


Figure 19. View information for N Series Gateway server

- c. Click a file system (for example, /vol1/tpc2) to view file system information.

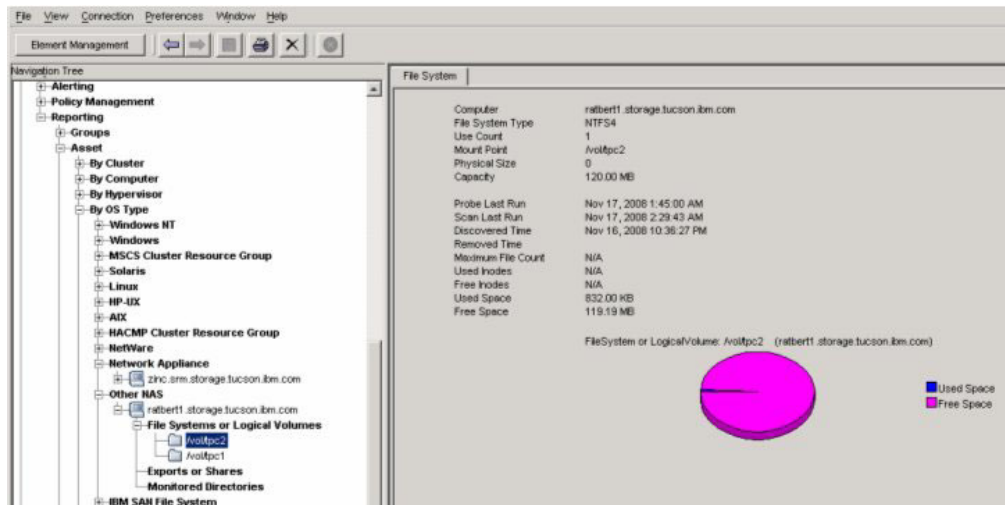


Figure 20. View information for a file system

7. You can see the following information and reports for **Other NAS**:
 - a. Dashboard
 - b. System reports
 - c. Batch reports
 - d. Asset reports
 - e. Capacity reports
 - Capacity reports - Charting
 - f. Usage reports
 - Usage reports - Filesystem Quota Violation
 - Capacity reports - Charting
8. You can monitor the N Series Gateway servers and create and monitor alerts for **Other NAS**:
 - Monitoring
 - Monitoring - Assign a Scan agent to a Filesystem
 - Monitoring - Grouping
 - Monitoring - Scan job
 - Monitoring - Filesystem Quotas
 - Alerting
 - Filesystem Alert definition
 - Alerting - Directory Alerts
 - Alerting - Alert Log
 - Alerting - Alerts shown on the topology

Chapter 5. Uninstalling Tivoli Storage Productivity Center

You can uninstall Tivoli Storage Productivity Center by using the uninstallation wizard or by using silent mode from the command line.

You must uninstall the components and related software in the following order:

1. If you installed a Storage Resource agent on many systems, uninstall those agents before you uninstall the stand-alone GUI.
2. If all Tivoli Storage Productivity Center components are installed on one server, uninstall all components.
3. If Tivoli Integrated Portal and Tivoli Common Reporting are installed on a separate server, uninstall those components.
4. If the database repository is installed on a separate server, uninstall the database repository.
5. Uninstall DB2.

Note: On the Windows operating system, if you have installed the Monitoring Agent service, you must stop the **Monitoring Agent for Windows OS - Watchdog** and **Monitoring Agent for Windows OS - Primary** services before you uninstall Tivoli Storage Productivity Center.

Uninstalling Tivoli Storage Productivity Center in a single-server environment by using the wizard

In a single-server environment, you can uninstall Tivoli Storage Productivity Center by using a wizard. With the wizard, the uninstallation process requires minimal user interaction.

Before you start the Tivoli Storage Productivity Center uninstallation program, you must source the DB2 profile `db2profile` for the instance owner of the DB2 database. For example:

```
. /home/db2inst1/sqllib/db2profile
```

To uninstall the software in a single-server environment:

1. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
2. Start the Tivoli Storage Productivity Center uninstallation wizard in one of the following operating systems:
 - For the Windows operating system:
 - a. Click **Start > Control Panel > Programs > Programs and Features**.
 - b. Highlight **Tivoli Storage Productivity Center** and click **Uninstall/Change**.

- For the AIX or Linux operating systems:

Run the following command from the root directory:

```
TPC_installation_directory/_uninst/uninstall
```

where *TPC_installation_directory* is where Tivoli Storage Productivity Center is installed.

3. Review the message that indicates that all installed Tivoli Storage Productivity Center components is uninstalled.
4. Click **Uninstall** and click **Next**. You cannot cancel the uninstallation process after it starts.
5. Restart the system if needed.
 - For the AIX or Linux operating system, you do not have to restart the system.
 - For the Windows operating system, you can restart the system now or later. You must restart the system to reinstall Tivoli Storage Productivity Center.
6. If an error occurred during the installation process, review the installation log files.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Reviewing the log files to resolve installation issues” on page 165

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

Uninstalling Tivoli Storage Productivity Center in a multiple-server environment by using the wizard

Uninstalling Tivoli Storage Productivity Center by using the uninstallation wizard requires minimal user interaction.

For *Configuration A* and *Configuration B*, for the AIX or Linux operating systems, you must source the DB2 profile `db2profile` for the instance owner of the DB2 database. The following example sources the DB2 profile:

```
. /home/db2inst1/sqllib/db2profile
```

For this procedure, the terms *Server A* and *Server B* denote the two servers.

- *Configuration A*: The Tivoli Storage Productivity Center database repository is installed on Server A and all other Tivoli Storage Productivity Center components are installed on Server B.
- *Configuration B*: The Tivoli Storage Productivity Center database repository, Tivoli Storage Productivity Center Servers, Storage Resource agents, stand-alone GUI, and command-line interface (CLI) on Server A and Tivoli Integrated Portal, Tivoli Common Reporting, and web-based GUI on Server B.

To uninstall the software in a multiple-server environment:

1. Complete the following steps for *Configuration A*:
 - a. Complete these steps on Server B for *Configuration A*:
 - 1) Log on to Server B with the appropriate user privileges to uninstall the Tivoli Storage Productivity Center components.
 - 2) Start the Tivoli Storage Productivity Center uninstallation wizard. For the Windows operating system, click **Start > Control Panel > Programs > Programs and Features**. Highlight **Tivoli Storage Productivity Center** and click **Uninstall/Change**.

For the AIX or Linux operating system, run the following command from the root directory:

```
TPC_installation_directory/_uninst/uninstall
```

where *TPC_installation_directory* is the location where Tivoli Storage Productivity Center is installed.

A window is displayed that indicates that all Tivoli Storage Productivity Center components installed on the system will be uninstalled.

- 3) Click **Uninstall**. A confirmation is displayed.
- 4) Click **Next** to start the uninstallation process.

Remember the following for your operating system:

- For the Windows operating system, when the uninstallation process is finished, a window is displayed so you can choose to start your system now or later. You must restart the system before you can reinstall Tivoli Storage Productivity Center.
- For the AIX or Linux operating system, you do not have to restart your system.

b. Complete the following steps on Server A:

- 1) Log on to Server A with the appropriate user privileges to uninstall the Tivoli Storage Productivity Center database repository.
- 2) Start the Tivoli Storage Productivity Center uninstallation wizard.
- 3) Follow the instructions in the wizard to uninstall the database repository.

2. Complete these steps for *Configuration B*:

a. Do the following on Server B:

- 1) Log on to Server B with the appropriate user privileges to uninstall Tivoli Integrated Portal and Tivoli Common Reporting.
- 2) Start the Tivoli Storage Productivity Center uninstallation wizard.
- 3) Proceed through the wizard to uninstall Tivoli Integrated Portal and Tivoli Common Reporting.

b. Complete the following steps on Server A:

- 1) Log on to Server A with the appropriate user privileges to uninstall the Tivoli Storage Productivity Center components.
- 2) Start the Tivoli Storage Productivity Center uninstallation wizard.
 - For the Windows operating system, click **Start > Control Panel > Programs > Programs and Features**. Highlight **Tivoli Storage Productivity Center** and click **Uninstall/Change**.
 - For the AIX or Linux operating systems, run the following command from the root directory:

```
TPC_installation_directory/_uninst/uninstall
```

where *TPC_installation_directory* is the location where Tivoli Storage Productivity Center is installed.

A window is displayed that indicates that all Tivoli Storage Productivity Center components installed on the system will be uninstalled.

- 3) Click **Uninstall**. A confirmation is displayed.
- 4) Click **Next** to start the uninstallation process.
- 5) Follow the instructions in the wizard to uninstall Tivoli Storage Productivity Center.

3. If an error occurred during the installation process, review the installation log files.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Reviewing the log files to resolve installation issues” on page 165

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

Uninstalling Tivoli Storage Productivity Center by using silent mode

You can uninstall Tivoli Storage Productivity Center by using silent mode. This mode is useful if your system is running from a terminal that cannot display graphics.

In silent mode, the following server configurations are possible:

Tip: The terms *Server A* and *Server B* denote the two servers in a multiserver environment.

Option 1

All Tivoli Storage Productivity Center components are installed on one server (single-server environment).

Option 2

The Tivoli Storage Productivity Center database repository is installed on Server A and all other Tivoli Storage Productivity Center components are installed on Server B.

Option 3

Tivoli Integrated Portal, Tivoli Common Reporting, and the web-based GUI are installed on Server B and the other Tivoli Storage Productivity Center components are installed on Server A.

To uninstall Tivoli Storage Productivity Center in silent mode:

1. Complete the following steps for option 1:
 - a. Log on to the Tivoli Storage Productivity Center system with the appropriate user privileges.
 - b. Start the uninstallation program.
 - For the Windows operating system, run the following in the C:\ directory:

```
TPC_installation_directory\_uninst\uninstall.bat -i silent
```
 - For the AIX or Linux operating systems, source the DB2 profile, and run the following command from the root directory:

```
TPC_installation_directory/_uninst/uninstall -i silent
```
 - c. Restart the system if needed.
 - For the Windows operating systems, restart the system after the uninstallation process is complete.

- For the AIX or Linux operating systems, you do not have to restart your system.
2. Complete the following steps for option 2:
 - a. Log on to Server B with the appropriate user privileges to uninstall Tivoli Storage Productivity Center.
 - b. Start the uninstallation program.
 - For the Windows operating system, run the following from the C:\ directory:
`TPC_installation_directory_uninst\uninstall.bat -i silent`
 - For the AIX or Linux operating system, source the DB2 profile, and run the following command from the root directory:
`TPC_installation_directory/_uninst/uninstall -i silent`
 - c. Restart the system if needed.
 - For the Windows operating systems, restart the system after the uninstallation process is complete.
 - For the AIX or Linux operating systems, you do not have to restart your system.
 - d. Log on to Server A with the appropriate user privileges to uninstall the Tivoli Storage Productivity Center database repository.
 - e. Start the uninstallation program.
 - For the Windows operating system, run the following command from the C:\ directory:
`TPC_installation_directory_uninst\uninstall.bat -i silent`
 - f. Restart the system if needed.
 - For the Windows operating systems, restart the system after the uninstallation process is complete.
 - For the AIX or Linux operating systems, you do not have to restart your system.
 3. Complete the following steps for option 3:
 - a. Log on to Server B with the appropriate user privileges to uninstall Tivoli Integrated Portal and Tivoli Common Reporting.
 - b. Start the uninstallation program.
 - For the Windows operating system, run the following from the C:\ directory:
`TPC_installation_directory_uninst\uninstall.bat -i silent`
 - For the AIX or Linux operating system, source the DB2 profile, and run the following from the root directory:
`TPC_installation_directory/_uninst/uninstall -i silent`

where *TPC_installation_directory* is where Tivoli Storage Productivity Center is installed.
 - c. Restart the system if needed.
 - For the Windows operating systems, restart the system after the uninstallation process is complete.
 - For the AIX or Linux operating systems, you do not have to restart your system.
 - d. Log on to Server A with the appropriate user privileges to uninstall the Tivoli Storage Productivity Center components (including the database repository).

- e. Start the uninstallation program.
 - For the Windows operating system, run the following from the C:\ directory:

```
TPC_installation_directory\uninst\uninstall.bat -i silent
```

Remember: You must restart the system after the uninstallation process is finished. There is no automatic reboot.
 - For the AIX or Linux operating system, source the DB2 profile, and run the following from the root directory:

```
TPC_installation_directory/_uninst/uninstall -i silent
```

where *TPC_installation_directory* is where Tivoli Storage Productivity Center is installed.
- f. If an error occurred during the installation process, review the installation log files.

Related tasks:

“Manually uninstalling Tivoli Storage Productivity Center components on Windows” on page 387

If the Tivoli Storage Productivity Center uninstallation process fails, you must manually uninstall the components.

“Manually uninstalling Tivoli Storage Productivity Center components on AIX or Linux operating systems” on page 386

If the Tivoli Storage Productivity Center uninstallation process fails, you must manually uninstall the components.

Related reference:

“Planning for Tivoli Storage Productivity Center authentication and authorization” on page 26

An operating system user name is required to install and log on to Tivoli Storage Productivity Center for the first time. After you install Tivoli Storage Productivity Center, you can assign roles to users that enable users to log on and determine the components that the user can view in the graphical user interface (GUI).

“Reviewing the log files to resolve installation issues” on page 165

If an error occurs during Tivoli Storage Productivity Center installation, you can review the error log files to resolve issues and continue the installation.

Uninstalling the agents

You can uninstall locally or remotely installed agents. This section describes how to uninstall these agents locally or remotely.

The Storage Resource agents are uninstalled by using the Tivoli Storage Productivity Center GUI.

Restriction: Do not delete the Storage Resource agent component from the Tivoli Storage Productivity Center server by using the stand-alone GUI or the command-line interface (CLI). This component must exist on the server to complete the uninstall process.

Deleting or uninstalling an agent using the GUI

You can delete or uninstall a Storage Resource agent and all the data that is collected by that agent from the database repository. The information that is

collected by the agent is no longer available within Tivoli Storage Productivity Center reports. When you delete an agent, the agent is uninstalled and you can no longer activate the agent.

To delete or uninstall a Storage Resource agent:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select an agent and under the Select Action menu, click **Delete**.
3. In the confirmation dialog, select **Yes Delete the agent**.

Uninstalling the Storage Resource agent manually

You can manually uninstall the Storage Resource agent.

To uninstall the Storage Resource agent, complete the following steps:

1. If the Storage Resource agent is on the Virtual I/O system, complete the following steps:
 - a. Log in with the **padmin** user ID.
 - b. Run the following command to set up the AIX environment:
`oem_setup_env`
2. Go to the directory where the agent is installed:
`agent_install_location`

where `agent_install_location` is where the Storage Resource agent is installed.

3. Run the uninstallation command:

Restriction: Do not delete the Storage Resource agent component from the Tivoli Storage Productivity Center server by using the stand-alone GUI or the command-line interface (CLI). This component must exist on the server to complete the uninstall process.

```
▶ bin/Agent --uninstall [-force] --serverName TPC_server_name [-debug MAX] ▶
```

Where:

-uninstall

Uninstalls the agent.

-force

This optional parameter forces an uninstallation. If you use this parameter, do not provide the `serverName` parameter.

-serverName *TPC_server_name*

TPC_server_name is the Data server name as defined in IBM Tivoli Storage Productivity Center. You can check the configuration file for the server name:

```
agent_install_location/config/Agent.config
```

where `agent_install_location` is where the agent is installed.

-debug MAX

This optional parameter is for debugging purposes. If you set the **-debug** parameter, then some files are not deleted.

Remember: If you run the uninstallation program from the bin directory, the bin directory is not deleted.

If you run the uninstallation program outside of the agent installation directory, then you must specify the full path.

If the uninstallation fails, you must look at the return codes. For more information about return codes, see <http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for **Return codes used by Storage Resource agent**.

Uninstalling the Tivoli Storage Productivity Center Monitoring Agent

This procedure provides information about uninstalling the Tivoli Storage Productivity Center Monitoring Agent.

Uninstalling the IBM Tivoli Storage Productivity Center Monitoring Agent on Windows

You can uninstall the IBM Tivoli Storage Productivity Center Monitoring Agent on Windows.

To uninstall the Tivoli Storage Productivity Center Monitoring Agent on Windows:

1. Stop the IBM Tivoli Monitoring agent before uninstalling the agent. For information about stopping the agent for Windows, go to http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/index.jsp?topic=/com.ibm.omegamon.mes.doc_7.0.1/t_stop_win.html.
2. Log on to the system as a local administrator with administrative authority on Windows.
3. Exit the Tivoli Storage Productivity Center GUI.
4. Click **Start > Settings > Control Panel > Add/Remove Programs**.
5. Select **IBM Tivoli Storage Productivity Center** and click **Remove**.
6. On the language page, select a language and click **OK**.
7. On the Welcome page, click **Next**.
8. Select **IBM Tivoli Storage Productivity Center Monitoring Agent** and click **Next**.
9. On the Summary Information page, review the information and click **Uninstall**.
10. Click **Finish**.
11. On the Restart Your Computer page, click **Yes, restart my computer**.
12. Click **Finish**.

Uninstalling the IBM Tivoli Storage Productivity Center Monitoring Agent on UNIX or Linux

You can uninstall the IBM Tivoli Storage Productivity Center Monitoring Agent on UNIX or Linux.

To uninstall the Tivoli Storage Productivity Center Monitoring Agent on UNIX or Linux:

1. Stop the IBM Tivoli Monitoring agent before uninstalling the agent. For information about stopping the agent on UNIX or Linux operating systems, go

to <http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/topic/com.ibm.docguide.doc/using.htm> and search for *Starting or stopping an agent on UNIX and Linux systems*.

2. For UNIX and Linux, ensure that you have root authority.

Tip: For UNIX and Linux systems, you must source the DB2 profile before uninstalling the Tivoli Storage Productivity Center Monitoring Agent.

3. Open a terminal session and go to the Tivoli Monitoring Services installation directory.
4. Run the uninstallation program from this default directory:
`/opt/IBM/ITM/bin/uninstall.sh`
5. On the Manage Tivoli Enterprise Monitoring Services window, select **Monitoring Agent for TPC** and press **Enter**.
6. Confirm your selection by entering **1** and pressing **Enter** again.
7. Wait for the uninstallation program to complete.

For more information about uninstalling an agent from Tivoli Monitoring Services, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.1/itm_install427.htm.

Uninstalling DB2

Uninstall DB2 after you have uninstalled IBM Tivoli Storage Productivity Center and Agent Manager. Use the DB2 uninstallation program to uninstall DB2.

For information about uninstalling DB2 V9.1 on Windows systems, go to <http://publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp?topic=/com.ibm.db2.udb.uprun.doc/doc/t0007436.htm>.

For information about uninstalling DB2 V9.5 on Windows systems, go to <http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/index.jsp?topic=/com.ibm.db2.luw.qb.server.doc/doc/t0007436.html>.

Uninstalling DB2 on UNIX or Linux

You can uninstall DB2 on UNIX or Linux systems.

To uninstall DB2 on UNIX and Linux:

Note: This example applies to DB2 V9.7.

1. As the DB2 instance user (typically `db2inst1`), drop the databases, and stop DB2.
 - a. If you are the user with root authority, run `su - db2inst1` to switch the user to the DB2 instance user. Otherwise, log in as the DB2 instance user.
 - b. Run `db2 list db directory` to list the databases that might not have been removed by the Tivoli Storage Productivity Center uninstallation program.
 - c. For each database name listed by the command in the previous step, run `db2 drop db database_name`. For example:

Command:

```
$ db2 drop db tcpdb
```

Response:

```
DB20000I The DROP DATABASE command completed successfully.
```

```
$ db2 drop db ibmcdb
```

Response:

```
DB20000I The DROP DATABASE command completed successfully.
```

- d. Run **db2stop force**. The following messages are displayed:


```
08/10/2010 21:46:13    0    0    SQL1064N  DB2STOP processing was successful.
SQL1064N  DB2STOP processing was successful.
```
 - e. Run **db2 terminate**. The following message is displayed:


```
DB20000I The TERMINATE command completed successfully.
```
 - f. Run **exit** to close the su shell or log out.
2. Drop the DB2 instance.

Note: This step must be performed as a user with root authority and run in a shell that has not had the db2profile sourced.

- a. Log in as a user with **root** authority.
 - b. Change to the DB2 instance directory. The default path depends on your operating system.

On AIX, run **cd /opt/IBM/db2/V9.7/instance**.

On Linux, run **cd /opt/ibm/db2/V9.7/instance**.
 - c. Run **./db2ilist** to list the db2instance name (typically db2inst1).
 - d. Run **./db2idrop -f db2inst1** where "db2inst1" is the instance name reported by the **db2ilist** command in the previous step. The following messages are displayed:


```
DBI1324W Support of the -f command is deprecated. For more information,
see the DB2 Information Center.
DBI1070I Program db2idrop completed successfully.
```
3. As the DB2 administration server (DAS) user (typically dasusr1), stop DAS.
- a. If you are the user with root authority, run **su - dasusr1** to switch user to the DAS user. Otherwise, log in as the DAS user.
 - b. Run **db2admin stop**. The following message is displayed:


```
SQL4407W The DB2 Administration Server was stopped successfully.
```
 - c. Run **exit** to close the su shell or log out.
4. Remove the DB2 administration server. You must remove the DB2 administration server before you remove the DB2 product.
- a. Log in as a user with **root** authority.
 - b. Change to the DB2 instance directory. The default path depends on your operating system.

On AIX, run **cd /opt/IBM/db2/V9.7/instance**.

On Linux, run **cd /opt/ibm/db2/V9.7/instance**.
 - c. Run **./dasdrop**. The following messages are displayed:


```
SQL4410W The DB2 Administration Server is not active.
DBI1070I Program dasdrop completed successfully.
```
5. Uninstall DB2. Use the **db2_deinstall** command to remove DB2 products or DB2 components. The **db2_deinstall** command removes all DB2 products from your system.
- a. Change to the DB2 installation directory. The default path depends on your operating system.

On AIX, run **cd /opt/IBM/db2/V9.7/instance**.

On Linux, run **cd /opt/ibm/db2/V9.7/instance**.
 - b. Run **./db2_deinstall -a**. The following messages are displayed:


```
DBI1016I Program db2_deinstall is performing uninstallation.
Please wait.
The execution completed successfully.
For more information see the DB2 uninstallation log at
"/tmp/db2_deinstall.log.pid".
```

6. Delete DB2 users and groups. The default user and group names are shown in the following commands.
 - a. Run the following commands to remove the users:

```
userdel -r dasusr1
userdel -r db2inst1
userdel -r db2fenc1
```
 - b. Run the appropriate commands for your operating system to remove the groups.

On AIX run the following commands:

```
rmgroup dasadm1
rmgroup db2iadm1
rmgroup db2fadm1
```

On Linux run the following commands:

```
groupdel dasadm1
groupdel db2iadm1
groupdel db2fadm1
```
7. Remove the DB2 directory.
 - a. Run `cd`.
 - b. Run the command to remove the DB2 directory on your operating system:

On AIX, run `rm -r /opt/IBM/db2`.

On Linux, run `rm -r /opt/ibm/db2`.
8. Remove any DB2 definitions from `/etc/services`.

Deleting Storage Resource agent registry entries after a failed installation or uninstallation

If your Tivoli Storage Productivity Center Version 5.1 installation or uninstallation fails, you must find and manually delete the Storage Resource agent registry entries that remain.

Deleting registry files in the Windows operating system

To delete registry entries after a failed installation or uninstallation in the Windows operating system:

1. In Windows Registry Editor, navigate to HKEY_LOCAL_MACHINE > SOFTWARE > IBM > TPCAGENT > NUMBER or HKEY_LOCAL_MACHINE > SOFTWARE > Wow6432Node > IBM > TPCAGENT > NUMBER.
2. Delete this entry:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\IBM\TPCAGENT\1]
"Home"="C:\Program Files\IBM\TPC"
"Port"="9510"
```
3. If the Storage Resource agent is running as a daemon service, delete the service. For example, `sc delete AgentX`, where AgentX is the name of the Storage Resource agent.
4. Delete the folder in which the agent was installed.

Display the agent name and folder by running `services.msc` and right-click the service to see the properties.

- For daemon agents, verify that the associated CAP file was deleted.
The CAP file is located in the %ALLUSERSPROFILE%\Application Data\IBM\CAP folder. The naming convention for this file is `srainstance_default.xml`.

Remember: After uninstalling a daemon agent, you must restart your system.

Deleting registry files on the UNIX or Linux operating systems

To delete registry entries after a failed installation or uninstallation on the UNIX or Linux operating systems:

- Delete the Storage Resource agent entry that is associated with the Tivoli Storage Productivity Center installation on the local server.

The terms *Server A* and *Server B* are used in this procedure to denote the different servers.

When you install Tivoli Storage Productivity Center on Server A, a Storage Resource agent is also installed on Server A. When you deploy another Storage Resource agent from Tivoli Storage Productivity Center Server B to Server A, an entry is added to the file `/etc/Tivoli/TSRM/registryNA` on Server A. This entry is associated with the Storage Resource agent deployed from Server B.

If a failure occurs when you install or uninstall the Storage Resource agent from Server A, you must delete the entry from the `/etc/Tivoli/TSRM/registryNA` file that is associated with the agent on Server A. For example, you would delete these lines from the file:

```
Instance=1
Home=/opt/IBM/TPC
Port=9567
```

- If the agent was a daemon, verify that the process is running and stop it.

- To list all the running processes, run the following command:

```
ps -ef
```

- To stop a process, run the following command:

```
kill -9 process_number
```

- For daemon agents, verify that the associated CAP File was deleted.

The CAP file is located in the `/opt/IBM/CAP` directory. The naming convention for this file is `srainstance_default.xml`.

Manually uninstalling Tivoli Storage Productivity Center components on AIX or Linux operating systems

If the Tivoli Storage Productivity Center uninstallation process fails, you must manually uninstall the components.

This procedure assumes the following conditions:

- You installed Tivoli Storage Productivity Center in the default directory.
- You used default ports.
- The Storage Resource agent service name is *Agent1*.
- You used the default Tivoli Storage Productivity Center database repository named TPCDB.

To manually uninstall Tivoli Storage Productivity Center components:

- Stop the Deployment engine by running the following command:

```
kill -9 `ps -aef | grep acsi | grep -v grep | \
awk '{ print $2 }'` 2> /dev/null
```

2. Uninstall Tivoli Integrated Portal and the deployment engine by using the following commands:


```
. /var/ibm/common/acsi/setenv.sh
/usr/ibm/common/acsi/bin/si_inst.sh -r -f
rm -rf /var/ibm/common/acsi /usr/ibm/common/acsi
/opt/IBM/Tivoli/ /opt/IBM/tipv2/ /opt/IBM/tipv2Components/
```
3. Stop Tivoli Integrated Portal and Tivoli Common Reporting by using the following command:


```
kill -9 `ps -aef | grep tip | grep -v grep | \
awk '{ print $2 }'` 2> /dev/null
```
4. Stop the Storage Resource agent by using the following command:


```
kill -9 `ps -aef | grep /opt/IBM/TPC/agent | \
grep -v grep | awk
'{ print $2 }'` 2> /dev/null
```
5. Stop the Data server by using the following command:


```
kill -9 `ps -aef | grep data | grep -v grep | \
awk '{ print $2 }'` 2> /dev/null
```
6. Stop the Device server by using the following command:


```
kill -9 `ps -aef | grep device | grep -v grep | \
awk '{ print $2 }'` 2> /dev/null
```
7. Stop the Replication server by using the following command:


```
kill -9 `ps -aef | grep replication | grep -v grep | \
awk '{ print $2 }'` 2> /dev/null
```
8. Stop the Tivoli Storage Productivity Center GUI by running the following command:


```
kill -9 `ps -aef | grep TSRMgui | grep -v grep | \
awk '{ print $2 }'` 2> /dev/null
```
9. Delete the following files and folders:
 - `rm `find . -name *tpcdsrv* -print``
 - `rm `find . -name *itsanm* -print``
 - `rm /etc/Tivoli/TSRM/registry`
 - `rm /etc/Tivoli/TSRM/registryNA`
 - `rm -fr /etc/Tivoli/TSRM/lock`
 - `rm -fr /opt/IBM/TPC`
 - `rm -fr /root/.coi (Linux)`
 - `rm -fr /home/root/.coi (AIX)`
10. Drop the database by completing the following steps:
 - a. Source the db2profile.
 - b. Enter `db2 drop db TPCDB`.
11. Navigate to the `/db2_admin_home/db2_instance_name/` directory and ensure that the database directory is deleted.

Manually uninstalling Tivoli Storage Productivity Center components on Windows

If the Tivoli Storage Productivity Center uninstallation process fails, you must manually uninstall the components.

This procedure assumes the following conditions have been met:

- You installed Tivoli Storage Productivity Center in the default directory.
- You used default ports.

- The Storage Resource agent service name is *Agent1*.
- The service name for the Tivoli Integrated Portal service is created using the Tivoli Integrated Portal port. For example, if the port for IBM Tivoli Integrated Portal V2.2 is 9568, the service name is IBMWAS70Service - V2.2_TIPProfile_Port_9568.
- You used the default Tivoli Storage Productivity Center database repository name TPCDB.

To manually uninstall Tivoli Storage Productivity Center components:

1. Uninstall IBM Tivoli Integrated Portal and the Deployment engine:
 - a. Navigate to one of the following systems:
 - On 32-bit systems, navigate to the following directory:
C:\Program Files\IBM\Common\acsi\

and run the following command:
setenv.cmd
 - On 64-bit systems, navigate to the following directory:
C:\Program Files (x86)\IBM\Common\acsi\

and run the following command:
setenv.cmd
 - b. Navigate to C:\Program Files\IBM\Common\acsi\bin\ and run the following command:
si_inst.bat -r -f
 - c. If the directory exists, in a command prompt window, enter the following command:
rmdir /s /q "C:\Program Files\IBM\Common\acsi"
 - d. If the directory exists, in a command prompt window, enter the following command:
rmdir /s /q "C:\Program Files (x86)\IBM\Common\acsi"
 - e. Navigate to the C:\Program Files\ibm\tipv2\profiles\TIPProfile\bin folder.
 - f. Run the following command:
stopserver.bat server1
 - g. Optional: Run the following command:
sc stop "IBMWAS70Service - V2.2_TIPProfile_Port_9568"

If you run this command with a user ID that has administrative authority, you do not have to provide a user name and password.
 - h. In a command prompt window, enter the following command:
sc delete "IBMWAS70Service - V2.2_TIPProfile_Port_9568"
 - i. In a command prompt window, enter the following command:
rmdir /s/q "C:\Program Files\IBM\tipv2"
2. Uninstall Tivoli Common Reporting by entering the following commands:
sc stop "IBM Cognos Content Database"
sc delete "IBM Cognos Content Database"
rmdir /s/q "C:\Program Files\IBM\tipv2Components"
3. Uninstall the Storage Resource agent by using these commands:
sc stop Agent1
sc delete Agent1

4. Remove the Storage Resource agent registry entries by using the following commands:

- On a 32-bit system:

```
reg delete HKEY_LOCAL_MACHINE\SOFTWARE\IBM\TPCAGENT\1
```

where 1 is the instance of the Agent.

If you do not have multiple agents, you can also type the `reg delete HKEY_LOCAL_MACHINE\SOFTWARE\IBM\TPCAGENT` command. If you have multiple agents, and you run this command, all of the agents are deleted.

- On a 64-bit system:

```
reg delete HKLM\SOFTWARE\Wow6432Node\IBM\TPCAGENT\1
```

where 1 is the instance of the Agent.

If you do not have multiple agents, you can also type the `reg delete HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\IBM\TPCAGENT` command. If you have multiple agents and you run this command, all of the agents are deleted.

5. Uninstall the Data server by using these commands:

```
sc stop TSRMsrv1  
sc delete TSRMsrv1
```

6. Uninstall the Device server by using these commands:

```
sc stop "IBM WAS80 Service - TPCDeviceServer"  
sc delete "IBM WAS80 Service - TPCDeviceServer"
```

7. Uninstall the Replication server by using these commands:

```
sc stop "IBM WAS80 Service - TPCReplicationServer"  
sc delete "IBM WAS80 Service - TPCReplicationServer"
```

8. Drop the database by using this command:

```
db2cmd /C db2 drop db TPCDB
```

9. Navigate to the `C:\DB2\` directory and ensure that the database directory is deleted.

10. In a command window, enter the following command:

```
rmdir /s/q "c:\Program files\ibm\tpc"
```

11. Restart your system.

Chapter 6. Reference

agent.sh command

The **agent.sh** command lets you start, stop, and restart the Storage Resource agent. You can also display the status and version of the Storage Resource agent.

Note:

- You must have root authority to run this command.
- For Windows, the status, stop, and start functions are handled through the Windows Services panel.

```
►► agent.sh [ status ] [ start ] [ stop ] [ restart ] [ version ] [ versionall ] ◄◄
```

Parameters:

status

Displays the current status of the Storage Resource agent. The status is returned is running or not running.

start

Starts the Storage Resource agent.

stop

Stops the Storage Resource agent.

restart

Stop and then start the Storage Resource agent.

version

Displays the current version of the Storage Resource agent.

versionall

Displays the version of the Storage Resource agent and its related components.

Appendix A. Accessibility features for Tivoli Storage Productivity Center

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

The following list includes the major accessibility features in IBM Tivoli Storage Productivity Center:

- Tivoli Storage Productivity Center functions are available using the keyboard for navigation instead of the mouse. You can use keys or key combinations to perform operations that can also be done using a mouse. However, you must use the mouse to navigate the Topology Viewer and report graphs. Standard operating system keystrokes are used for standard operating system operations.
- You can use screen readers to read the user interface.
- The user interface communicates all information independently of color.
- The *IBM Tivoli Storage Productivity Center Information Center*, and its related publications are accessibility-enabled and include the following accessibility features:
 - The information center is provided in XHTML 1.0 format, which is viewable in most Web browsers. XHTML allows you to view documentation according to the display preferences set in your browser. With XHTML, you can use screen readers and other assistive technologies.
 - All documentation for Tivoli Storage Productivity Center is available in Adobe Portable Document Format (PDF) using the Adobe Acrobat Reader. You can access the PDFs from the Printable PDFs topic in the information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v59r1/index.jsp>.
 - All images are provided with alternative text, so that users with vision impairments can understand the contents of the images.

Keyboard navigation

This product uses standard Microsoft Windows navigation keys.

Displaying the interface

Click **Preferences > Look and Feel** to select how to display the Tivoli Storage Productivity Center user interface. To do so, complete the following steps:

1. Start the Tivoli Storage Productivity Center user interface.
2. Click **Preferences > Look and Feel** to change the visual appearance of the user interface to best suit your visual needs:
 - Windows Classic
 - Windows
 - CDE/Motif
 - Metal

Limitations for sight-impaired users on IBM Tivoli Storage Productivity Center for Replication

There are some limitations in the product graphical user interface (GUI) that require special attention before you start using Tivoli Storage Productivity Center for Replication with a screen reader:

Dojo components are not read by all screen readers.

The JAWS screen reader does not read some Dojo components on Internet Explorer 7. Use the command-line interface instead of the GUI with JAWS on Internet Explorer 7.

IBM and accessibility

See the IBM Human Ability and Accessibility Center Web site for more information about the commitment that IBM has to accessibility.

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

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

The following list includes the major accessibility features in Tivoli Storage Productivity Center for Replication:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

See the IBM Human Ability and Accessibility Center website at www.ibm.com/able for more information about the commitment that IBM has for accessibility.

Accessibility and keyboard shortcuts in the information center

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. Using the major accessibility features in this product, users can perform these tasks:

- Use assistive technologies, such as screen-reader software and digital speech synthesizer, to hear what is displayed on the screen. Consult the product documentation of the assistive technology for details on using those technologies with this product.
- Operate specific or equivalent features by using only the keyboard.
- Magnify what is displayed on the screen.

In addition, the documentation was modified to include the following features to aid accessibility:

- All documentation is available in HTML formats to give the maximum opportunity for users to apply screen-reader software technology.
- All images in the documentation are provided with alternative text so that users with vision impairments can understand the contents of the images.

Use the following key combinations to navigate the interface by keyboard:

- To go directly to the Topic pane, press Alt+K, and then press Tab.
- In the Topic pane, to go to the next link, press Tab.
- To go directly to the Search Results view, press Alt+R, and then press the Enter or Up-Arrow key to enter the view.
- To go directly to the Navigation (Table of Contents) view, press Alt+C, and then press the Enter or Up-Arrow key to enter the view.
- To expand and collapse a node in the navigation tree, press the Right and Left-Arrow keys.
- To move to the next topic node, press the Down-Arrow or Tab key.
- To move to the previous topic node, press the Up-Arrow key or Shift+Tab.

- To go to the next link, button, or topic node from inside on of the views, press Tab.
- To scroll all the way up or down in a pane, press Home or End.
- To go back, press Alt+Left Arrow; to go forward, press Alt+Right Arrow.
- To go to the next pane, press F6.
- To move to the previous pane, press Shift+F6.
- To print the active pane, press Ctrl+P.

Related accessibility information for sight-impaired users

The following list contains hints and tips that can help you more fully use the graphical user interface:

Drop-down lists are positioned directly over or before the radio button that activates it.

If you use a screen reader, you should be aware that there are radio buttons to activate drop-down lists for several GUI pages. The way to activate the drop-down list is by selecting the associated radio button. The drop-down list is positioned directly over or before the radio button that activates it. When you use a screen reader that processes the fields and controls of a page sequentially, you might select the radio button, but not know that the associated drop-down list has been activated. The screen reader processes inactive drop-down lists first, and then processes the next radio button. The drop-down list is activated if you select the radio button.

On the following pages, keep in mind that radio buttons activate a drop-down list:

- Administration
- ESS/DS Paths
- Sessions
- Session Details
- Storage Systems

Tables are best understood by reviewing the surrounding text and the table row and column number of the table.

On some graphical user pages, tables use the header or row ID attributes when reading a single cell. The screen reader reads the table row and column number, along with cell data. Therefore, you can infer the column header and row ID.

Experiment with and fine-tune the way your screen reader pronounces some of the product abbreviations.

Your screen reader might pronounce abbreviations as if they were words. For example, the common abbreviation for Enterprise Storage Server is ESS. Your screen reader might read ESS as the word "ess". With some screen readers you can hear alternate pronunciations. If you frequently use the software you might prefer to fine-tune such associations in your settings. When an association is created, the screen reader can recognize the abbreviation as a word. If you can add dictionary words with your screen reader, replace the capitalized character sequence with the sequence E space S space S.

Typically, this abbreviation is used in the combination form of ESS/DS. This term refers to the Enterprise Storage Server 800, the DS6000, or the DS8000.

Some decorative artifacts might persist if the cascading style sheet is disabled.

Enable cascading style sheets when possible; otherwise, some decorative elements might persist in the web browser GUI. These artifacts do not affect performance. If they become too distracting, consider using the command-line interface instead.

For efficiency, confirmation dialogs place initial focus on the Yes button.

When a confirmation dialog box is displayed, focus is given to the **Yes** button. Therefore, the screen reader reads “Yes” but does not read the confirmation text. The software processes the information in this way when you do the following types of tasks:

- Perform an action on a session
- Remove a connection to a storage system
- Click the **About link**
- Create a high-availability connection

To read the confirmation text before clicking the **Yes**, **No**, or **OK** button, view the previous heading before the button.

Dojo components are not read by all screen readers.

The Job Access for Windows and Speech (JAWS) screen reader does not read some Dojo components on Internet Explorer 7. Use the command-line interface instead of the GUI with JAWS on Internet Explorer 7.

Firefox is the preferred browser for use with a screen reader.

Use Firefox as the screen reader because other browsers might not fully expose assistive technology content to the screen reader.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
U.S.A*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE: This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Intel Xeon, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Glossary

This glossary includes terms and definitions for IBM Tivoli Storage Productivity Center.

The following cross-references are used in this glossary:

- *See* refers the reader from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- *See also* refers the reader to a related or contrasting term.

To view glossaries for other IBM products, go to: <http://www.ibm.com/software/globalization/terminology/>.

active management server

A management server from which the storage environment can be monitored and managed. The active management server replicates its database to the standby server.

business continuity

The capability of a business to withstand outages and to operate mission-critical services normally and without interruption in accordance with predefined service-level agreements.

CIM See *Common Information Model*.

CIM agent

The code that consists of common building blocks that can be used instead of proprietary software or device-specific programming interfaces to manage devices that are compliant with the Common Information Model (CIM).

CKD See *count key data*.

cluster

1. In SAN Volume Controller or Storwize V7000, a pair of nodes that provides a single configuration and service interface.
2. In IBM System Storage™ DS8000®, a partition capable of performing all DS8000 series functions. With two clusters in the DS8000 storage unit, any operational cluster can take over the processing of a failing cluster.

Common Information Model (CIM)

An implementation-neutral, object-oriented schema for describing

network management information. The Distributed Management Task Force (DMTF) develops and maintains CIM specifications.

community name

The part of a Simple Network Management Protocol (SNMP) message that represents a password-like name and that is used to authenticate the SNMP message.

consistency group

A group of copy relationships between virtual volumes or data sets that are maintained with the same time reference so that all copies are consistent in time.

copy set

The set of source volumes or target volumes involved in a FlashCopy operation.

count key data (CKD)

In mainframe computing, a data-record format employing self-defining record formats in which each record is represented by up to three fields: a count field identifying the record and specifying its format, an optional key field that can be used to identify the data area contents, and an optional data field that typically contains the user data.

data collection

The process of obtaining performance and availability monitoring data and providing that data to a metric evaluator. Examples of data collectors include Domain Name System (DNS) probes, web page analyzers, or database analyzers. See also *discovery*.

discovery

The process of finding resources within an enterprise, including finding the new location of monitored resources that were moved. See also *schedule*.

enterprise repository

A component of the Data server that records and stores all information about

the monitored computers' storage assets and their usage over time. The repository is organized into relational database tables and is accessed by Data server using Java Database Connectivity (JDBC).

extended count key data (ECKD™)

An extension of the count key data (CKD) architecture.

FlashCopy

In storage systems, a feature that creates a point-in-time copy of a volume or set of volumes on a specified target volume.

flush-through mode

A process in which data is written to a storage device at the same time as the data is cached.

freeze An operation in which a storage system blocks I/O from the host system to the affected volumes on the primary site. A freeze operation stops mirroring between the primary and secondary volumes to ensure data consistency at the secondary site. See also *thaw*.

GiB See *gibibyte*.

gibibyte (GiB)

A base-2 unit of measurement that is equal to 1,073,741,824 bytes (2 to the 30th power). See also *gigabyte*.

gigabyte (GB)

In decimal notation, 1,073,741,824 bytes, when referring to memory capacity; in all other cases, it is defined as 1,000,000,000.

Global Copy

A non-synchronous long-distance copy option for data migration and backup. See also *remote mirror and copy* and *Metro Mirror*.

Global Mirror

An optional capability of the remote mirror and copy feature that provides a two-site extended-distance remote copy. Data that is written by the host to the storage unit at the local site is automatically maintained at the remote site. See also *Metro Mirror* and *remote mirror and copy*.

Globally Unique Identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system.

grain size

The unit size for allocating space on thin-provisioned volumes, such as 32, 64, 128, and 256 kibibyte. The grain size is defined when a volume is created.

GUID See *Globally Unique Identifier*.

Hardware Management Console (HMC)

In a system storage environment, a system that acts as the focal point for configuration, management of Copy Services functions, and maintenance.

heat map

A color-coded data chart where colors are used to differentiate values in a data set.

high availability

Pertaining to a clustered system that is reconfigured when node or daemon failures occur, so that workloads can be redistributed to the remaining nodes in the cluster.

high-availability solution

A combination of hardware, software, and services that fully automates the recovery process and does not disrupt user activity. High-availability solutions must provide an immediate recovery point with a fast recovery time.

HMC See *Hardware Management Console*.

host volume

A volume that represents the volume functional role from an application point of view. The host volume can be connected to a host or server. It receives read, write, and update application I/O, depending on the site to which the application is writing.

in-band discovery

The process of discovering information about the storage area network (SAN), including topology and attribute data, through the Fibre Channel data paths. Contrast with *out-of-band discovery*.

job

A unit of work within a run for a schedule. One run can have multiple jobs. For example, a system probe schedule has one job created per system in its definition.

job scheduler

A component of the Data server that deploys all monitoring activities.

journal volume

A volume that holds a consistent copy of data until a new consistent copy is formed. The journal volume restores the last consistent point during a recovery.

KiB See *kibibyte*.

kibibyte (KiB)

A base-2 unit of measurement that is equal to 1,024 bytes (2 to the 10th power). See also *kilobyte*.

kilobyte

For processor storage, real and virtual storage, and channel volume, 2 to the power of 10 or 1,024 bytes. For disk storage capacity and communications volume, 1,000 bytes.

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier that is used to differentiate devices, each of which is a logical unit (LU).

LUN See *logical unit number*.

Management Information Base (MIB)

In the Simple Network Management Protocol (SNMP), a database of objects that can be queried or set by a network management system.

management server

A system that provides a central point of control for managing data replication.

management server relationship

A connection between two replication servers, where one server acts as the active server and replicates the data that is necessary for the standby server to take control of the replication environment.

mebibyte (MiB)

A base-2 unit of measurement that is equal to 1,048,576 bytes (2 to the 20th power). See also *megabyte*.

megabyte (MB)

For processor storage, real and virtual storage, and channel volume, 2 to the 20th power or 1,048,576 bytes. For disk storage capacity and communications volume, 1,000,000 bytes.

Metro Mirror

A method of synchronous remote replication that operates over distances up to 300 kilometers. This method constantly

updates a secondary copy of a volume to match changes that are made to a source volume. See also *remote mirror and copy* and *Global Copy*.

Metro Global Mirror

A three-site, high-availability, disaster recovery solution. Metro Global Mirror uses synchronous replication to mirror data between a local site and an intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site.

MiB See *mebibyte*.

MIB See *Management Information Base*.

native interface

An interface that is specific to a system or subsystem.

out-of-band discovery

The process of discovering storage area network (SAN) information, including topology and device data, without using the fibre-channel data paths. A common mechanism for out-of-band discovery is the use of Simple Network Management Protocol (SNMP) Management Information Base (MIB) queries, which are invoked over an IP network. Contrast with *in-band discovery*.

ping

A job that tracks the availability of assets and that is performed by an agent. Several ping jobs can be used to monitor the availability of any computer or subset of computers in the network. See also *schedule*.

pool

A storage pool or pool is a grouping of storage space that consists of volumes, logical unit numbers (LUNs), or addresses that share a common set of administrative characteristics.

practice volume

A volume that can be used to test disaster-recovery actions while maintaining disaster-recovery capability.

primordial pool

Storage capacity that is deallocated on a storage device. Storage pools are created by allocating storage capacity from primordial pools.

probe A job that itemizes and creates an

inventory of assets, such as computers, controllers, disk drives, file systems, and logical units. See also *schedule*.

recovery point objective

The maximum amount of data loss that can be tolerated during a service interruption.

remote mirror and copy

A feature of a storage server that constantly updates a secondary copy of a logical volume to match changes made to a primary logical volume. The primary and secondary volumes can be on the same storage server or on separate storage servers. See also *Metro Mirror* and *Global Mirror*.

resource

In a storage environment, an entity that is monitored. Resources can include fabrics, switches, computers, storage systems, and tape libraries.

role

A job function that identifies the tasks that a user can perform and the resources to which a user has access. A user can be assigned one or more roles.

role pair

The association of two roles in a replication session that take part in a copy relationship. For example, in a Metro Mirror session, the role pair can be the association between the volume roles of Host1 and Host2.

scan job

A job that monitors the usage and consumption of storage. See also *schedule*.

schedule

A monitoring task that collects data about storage resources. Monitoring tasks include scans, probes, pings, discovery, and performance monitors. For example, define a probe schedule to collect data about storage systems; define a scan schedule to collect file-level information about hosts.

session

A collection of volumes within a logical subsystem that are managed together during the creation of consistent copies of data. All volumes in a session must transfer their data successfully to the remote site before the increment can be called complete.

site awareness

The association of a location (site) with each storage system and each session's volume role. When a copy set is added, only the storage systems with locations that match those locations of the volumes' roles can be selected. This constraint prevents the reversing of a hardware relationship or the selection of a volume at the wrong location.

SMI-S See *Storage Management Initiative Specification*.

standby management server

A management server that is a backup for the active server. The replication environment cannot be monitored or managed from the standby server.

Storage Management Initiative Specification (SMI-S)

A design specification developed by the Storage Networking Industry Association (SNIA) that specifies a secure and reliable interface with which storage management systems (SMSs) can identify, classify, monitor, and control physical and logical resources in a storage area network (SAN). The interface integrates the various devices to be managed in a SAN and the tools used to manage them.

storage pod

A logical entity of components of a system that consists of two storage nodes, and one or two storage subsystems directly connected with these storage nodes.

storage resource group

A named collection of logically related resources that are monitored by Tivoli Storage Productivity Center. Monitored resources can include fabrics, switches, computers, storage systems, tape libraries, and other storage resource groups.

target volume

A volume that receives data from a host volume or another intermediate volume. It is used only in FlashCopy sessions.

thaw

An operation in which a storage system releases the block of the I/O from the host system to the affected volumes on the primary site. A thaw operation can

occur after a freeze operation ends and consistency is formed on the secondary site. See also *freeze*.

track space-efficient (TSE) volume

A volume in which storage space is allocated on an as-needed basis by using space on the target volume only when tracks are copied from the source volume to the target volume.

virtualization

The substitution of virtual resources for actual resources, where the virtual resources have the same functions and external interfaces as their counterparts, but differ in attributes, such as size, performance, and cost. Virtualization is commonly applied to physical hardware resources by combining multiple physical resources into shared pools from which users receive virtual resources.

virtual storage area network (VSAN)

A fabric within the storage area network (SAN).

VSAN See *virtual storage area network*.

zone A logical grouping of switches, switch ports, and their attached devices in a fabric.

zone alias

A name that is given to a collection of one or more zone members that you want to manage together.

zone set

A group of zones that function together on a fabric.

Index

Numerics

64-bit environment supported 23

A

about this document xiv
accessibility features
 for Tivoli Storage Productivity Center 393
accessibility features for users with disabilities 395
adding
 installation license 167, 168
agent.sh command
 syntax 391
agents
 assigning agents for scan and probe jobs 276
 configuration file 280
 configuring for Storage Resource agent upgrade 295
 deploying Storage Resource agents 281
 deployment considerations for Storage Resource agents 282
 hardware requirements 3
 overview 220
 planning for CIM agents 68
 planning for Storage Resource agents 44
 planning for TagmaStore CIM agents 69
 preparing Storage Resource agents for installation 169, 282
 upgrading Storage Resource agent 364
AIX 386
alerts, overview 223
as a stand-alone system 161
assign
 roles 220

B

back up
 migrating 349
Baroc file location 274
Brocade
 planning for Brocade Network Advisor 88
 planning for DCFM 88, 89
Brocade Network Advisor
 planning for 88
browsers supported 13

C

CAP file 127, 128

certificates
 creating for Storage Resource agents 286
 replacing for Storage Resource agents 291
changing timeout value for web-based GUI 338
CIM agents
 planning for 68
 planning for TagmaStore 69
 software requirements 15, 51
cluster resource group
 for PowerHA 109
cluster resource groups 107
commands
 agent.sh 391
comments, sending xiv
common user, about 27, 201
components 386, 387
configuration
 analysis tool 54
 authorizing users 266
 editing nas.config file 280
 for Device server 96
 for switches 304
 guidelines 25
 History Aggregator 280
 history settings 303
 history tool 56
 NetApp 114
 overview 220
 planning for Fabric Manager 97
 planning for HBAs 100
 planning for roll up reports 53
 planning for VMware 120
 Removed Resource Retention 301
 Resource History Retention 300
 role-to-group mapping 266
 the Device server 304
configuration files
 agent.config file 279
 server.config file 279
configuration utility
 accessing Tivoli Storage Productivity Center for Replication 239
 how to create alert definitions 237
 how to display availability reports 235
 how to manage tape library reports 238
 how to monitor storage assets, storage systems and fabrics 235
 learn how to manage element managers and external tools 238
 learn how to manage servers and agents 235
 show managed devices 235
 tasks 234
configuring 204
 Data server for Windows domain 310, 311

configuring (*continued*)
 inactive zone sets for fabric data source 306
 Monitoring agent 332
 Storage Resource agent for upgrade 295
credentials, migrating 345

D

Data agents
 migrating using upgrade wizard 368
 uninstalling using GUI 381
Data Manager
 agent.config file 279
 databases that can be monitored 39
Data Path Explorer view 63
database repository
 planning to retrieve information from DB2 views 124, 125
 software requirements 12
databases
 histories, specifying how long to keep 302
 specifying how long to keep information 303
 that Data Manager can monitor 39
DB2
 installing on AIX 186, 189
 installing on Linux 184
 installing on UNIX 184
 installing on Windows 182
 monitoring the DB2 environment 126
 preparing to install on UNIX 181
 preparing to install on Windows 181
 uninstalling 383
 upgrading 345
 verifying installation 191
 verifying installation using CLP 191
 verifying installation using First Steps 191
DCFM
 planning for 88, 89
deleting duplicate DB2 entries 363
Device server
 configuration 96, 304
 setting timeout values 306
disabilities, accessibility features for users with 395
Discovery Library Adapter 126
domain account
 how to grant local administrative privileges 311
drop-down lists, limitations of to the sight-impaired 395
DS8000
 how to manage DS8000 element managers 238
 planning for 73
 planning for performance 74

E

editing configuration files 279
EMC PowerPath Multipathing
 planning for 43

F

Fabric agents
 migrating using upgrade wizard 368
fabric management
 planning for
 Brocade 86
 QLogic and Cisco 89
Fabric Manager
 directors 99
 planning for 85
 planning for configuration 97
 planning for HBAs 100
 switches 99
fabrics 56
federated repositories
 adding an LDAP repository 206
 removing an LDAP repository 213
file systems
 supported 40
fully qualified host names
 host names
 fully qualified host names
 required 135
 required 135

G

glossary 403

H

hardware requirements 1
 Tivoli Storage Productivity Center for
 Replication 1
 Tivoli Storage Productivity Center
 GUI 4
HBA 56
HBAs
 planning for Fabric Manager 100
History Aggregator, configuring 280
host port 56
hosts 56

I

installation
 components 133
 reinstalling 192
 silent mode, multiple-server
 environment 154, 157
 silent mode, single-server
 environment 142
 upgrading using upgrade
 wizard 352, 355, 356
installation failure
 deleting registry files 385
installing 161
 agents 168
 DB2 on AIX 186

installing (*continued*)
 DB2 on AIX POWER7 Systems 189
 DB2 on Linux 184
 DB2 on UNIX 184
 DB2 on Windows 182
 DB2, preparing for installation on
 UNIX 181
 DB2, preparing for installation on
 Windows 181
GUI
 TIP 135
Monitoring agent 173
new installation 135
planning 130
single-server environment 141
starting installation 138
Storage Resource agent
 considerations 282
Storage Resource agents 281
Storage Resource agents locally 169
Storage Resource agents
 remotely 169, 282
TIP
 installing 135
Tivoli Integrated Portal 369
wizard
 remote database 150

IPv6

configuring
 AIX for IPv6 337
 DB2 on AIX for IPv6 337
 DB2 on Linux for IPv6 338
configuring DB2 on Linux 338
configuring for AIX 337
configuring for DB2 on AIX 337
limitations 48
planning for 46

J

jobs, overview 221

L

LDAP
 planning for 50
LDAP repository
 adding to federated repositories 206
 removing 213
LDAP servers
 software requirements 15
license
 adding 167, 168
licenses 22
Linux 386
log files
 installation error 166
 setting retention period 275

M

MIB file location 274
migrating
 Data agents 366
 Fabric agents 366

monitoring
 databases 39
 file systems 40
 MSCS clustered nodes and cluster
 resource groups 107
 networked file systems 41
 PowerHA clustered nodes and cluster
 resource groups 111
 volume managers 41
Monitoring agent
 configuring 332
 configuring for Windows 332
 configuring on UNIX 333
 installing 173
 installing and configuring 173, 176
 starting 333
 stopping 333
Monitoring Agent
 uninstalling 382
 uninstalling on UNIX 382
 uninstalling on Windows 382
MSCS support 106, 107
multipath 56
multipath subsystem device drivers
 planning for 42
 upgrading 43

N

NAS
 monitor 277
 support of 113
NAS configuration file, editing 280
nas.config file, editing 280
native interfaces
 planning for 48
NetApp
 configuration 114
 identifying in user interface 116
 limitations 116
 planning for 114
 removing a device 115
networked file systems
 supported 41

O

operating systems supported 4

P

passwords
 changing using password tool 225,
 227, 229, 231
 valid characters 29
path planning 56
performance
 configuration 104
 planning for DS8000 74
performance metrics 102
performance monitor 56
planning
 authorization for users 27, 201
 for an MSCS environment 107
 for CIM agents 68
 for Data Path Explorer view 63

- planning (*continued*)
 - for DS8000 73
 - for MSCS support 106
 - for NAS system support 113
 - for retrieving information from DB2 views 124, 125
 - for roll up reports configuration 53
 - for rollup reports 52
 - for SAN configuration tools 54
 - for Solaris Storage Resource agents 66
 - for System Storage N Series Gateway servers 117
 - for TagmaStore CIM agents 69
 - for tape library support 38
 - for VMware configuration 120
 - for VMware support 117
 - for volume manager support 41
 - for XIV system 83
 - IPv6 support 46, 48
 - NetApp support 114
 - overview 1
 - storage system support 37
 - switches 38
 - Virtual I/O Server support 122
- PowerHA
 - environment 111
 - planning for 109
 - upgrading agents 111
- probe 56

Q

- quota and constraint e-mail address rules 276

R

- RAID 56
- reader feedback, sending xiv
- replication planning 56
- Replication server
 - upgrading 344
- reports
 - configuring for data aggregation 280
 - specifying how long to keep information 303
 - VMware capacity 120
- resource histories
 - setting retention periods 300
- resources
 - specifying how long to keep 301
- roles 219
- roll up reports
 - planning for configuration 53
- rollup reports
 - planning for 52

S

- SAN configuration tools
 - planning 54
- SAN Planner 56
- SAN Volume Controller
 - publications xi
- SDD 56

- sending comments xiv
- sight-impaired 395
- silent mode installation
 - multiple-server environment 154, 157
 - single-server environment 142
- silent-mode upgrade 352, 357, 360
- single sign-on
 - planning for 52
 - user authentication method used with 52
- SNMP
 - sending traps 273
- software requirements
 - database repository 12
 - for CIM agents 15, 51
 - for LDAP servers 15
 - operating systems 4
- SSH protocol
 - creating a certificate 286
- SSL certificates
 - replacing a certificate 291
- starting
 - through Tivoli Integrated Portal
 - Tivoli Storage Productivity Center 200
 - Tivoli Storage Productivity Center for Replication 200
 - Tivoli Storage Productivity Center 195
 - Tivoli Storage Productivity Center for Replication 199
 - Tivoli Storage Productivity Center stand-alone GUI 195
 - Tivoli Storage Productivity Center web-based GUI 198
- storage capacity of volumes 73
- Storage Optimizer
 - analysis report 62
 - optimization report 62
 - supported storage systems 62
- storage pools 56
- storage port 56
- Storage Resource agents
 - configuring for upgrade 295
 - creating certificates for 286
 - deploying 281
 - deployment considerations 282
 - importing authentication information 321
 - installing 281
 - installing locally 169
 - monitoring by Tivoli Agent Management Services 127, 128
 - planning for 44
 - planning for Solaris 66
 - preparing for 169
 - protocol support 44
 - replacing SSL certificates 291
 - uninstalling manually 381
 - uninstalling using GUI 381
 - upgrading 364
 - upgrading manually 365, 366
- storage resource groups 56
- Storage Subsystem Credential Migration tool 345
- storage system supported 37

- subsystem device drivers
 - planning for 42
 - upgrading 43
- switch 56
- switch zone 56
- switches
 - configuring 304
 - planning for 99
 - supported 38
 - using private networks 101
- System Storage N Series Gateway servers
 - planning for 117
 - upgrading 370

T

- tape library support 38
- TCP/IP ports
 - planning for 33, 34
 - used by Tivoli Storage Productivity Center for Replication 34
- TEC events
 - sending 273
- time zone 56
- timeout values
 - setting for Device server 306
- Tivoli Agent Management Services
 - monitoring Storage Resource agents 127, 128
- Tivoli Common Reporting
 - installing and configuring 161
- Tivoli Integrated Portal 220
 - Tivoli Common Reporting 219
- Tivoli Storage Productivity Center
 - hardware requirements 1
 - installing 135
 - Lightweight Directory Access Protocol SAN Volume Controller remote authentication support 329
 - software requirements 4
 - uninstalling 375, 376, 378
 - uninstalling Monitoring Agent 382
 - uninstalling Monitoring Agent on UNIX 382
 - uninstalling Monitoring Agent on Winsows 382
 - uninstalling the agents 380
- Tivoli Storage Productivity Center for Replication
 - accessing through Tivoli Storage Productivity Center 239
 - software requirements 4
 - uninstalling 375, 376, 378
- topology 56
- trademarks 401
- translations
 - browser locale requirement xiii

U

- uninstall 386, 387
- uninstalling
 - Data agent 381
 - DB2 383
 - Storage Resource agent 381
 - Storage Resource agent manually 381

- uninstalling *(continued)*
 - Tivoli Storage Productivity Center 375, 376, 378
 - Tivoli Storage Productivity Center agents 380
 - Tivoli Storage Productivity Center for Replication 375, 376, 378
 - Tivoli Storage Productivity Center Monitoring Agent 382
 - Tivoli Storage Productivity Center Monitoring Agent on UNIX 382
 - Tivoli Storage Productivity Center Monitoring Agent on Windows 382
- upgrading 363
 - in silent mode 352, 357, 360
 - multiple-server environment 354
 - Replication server 344
 - single-server environment 351
 - Storage Resource agents
 - manually 365, 366
 - System Storage N Series Gateway servers 370
- user authentication 28, 204
 - DB2 user names and passwords 28
 - general requirements 28
 - work sheets 31
- user names
 - valid characters 29
- zone sets
 - configuring for inactive zones
 - ets 306

V

- views
 - planning to retrieve information 124, 125
- Virtual I/O Server
 - planning for 122
- VMware
 - capacity reports 120
- VMware support 117
 - configuration 120
- volume managers
 - supported 41
- volume planning 56
- VPA 56

W

- windows 387
- Windows domain
 - adding Tivoli Storage Productivity Center 312
 - configuring Data server for 310, 311
- workload profile 56
- WWPN 56

X

- XIV system
 - planning for 83

Z

- zone 56
- zone planning 56
- zone set 56



Product Number: 5725-F93, 5725-G33, 5725-F98

Printed in USA

SC27-4047-00

