IBM Tivoli Storage Productivity Center

Problem Determination Guide

Version 4.1

GC27-2342-00

IBM Tivoli Storage Productivity Center

Problem Determination Guide

Version 4.1

GC27-2342-00

Note:
Before using this information and the product it supports, read the information in "Notices" on page 93.

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

Contents

Preface

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

Who should read this guide

This publication is intended for administrators or users who are troubleshooting problems with IBM Tivoli Storage Productivity Center.

Administrators and users should be familiar with the following topics:

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

If you are using these storage systems or applications, you should also be familiar with these products:

- IBM® System Storage™ DS8000®
- IBM System Storage DS6000™
- IBM System Storage DS5000
- IBM System Storage DS4000®
- IBM System Storage DS3000
- IBM XIV Storage System
- IBM System Storage SAN Volume Controller

Publications

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

The publications are available from the IBM publications center at http://www.ibm.com/shop/publications/order

Tivoli Storage Productivity Center publications

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

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

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

Publication Title	Order Number
IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide	SC27-2337
IBM Tivoli Storage Productivity Center User's Guide	SC27-2338
IBM Tivoli Storage Productivity Center Messages	SC27-2340
IBM Tivoli Storage Productivity Center Command-Line Interface Reference	SC27-2339
IBM Tivoli Storage Productivity Center Problem Determination Guide	GC27-2342
IBM Tivoli Storage Productivity Center Workflow User's Guide	SC27-2341

Tivoli Storage Productivity Center for Replication publications

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

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

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

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

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

System Storage Productivity Center publications

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

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

Click **System Storage Productivity Center**.

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

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

System Storage DS3000 publications

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

To see the IBM DS3000 publications, follow these steps:

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

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

System Storage DS4000 and DS5000 publications

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

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

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

Publication Title	Order Number
IBM System Storage DS4000 Concepts Guide	GC26-7734

Publication Title	Order Number
IBM System Storage DS4000/DS5000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview	GC53-1137
IBM System Storage DS Storage Manager Version 10 IBM System Storage DS Storage Manager Installation and Host Support Guide	GC53-1135
IBM System Storage DS Storage Manager Version 10.50 Copy Services User's Guide	GC53-1136
IBM System Storage DS4800 Storage Subsystem Quick Start Guide	GC27-2148
DS5100, DS5300 and EXP5000 Quick Start Guide	GC53-1134
IBM System Storage DS4800 Storage Subsystem Installation, User's, and Maintenance Guide	GC26-7845
IBM System Storage DS5100 and DS5300 Storage Subsystems Installation, User's, and Maintenance Guide	GC53-1140
IBM System Storage DS4000/DS5000 Hard Drive and Storage Expansion Enclosure Installation and Migration Guide	GC53-1139
IBM System Storage DS5000 EXP5000 Storage Expansion Enclosure Installation, User's, and Maintenance Guide	GC53-1141
IBM System Storage DS3000, DS4000, and DS5000 Command Line Interface and Script Commands Programming Guide	GC52-1275
IBM System Storage DS4000/DS5000EXP810 Storage Expansion Enclosure Installation, User's and Maintenance Guide	GC26-7798

System Storage DS6000 publications

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

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

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

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

System Storage DS8000 publications

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

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

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

System Storage DS Open Application Programming Interface publications

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

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

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

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

System Storage SAN Volume Controller publications

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

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

Publication Title	Order Number
IBM System Storage SAN Volume Controller CIM Agent Developer's Guide	SC26-7904
IBM System Storage SAN Volume Controller Command-Line Interface User's Guide	SC26-7903
IBM System Storage SAN Volume Controller Software Installation and Configuration Guide	SC23-6628
IBM System Storage SAN Volume Controller Host Attachment Guide	SC26-7905
IBM System Storage SAN Volume Controller Planning Guide	GA32-0551
IBM System Storage SAN Volume Controller Troubleshooting Guide	GC27-2227
IBM System Storage SAN Volume Controller Hardware Maintenance Guide	GC27-2226

Publication Title	Order Number
IBM System Storage SAN Volume Controller Model 2145-8G4 Hardware Installation Guide	GC27-2220
IBM System Storage SAN Volume Controller Model 2145-8A4 Hardware Installation Guide	GC27-2219
IBM System Storage SAN Volume Controller Model 2145-4F2 Hardware Installation Guide	GC27-2222
IBM System Storage SAN Volume Controller Models 2145-8F2 and 8F4 Hardware Installation Guide	GC27-2221

XIV Storage System publications

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

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

DB2 Database for Linux, UNIX, and Windows publications

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

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

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

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

IBM International Technical Support Organization publications

The IBM International Technical Support Organization (ITSO) publishes IBM Redbooks®, which are books on specialized topics.

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

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

For information about IBM Tivoli Storage Productivity Center, see *IBM Tivoli Storage Productivity Center V4.1 Update Guide*. Search for **SG24-7725**.

Note: This publication is not currently available. Check back at a future time for this publication.

Translations

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

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

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

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

Accessing publications online

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

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

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

Note: If you print PDF documents on other than letter-sized paper, select the **Fit to page** check box in the **Adobe Acrobat Print** dialog. This option is available when you click **File → Print**. **Fit to page** ensures that the full dimensions of a letter-sized page print on the paper that you are using.

Ordering publications

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

You can order many IBM publications online at http://www.ibm.com/shop/publications/order.

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

Providing feedback about publications

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

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

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

Contacting IBM Support Center

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

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

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

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

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

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

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

Reporting a problem

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

Have the following information ready when you report a problem:

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

Conventions used in this guide

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

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

The following typeface conventions are used in this publication:

Bold

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

Italic

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

monospace

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

For syntax notation, these conventions are used:

- < > (less than, greater than symbols) are used to indicate a variable value. Do not type the < > symbols.
- # is the prompt for the root user on UNIX platforms.
- Uppercase and lowercase characters do matter. Type in commands exactly as shown.

New for IBM Tivoli Storage Productivity Center Version 4.1

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

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

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

Name change

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

Licensing changes

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

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

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

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

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

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

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

Integration features

Tivoli Storage Productivity Center provides these integration features.

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

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

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

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

This integration enables you to:

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

Integration of Tivoli Storage Productivity Center and IBM Tivoli Integrated Portal

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

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

Launch in context feature

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

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

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

There are three levels of launch in context ability:

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

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

Launch with single sign-on

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

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

Single sign-on

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

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

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

Tivoli Storage Productivity Center uses Lightweight Third Party Authentication (LTPA) tokens to pass the user information between applications. To use LTPA tokens for single sign-on, each participating application must possess the same set of keys to encode and decode the user information contained in the token. As

an additional security feature, the LTPA tokens expire after a determined amount of time. When the tokens expire, you must re-enter your user ID and password information.

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

Storage Resource agents

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

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

You can use the following functions:

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

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

SQL access to Tivoli Storage Productivity Center database

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

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

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

Storage entity views

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

Entities defined by Tivoli Storage Productivity Center

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

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

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

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

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

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

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

A storage resource group is a set of entities managed by Tivoli Storage Productivity Center. These entities can be servers, switches, storage subsystems, fabrics, storage pools, and storage volumes. Storage resource groups can be a group of heterogeneous objects and can also contain other storage resource groups without any connectivity.

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

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

Storage resource groups also work with these profiles:

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

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

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

IBM General Parallel File System™

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

Installation changes

IBM Tivoli Storage Productivity Center for Replication

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

IBM DB2 Database for Linux, UNIX, and Windows

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

Installation of IBM Tivoli Integrated Portal

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

Embedded WebSphere® 6.1 and JRE 1.5

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

Silent installation

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

New device and application support

IBM System Storage DS8000 4.2

This release supports DS8000 4.2 with these additional items:

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

IBM System Storage SAN Volume Controller 4.3.1

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

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

Microsoft SQL Server 2005 and Microsoft SQL Server 2008 databases

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

EMC PowerPath

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

EMC PowerPath version 4.0 or later is supported.

Network Appliance (NetApp)

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

IBM XIV Storage System

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

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

Multipath subsystem device drivers

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

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

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

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

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

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

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

Chapter 1. Introduction to problem determination

Troubleshooting is a systematic approach to solving a problem. The goal of troubleshooting is to determine why something does not work as expected and explain how to resolve the problem.

Overview

The first step in the troubleshooting process is to describe the problem completely. Problem descriptions help you and the IBM Support person know where to start too find the cause of the problem. This step includes asking yourself basic questions:

- What are the symptoms of the problem?
- Where does the problem occur?
- When does the problem occur?
- Under which conditions does the problem occur?
- Can the problem be reproduced:

The answers to these questions typically lead to a good description of the problem, and that is the best way to start down the path of problem resolution.

What are the symptoms of the problem?

When starting to describe a problem, the most obvious question is "What is the problem?" This might seem like a straightforward question; however, you can break it down into several more-focused questions that create a more descriptive picture of the problem. These questions can include:

- Who, or what, is reporting the problem?
- What are the error codes and messages?
- How does the system fail? For example, is it a loop, hang, crash, performance degradation, or incorrect result?
- What is the business impact of the problem?

Where does the problem occur?

Determining where the problem originates is not always easy, but it is one of the most important steps in resolving a problem. Many layers of technology can exist between the reporting and failing components. Networks, disks, and drivers are only a few components to consider when you are investigating problems.

The following questions help you to focus on where the problem occurs to isolate the problem layer:

- Is the problem specific to one platform or operating system, or is it common across multiple platforms or operating systems?
- Is the current environment and configuration supported?
- Is the application running locally on the database server or on a remote server?
- Is there a gateway involved?
- Does the database reside on a local or remote computer?

Remember that if one layer reports the problem, the problem does not necessarily originate in that layer. Part of identifying where a problem originates is understanding the environment in which it exists. Take some time to completely describe the problem environment, including the operating system and version, all corresponding software and versions, and hardware information. Confirm you are running within an environment that is a supported configuration; many problems can be traced back to incompatible levels of software that are not intended to run together or have not been fully tested together.

When does the problem occur?

Develop a detailed timeline of events leading up to a failure, especially for those cases that are one-time occurrences. Yo can most easily do this by working backward: Start at the time an error was reported (as precisely as possible, even down to the millisecond), and work backward through the available logs and information. Typically you need to look only as far as the first suspicious event that you find in any diagnostic log; however, this is not always easy to do and takes practice. Knowing when to stop looking is especially difficult when multiple layers of technology are involved, and when each has its own diagnostic information.

To develop a detailed timeline of events, answer these questions:

- Does the problem happen only at a certain time of day or night?
- How often does the problem happen?
- What sequence of events leads up to the time that the problem is reported?
- Does the problem happen after an environment change such as upgrading or installing software or hardware?

Responding to questions like this helps you with a frame of reference in which to investigate the problem.

Under which conditions does the problem occur?

Knowing which systems and applications are running at the time that a problem occurs is an important part of troubleshooting. These questions about your environment can help you to identify the root cause of the problem:.

- Does the problem always occur when the same task is being performed?
- Does a certain sequence of events need to occur for the problem to surface?
- Do any other applications fail at the same time?

Answering these types of questions can help you explain the environment in which the problem occurs, and correlate any dependencies. Remember that just because multiple problems might have occurred around the same time, the problems are not necessarily related.

Can the problem be reproduced?

From a troubleshooting standpoint, the *ideal* problem is one that can be reproduced. Typically, problems that can be reproduced have a larger set of tools or procedures at your disposal to help you investigate. Consequently, problems that you can reproduce are often easier to debug and solve. However, problems that you can reproduce can have a disadvantage: If the problem is of a significant

business impact, you do not want it to recur. If possible, recreate the problem in a test or development environment, which typically offers you more flexibility and control during your investigation.

However, reproducible problems can have a disadvantage: if the problem is of significant business impact, you do not want it recurring. If possible, recreating the problem in a test or development environment is often preferable in this case.

- Can the problem be recreated on a test system?
- Are multiple users or applications encountering the same type of problem?
- Can the problem be recreated by running a single command, a set of commands, or a particular application?

Tivoli Storage Productivity Center problem determination resources

A wide variety of troubleshooting and problem determination information is available to assist you in using Tivoli Storage Productivity Center.

Tivoli Storage Productivity Center documentation

Troubleshooting information can be found in the Information Center. For information about theTivoli Storage Productivity Center library, see http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp.

Tivoli Storage Productivity Center Technical Support Web site

Refer to the Tivoli Storage Productivity Center Technical Support Web site if you are experiencing problems and want help finding possible causes and solutions. The Technical Support site has links to the latest Tivoli Storage Productivity Center publications, technical notes, Authorized Program Analysis Reports (APARs), fix packs, and other resources. You can search through this knowledge base to find possible solutions to your problems.

Check the following Web site: http://www.ibm.com/software/sysmgmt/products/support/IBMTotalStorageProductivityCenterStandardEdition.html.

Chapter 2. Tivoli Storage Productivity Center tools

This section describes the tools that are available to help you diagnose problems with Tivoli Storage Productivity Center. The tools are Agent Manager toolkit, Service tool, Repocopy tool, and Tracing tool.

Agent Manager toolkit for administrators

In addition to the commands in <Agent_Manager_install_dir>, the Agent Manager provides a toolkit of "as-is" administration tools. The toolkit commands are in the following directory: <Agent_Manager_install_dir>/toolkit/bin. Use caution when using these tools because they have not been translated and the output of the commands are not translated. They are also subject to change in the future.

The commands provided are:

HealthCheck

HealthCheck is a command to verify the state of the basic functions of the Agent Manager. This command indicates if the Agent Manager is or is not operational. This command must be run from the machine on which the Agent Manager is installed. This command requires certain parameters depending on if you used the installation defaults and if you are using generated or demonstration certificates. The files are:

HealthCheck.sh
The script that runs the **HealthCheck** command on UNIX and Linux systems.

HealthCheck.bat
The script that runs the **HealthCheck** command on Windows systems.

HealthCheck.jar
The JAR file that contains the Java™ code for the command.

An example of using the **HealthCheck** command is to verify that the Common agent password you are using is still valid:

```
For Windows:
   HealthCheck -registrationPw changeMe
For UNIX:
   ./HealthCheck.sh -registrationPw changeMe
```

Where **changeMe** is the password you are checking for.

LogCollector

LogCollector gathers logs and other information needed to debug problems with the Agent Manager. **LogCollector** creates the **LogCollector.zip** file, which is located in the root directory of where the Agent Manager is installed. This command must be run from the machine on which the Agent Manager is installed. Review the **LogCollector** README file for instructions on running the commands and the parameters required for this command. The **LogCollector** README file also lists known issues and limitations. The files are:

LogCollector.sh
The script that runs the **LogCollector** command on UNIX and Linux systems.

LogCollector.bat
The script that runs the **LogCollector** command on Windows systems.

LogCollector.jar
The JAR file that contains the Java code for the command.

Repository Copy tool

The Repository Copy tool enables you to export all the tables in the IBM Tivoli Storage Productivity Center repository (Data Manager data) for purposes of debugging problems.

You can send the output to the IBM Support Center to help debug problems.

Note: With the addition of partitioned tables in IBM Tivoli Storage Productivity Center 4.1, there is a limitation on what database schemas can be imported between one another. For example, if the source database does not have partitioned tables (for example, a IBM Tivoli Storage Productivity Center 3.3 database), you cannot import the database into a new installation of IBM Tivoli Storage Productivity Center 4.1 (because the database has partitioned tables).

You can successfully import a IBM Tivoli Storage Productivity Center 3.3 database into an upgraded IBM Tivoli Storage Productivity Center 4.1 database if the database migration tool has not been run (the database does not have partitioned tables).

Before you can copy data, you must meet the following conditions:

- Both the source and target databases use the old database schema (with non-partitioned tables)
- Both the source and target databases use the new database schema (with partitioned tables)

You can determine if the database migration tool has been run if you see the `migrateTable.log` file in the `<TPC_install_directory>\data\server\tools` directory.

Exporting repository data

Use the Repository Copy tool to export data from an existing repository into a comma delimited file.

To export repository data, follow these steps:

1. Stop the Device server and Data server.
2. Go to the following default directory:
   - If on Windows: `c:\Program Files\IBM\TPC\data\server\tools`
   - If on Linux or UNIX: `/opt/IBM/TPC/data/server/tools`
3. Run the **repocopy.bat** or **repocopy.sh** command. A window is displayed prompting you for what you would like to do.
4. Select **Export data from repository tables**, and then click **Next**. The Options for Import/Export window is displayed.

5. In the Options for Import/Export window, enter information in the following fields:

   Directory for Export
   Enter the directory where the comma delimited file will be saved.

   Delimiter
   Enter a delimiter for the delimited file format (comma is the default).

   Quote Enter the symbol that will contain string data (double quotes is the default).

   IBM Tivoli Storage Productivity Center will export the data into the comma delimited file you specify, and place it in a file named `<tablename>.txt`. Click **Next**.
6. The Choices for DB2 Export window is displayed. You have the following choices:
   - Export Using DB2 native format
   - Export Using text files

   Make a choice and click **Next**.
7. The Table Export Choices windows is displayed. You have the following choices:
   - Export base tables (always)
   - Export Performance Manager tables
   - Export history tables used for IBM Tivoli Storage Productivity Center for Data history reports

   Make a choice and click **Next**.
8. The Connection Properties window is displayed.

   In this step, Tivoli Storage Productivity Center reads the `server.config` file and determines your current repository and the information for accessing the database.

   The information detected in the `server.config` file is displayed in the Connection Properties window within the following fields: **Database Types**, **User name**, **Password**, **Driver Class**, **Driver URL**, **Database**, **DB Creator**, **Classpath**.

   Note: If you want to export data from a different database from the one listed in the `server.config` file, you can select the database from the **Database Types** list box and manually enter the database information. Click **Finish**.
9. Tivoli Storage Productivity Center will then connect to the database and display the database and connection information. Click **Run** to begin the export process.
10. A window is displayed containing the export progress log. As you progress through the export, messages are written to this progress log, allowing you to keep track of the steps as they happen.

    Note: When you run the **repocopy** tool on a server that has a remote database, the **repocopy** tool displays a message. When the **repocopy** tool is used with a remote database, the DB2 shared library is not available for loading `libTSRMinsudb.so`. You can ignore this message. Just click **OK** and continue.

Importing repository data

Use the Repository Copy tool to import data from a comma delimited file. You must have previously exported data from an existing repository.

To import data into repository tables, follow these steps:

1. Stop the Device server and Data server.
2. Go to the following directory:
   - If on Windows: `c:\Program Files\IBM\TPC\data\server\tools`
   - If on Linux or UNIX: `/opt/IBM/TPC/data/server/tools`
3. Run the **repocopy.bat** or **repocopy.sh** command. A window is displayed prompting you for what you would like to do
4. Select **Import data into repository tables**, and then click **Next**. The Options for Import/Export window is displayed.
5. In the Options for Import/Export window, enter information in the following fields:

   Directory for Import
   Enter the directory where the comma delimited files are stored.

   Delimiter
   Enter a delimiter used for the delimited file format (comma is the default). You must use the same delimiter that was used in the exported file.

   Quote Enter the symbol used to contain string data (double quotes is the default). You must use the same quote that was used in the exported file.

   Delete before inserting
   Check this option if you want to delete any existing data in the repository tables before importing new data.

   IBM Tivoli Storage Productivity Center will import the data from the comma delimited file that you specify. Click **Next**.
6. The Connection Properties window is displayed.
7. Enter the database and the access information of the database to which you want to import data. The Connection Properties window contains the following fields: **Database Types**, **User name**, **Password**, **Driver Class**, **Driver URL**, **Database**, **DB Creator**, **Classpath**. Click **Finish**.
8. Tivoli Storage Productivity Center will then connect to the database and display the database and connection information. Click **Run** to begin the import process.
9. A window is displayed containing the import progress log. As you progress through the import, messages are written to this progress log, allowing you to keep track of the steps as they happen.

   Note: When you run the **repocopy** tool on a server that has a remote database, the **repocopy** tool displays a message. When the **repocopy** tool is used with a remote database, the DB2 shared library is not available for loading `libTSRMinsudb.so`. You can ignore this message. Just click **OK** and continue.

Service tool: collecting information

This section provides information about the information the **Service** tool collects from all installed IBM Tivoli Storage Productivity Center components. This tool detects system configuration, collects the applicable information, and creates a ZIP file that can be sent to the IBM support center.

The **Service** tool collects the following information:

- Host name
- IP address and configuration information
- Operating system and version
- Java home, version, and classpath
- Java Virtual Machine (JVM) implementation name and version
- Protocol statistics and current TCP/IP network connections, including listening ports
- Diagnostic information regarding the system and its services
- Listing of all library files, for example, server and library, agent and library, and GUI library

When the **Service** tool is run on the system where the Data server or the Device server is installed, it also collects the following information:

- For the Data Server, information on all the GUIs (remote and local) associated with this Data Server
- All e-fixes applied
- Installation logs
- Everything in the **log** directory, including subdirectories
- Everything in the **conf** directory
- Directory listing of the **lib** directory
- Everything in the WebSphere log directory
- Information in the **server.xml** file:

  ```
  manager\apps\was\config\cells\DefaultNode\nodes\DefaultNode\
  servers\server1\server.xml
  ```
- The version of DB2 (output from the **db2level** command)
- Output from the **ipconfig /all** command (Windows)
- Output from the **ipconfig -a** command (Linux or UNIX)
- Output from the **netstat -an** command

When the **Service** tool runs on the Data agent or Fabric agent computer, it collects this information:

- All e-fixes applied
- `traceNative.log` files from `<Agent_install_dir>\servlet\bin` directory.
- Scanner benchmark files from `<Agent_install_dir>\agent\bin\<platform>` directory.
- Core files on the agent computer (if any).
- Output from scanners run manually from the command line.
- Everything in the `log` directory (including subdirectories).
- Everything in the `conf` directory
- Directory listing of the `lib` directory
- Output from the **ipconfig /all** command (for Windows).

- Output from the **ifconfig -a** command (for UNIX or Linux).
- Output from the **netstat -an** command

Note: the **Service** tool does not collect service logs or other service information for IBM Tivoli Storage Productivity Center for Replication. If you need to collect the log and trace files for the replication server, see the **Obtaining logs and tracing** topic in the **Troubleshooting** section of the Tivoli Storage Productivity Center for Replication Information Center.

Running the Service tool for the servers

This topic describes how to run the **Service** tool for the Data Server and the Device server.

To run the tool for the servers, complete the following steps:

1. Log onto the system. You must have administrator authority on Windows or root authority on UNIX or Linux.
2. If you used the default directory, go to the following directory:
   - For Windows: `C:\Program Files\IBM\TPC\service\`.
   - For Linux or UNIX: `/<usr or opt>/IBM/TPC/service/`.
3. Run the following program:
   - For Windows: `service.bat`.
   - For Linux or UNIX: `service.sh`.
4. One or more zip files are created in the directory where you ran the **Service** tool. If you have a Data agent or Fabric agent installed on the same computer as the Data Server or the Device server, that information is collected when you run the tool for the Data Server or the Device server. The following zip files are created:

   ```
   TPCDATAservice.zip (for the Data Server)
   TPCDEVservice.zip (for the Device Server)
   TPCDATACAservice.zip (for the Data agent)
   TPCDEVCAservice.zip (for the Fabric agent)
   ```

Running the Service tool for the agents

This topic describes how to run the **Service** tool for the remote agents.

To run the **Service** tool for the remote Data agent or Fabric agent, complete the following steps:

1. Log onto the system. You must have administrator authority on Windows or root authority on UNIX or Linux.
2. If you used the default directory, go to the following directory:
   - For Windows: `C:\Program Files\IBM\TPC\ca\subagents\TPC\service`.
   - For Linux or UNIX: `/<usr or opt>/IBM/TPC/ca/subagents/TPC/service`.
3. Run the following program:
   - For Windows: `service.bat`.
   - For Linux or UNIX: `service.sh`.
4. A zip file is created in the directory where you ran the Service tool. The following zip files are created:

   ```
   TPCDATACAservice.zip (for the Data agent)
   TPCDEVCAservice.zip (for the Fabric agent)
   ```

Tracing the servers and agents

The tracing tool uses the IBM Logging Toolkit for C (CCLOG) for tracing.

Tracing the Data server

You can trace the following service components:

- Server
- Scheduler
- GUI
- Device
- Agent (local agent)

If you want to trace the Data server, you must turn all these components on.

To start the tracing function for the Data server services, follow these steps:

1. Open the Tivoli Storage Productivity Center GUI.
2. Click **Administrative Services** → **Services** → **Data Server**. Right-click **Server**. Click **Configure Tracing**.

   Note: Right-click on any of the other components to trace a component: Scheduler, GUI, Device, or Agent. For the Scheduler, GUI, Device, or Agent, you can only enable tracing and set the debugging level.
3. The tracing dialog window is displayed. Check the box for **Enable Trace** and enter a value for the following options:

   Level The level can be:

   DEBUG_MIN
   The minimum debugging level. This is the default and is always turned on. This includes trace information for ″key″ locations in the code which include first failure data capture. Currently, this setting does not capture any output data.

   DEBUG_MID
   The medium debugging level. This is reserved for future use and will include trace information for customized trace statements to aid in debugging. Currently, this setting does not capture any output data.

   DEBUG_MAX
   The maximum debugging level. This includes entry and exit information for every method in addition to all DEBUG_MID and DEBUG_MIN statements. This can be a performance hit.

   Maximum Number of Files
   This is the maximum number of files used before the Data server starts reusing the tracing files. The default is 5 files.

   Maximum File Size (kilobytes)
   This is the maximum file size for the tracing files. The default is 20 MB.

   Click **OK**.

Note: Changes made in the Agent or Server Tracing Configuration panel are not saved in the nativelog.config file. Settings made in the Agent or Server Tracing Configuration panel are effective only while the process is running. For example, if

the agent's default trace level was modified through the Agent Tracing Configuration panel, those settings will only be effective while the agent is running. If that agent is restarted, the defaults specified in the nativelog.config file will be reread and used.

Traces are written by default to <TPC_install_dir>\data\log\ServerTrace.log.

Tracing a remote Data agent or Storage Resource agent

To trace a remote Data agent or Storage Resource agent, follow these steps:

1. Open the Tivoli Storage Productivity Center GUI on the server.
2. Click **Administrative Services → Data Sources → Data/Storage Resource Agents**.
3. In the right pane, select a Data agent and click **Configure Tracing**.
4. The tracing dialog window is displayed. Check the box for **Enable Trace** and enter a value for the following options:

   Level The level can be:

   DEBUG_MIN
   The minimum debugging level. This is the default and is always turned on. This includes trace information for ″key″ locations in the code which include first failure data capture. Currently, this setting does not capture any output data.

   DEBUG_MID
   The medium debugging level. This is reserved for future use and will include trace information for customized trace statements to aid in debugging. Currently, this setting does not capture any output data.

   DEBUG_MAX
   The maximum debugging level. This includes entry and exit information for every method in addition to all DEBUG_MID and DEBUG_MIN statements. This can be a performance hit.

   Maximum Number of Files
   This is the maximum number of files used before the Data server starts reusing the tracing files. The default is 5 files.

   Maximum File Size (kilobytes)
   This is the maximum file size for the tracing files. The default is 20 MB.

   Click **OK**.

Tracing the Device server

There are several different trace loggers that allow tracing to be set for specific areas of the product. For example, one of these areas is the san.SanEventCorrelatorFactory trace logger that sets the tracing levels for functions that deal with SNMP traps. This section provides information on how to change the tracing levels for these areas. When collecting information for the investigation of problems, IBM support would provide information on the logger to configure. The following example shows how to change the tracing level for the logger.

The logging can be set to the following:

INFO This contains large amounts of data logged in the trace files and is used generally when investigating specific issues with the product.

WARN
This contains warning and error level messages logged in the trace files. This is the default level of tracing.

ERROR
This contains only error messages logged in the trace files.

ALL This contains all warning and error messages.

Tracing is already turned on by default. These instructions allow you to specify a different level of tracing. This example shows you how to turn tracing on for the Event Correlator.

To turn tracing on for the Device server, follow these steps:

1. Run the **srmcp** tracing command.

   Windows

   To run the **srmcp** tracing command on Windows, follow these steps:

   a. Go to this default directory on Windows:

      ```
      C:\Program Files\IBM\TPC\device\bin\w32-ix86
      ```

   b. Run the following command: **setenv**.

   c. Run this command:

      ```
      srmcp -u <user_id> -p <password> log set
      san.SanEventCorrelatorFactoryTraceLogger -filterkey INFO
      ```

      The filter key can be: INFO, WARN, ERROR, or ALL.

   UNIX

   To run the **srmcp** tracing command on UNIX, follow these steps:

   a. Go to this default directory on UNIX:

      ```
      /<usr or opt>/IBM/TPC/device/bin/aix (for AIX)
      /<usr or opt>/IBM/TPC/device/bin/linux (for Linux)
      ```

   b. Run this command:

      ```
      ./setenv.sh
      ```

   c. Run this command:

      ```
      ./srmcp.sh -u <user_id> -p <password>
      log set san.SanEventCorrelatorFactoryTraceLogger -filerkey INFO
      ```

      The filter key can be: INFO, WARN, ERROR, or ALL.

   The output will be logged in traceTPCDeviceServer.log.

Tracing the Fabric agent

To turn tracing on for the Fabric agent, follow these steps:

1. Run the **agentcli** command.

   Windows

   To run the **agentcli** command on Windows, follow these steps:

   a. Go to this default directory on Windows:

      ```
      C:\Program Files\IBM\TPC\ca
      ```

   b. Run this command:

```
agentcli TPCFabric log get san.InbandScannerTraceLogger
-filterkey INFO
```

The filter key can be: INFO, WARN, ERROR, or ALL.

UNIX To run the **agentcli** command on UNIX, follow these steps:

a. Go to this default directory on UNIX:

```
<usr or opt>/IBM/TPC/ca
```

b. Run this command:

```
agentcli TPCFabric log get san.InbandScannerTraceLogger
-filterkey INFO
```

The filter key can be: INFO, WARN, ERROR, or ALL.

The output will be logged in C:\Program Files\IBM\TPC\ca\subagents\TPC\Fabric\log for Windows and /<usr or opt>/IBM/TPC/ca/subagents/TPC/Fabric/log for UNIX.

Debugging SNMP traps

The **srmcp** tracing command can provide detailed information about the SNMP traps received (when at the INFO level) so that the filters can be defined. It also allows you to see the traps that have been caught by the filters after they have been added. The traps are logged in the trace output with a message below them indicating which filter caught the trap.

To turn tracing on for the Device server, follow these steps:

1. Run the **srmcp** tracing command.

Windows

To run the **srmcp** tracing command on Windows, follow these steps:

a. Run the following command: **setenv**.

b. Go to this default directory on Windows:

```
C:\Program Files\IBM\TPC\device\bin\w32-ix86
```

c. Run this command:

```
srmcp -u <user_id> -p <password>
log set san.SanEventCorrelatorFactoryTraceLogger -filterkey INFO
```

The filter key can be: INFO, WARN, ERROR, or ALL.

UNIX

To run the **srmcp** tracing command on UNIX, follow these steps:

a. Go to this default directory on UNIX:

```
/<usr or opt>/IBM/TPC/device/bin/aix (for AIX)
/<usr or opt>/IBM/TPC/device/bin/linux (for Linux)
```

b. Run this command:

```
./setenv.sh
```

c. Run this command:

```
./srmcp.sh -u <user_id> -p <password>
log set san.SanEventCorrelatorFactoryTraceLogger -filerkey INFO
```

The filter key can be: INFO, WARN, ERROR, or ALL.

The output will be logged in traceTPCDeviceServer.log.

Trace output location

The tracing output is stored in the locations shown in Table 1 and Table 2.

Table 1. Java trace output files

Class	Java trace output file
Agent	<TPC_install_dir>/data/ca/subagents/TPC/data/log/<hostname>/AgentTrace.log
Server	<TPC_install_dir>/data/log/ServerTrace.log

Table 2. Native trace output files

Class	Native trace output file
Agent	<TPC_install_dir>/data/ca/subagents/TPC/data/log/<hostname>/traceTSRMNativeAgent_<yyyymmddhhmmssLOGx>.log
Server	<TPC_install_dir>/data/log/traceTPCDNativeServer_<yyyymmddhhmmssLOGx>.log
UnixStopServer (only on UNIX)	<TPC_install_dir>/data/log/traceTPCDNativeStopServer_<yyyymmddhhmmssLOGx>.log
ImportExport	<TPC_install_dir>/data/log/traceTPCDNativeImportExport_<yyyymmddhhmmssLOGx>.log

Note:

- There are two copies of nativelog.config: one for the Agent and one for the Server. Depending on which component you want to trace, make sure the correct nativelog.config file is modified. For example, if you would like to trace the Agent, then make sure the nativelog.config file in `<TPC_install_dir>/data/ca/subagents/TPC/data/config` is configured for Agent tracing. Alternatively, if you wanted to perform ImportExport tracing, make sure that <TPC_install_dir>/data/config/nativelog.config is configured properly.
- Messages are logged before the creation of a user-defined installation directory. Because the log directory is not initially available, installation messages are written to standard error output. The standard error output is redirected to a temporary file. The temporary message log file is copied to the appropriate log directory, and the standard error output is redirected to reflect the new message log file location.

When an executable is run, it creates the tracing files shown in Table 3.

Table 3. Trace files for executables

Executable	Platform	Trace output location, file name, and description

udbexec.exe	Windows	`\<hostname>\traceTSRMNativeUdbexec_<PID>.log`
		`<TPC_install_dir>\data\ca\traceTSRMNativeUdbexec_<PID>_<yyyymmddhhmmss>.log`
		
		This exe is run whenever a DB2 probe or scan occurs. A DB2 probe occurs upon registration of a DB2 instance.

ExecSvc.exe	Windows	`<temp>\traceTSRMNativeExecSvc.log` This exe runs only on the remote computer during a Windows push installation.

Table 3. Trace files for executables (continued)

Executable	Platform	Trace output location, file name, and description
udbexec.exe	Windows and UNIX or Linux	<TPC_install_dir>/data/ca/traceTSRMNativeudbexec_<pid>_<yyyymmddhhmmss>.log This exe is run whenever a DB2 probe or scan occurs. A DB2 probe occurs upon registration of a DB2 instance.

Configuration files for tracing

Two configuration files have been added for Java tracing: ServerTraceLog.config and AgentTraceLog.config. The nativelog.config file will be used for native code tracing.

Table 4. Trace configuration files

File	Directory
Native code tracing configuration file (for server)	<TPC_install_dir>/data/config/nativelog.config
Native code tracing configuration file (for agent)	<TPC_install_dir>/data/ca/subagents/TPC/Data/Config/nativelog.config
Java tracing configuration file (for server)	<TPC_install_dir>/data/config/ServerTraceLog.config
Java tracing configuration file (for agent)	<TPC_install_dir>/data/ca/subagents/TPC/data/Config/AgentTraceLog.config

When changes are made to tracing configuration through the GUI, the appropriate configuration file will be updated. For example, if you change the Maximum Number of Files parameter for server tracing, the new value will appear in both ServerTraceLog.config and nativelog.config. Both ServerTraceLog.config and AgentTraceLog.config have additional parameters that are not listed below. Only the parameters listed below can be modified.

nativelog.config file

During installation or a maintenance upgrade of Tivoli Storage Productivity Center, a file named nativelog.config will be created in the <TPC_install_dir>/data/config directory. This file is used to specify the default levels of tracing for the Tivoli Storage Productivity Center processes (for example, the server and the agent). Depending on which platform Tivoli Storage Productivity Center is installed, some of the items listed in this file might vary. In particular, Windows installations will not have the #StopAgent section that UNIX platforms have. The reason for the difference is because the agent service is invoked differently on these platforms. Settings in these files are separated into sections representing Tivoli Storage Productivity Center processes. The trace settings per process are listed underneath the process they represent. For more information on these settings and their default values, see “Tracing the Data server ” on page 11.

The definitions for the entries in the nativelog.config file are:

 # Agent
: An Tivoli Storage Productivity Center process. In this case, the agent process is specified.

agt.level=DEBUG_MIN
The default trace level for the agent process.

agt.maxFiles=3
The maximum number of trace files to be created or used.

agt.maxFileSize=20480000
The maximum size of each trace file (in bytes).

The level of trace specified in this file is also applied to specific child processes executed from their parent process. For example, processes udbexec.exe and MiniProbe.exe can use the settings from the Agent section as their default trace level, because these processes are executed by the Agent process.

It is possible for you to edit the contents of this file if you want to change the default level of tracing for a particular process. However, these changes will not take effect until that process is stopped and restarted again. For example, if you want to change the tracing level of the server from DEBUG_MIN to DEBUG_MID, the nativelog.config file could be changed to reflect this. Once this change is made, the server would have to be stopped and restarted to make the change effective. By modifying this file, you can ensure that each time the server is stopped and restarted, the desired trace level will be applied.

Alternatively, trace settings can be changed during runtime through the GUI. For more configuration information, see "Tracing the Data server " on page 11.

Here is an example of the nativelog.config file for Windows:

```
# level: DEBUG_MIN || DEBUG_MID || DEBUG_MAX
# maxFiles: >= 2 && <= 100
# maxFileSize: >= 128000 && <= 102400000

# Agent
agt.level=DEBUG_MIN
agt.maxFiles=3
agt.maxFileSize=20480000

# Server
srv.level=DEBUG_MIN
srv.maxFiles=3
srv.maxFileSize=20480000

# ImportExport
impexp.level=DEBUG_MIN
impexp.maxFiles=3
impexp.maxFileSize=20480000
```

Here is an example of the nativelog.config file for UNIX:

```
# level: DEBUG_MIN || DEBUG_MID || DEBUG_MAX
# maxFiles: >= 2 && <= 100
# maxFileSize: >= 128000 && <= 102400000

# Agent
agt.level=DEBUG_MIN
agt.maxFiles=3
agt.maxFileSize=20480000

# ImportExport
impexp.level=DEBUG_MIN
impexp.maxFiles=3
impexp.maxFileSize=20480000
```

```
# StopAgent
stpagt.level=DEBUG_MIN
stpagt.maxFiles=3
stpagt.maxFileSize=20480000
```

AgentTraceLog.config file

The definitions for the entries in the AgentTraceLog.config file are:

ITSRM.logger.trace.Agent.logging=true
The state of tracing for the agent. In this case, tracing is enabled.

ITSRM.logger.trace.Agent.level=DEBUG_MAX
The trace level for the agent. DEBUG_MAX is currently the only level that will produce trace output.

ITSRM.handler.file.maxFiles=8
The maximum number of trace files to be created or used.

ITSRM.handler.file.maxFileSize=1024
The maximum size of each trace file (in bytes).

ServerTraceLog.config file

The definitions for the entries in the ServerTraceLog.config file are:

ITSRM.logger.trace.Server.logging=true
The state of tracing for the Server process. In this case tracing is enabled.

ITSRM.logger.trace.Server.level=DEBUG_MAX
The trace level for the Server process. DEBUG_MAX is currently the only level that will produce trace output.

ITSRM.handler.file.maxFiles=5
The maximum number of trace files to be created or used.

ITSRM.handler.file.maxFileSize=20480
The maximum size of each trace file (in bytes).

ITSRM.logger.trace.TivoliSRM-GUI.logging=false
The state of tracing for the GUI process. In this case tracing is disabled.

ITSRM.logger.trace.TivoliSRM-GUI.level=DEBUG_MAX
The trace level for the GUI process.

ITSRM.logger.trace.TivoliSRM-CIMOM.logging=false
The state of tracing for the CIMOM process. In this case tracing is disabled.

ITSRM.logger.trace.TivoliSRM-CIMOM.level=DEBUG_MAX
The trace level for the CIMOM process.

ITSRM.logger.trace.TivoliSRM-Agent.logging=false
The state of tracing for the Agent process. In this case tracing is disabled.

ITSRM.logger.trace.TivoliSRM-Agent.level=DEBUG_MAX
The trace level for the Agent process.

ITSRM.logger.trace.Scheduler.logging=false
The state of tracing for the Scheduler process. In this case tracing is disabled.

ITSRM.logger.trace.Scheduler.level=DEBUG_MAX
The trace level for the Scheduler process.

Chapter 3. Troubleshooting Agent Manager problems

This section provides information to help you troubleshoot Agent Manager problems.

Agent Manager log files

Agent Manager log files give important information about an installation, uninstallation, and so forth.

Agent Manager log files can be found in the following locations:

Table 5. Agent Manager Log file locations

Agent Manager log files	Location
Installation	`<Agent_Manager_install_dir>\logs`
Uninstallation	`<Agent_Manager_install_dir>\logs`
Run-time	`<Agent_Manager_install_dir>\AppServer\agentmanager\logs\AgentManager`

Agent Manager installation log files

Log files are created when you install Agent Manager.

The log files generated during the installation and initial configuration of the Agent Manager are located in the `<Agent_Manager_install_dir>\logs` directory. The following table lists the logs.

Table 6. Agent Manager installation log files

Log File	Description
am_install.log	InstallShield MultiPlatform (ISMP) log for the installation of the Agent Manager. Check this log first to verify that the Agent Manager installed properly and the Agent Manager server is started.
AMReturnValues.log	Summary of return values of the steps of the Agent Manager installation.
am_upgrade.log	Information about whether a new installation was performed or an existing version of the Agent Manager was found on the system and upgraded.
auth_stdout.log auth_stderr.log	Standard output and standard error logs for the AuthXMLUpdate program.
certGen_stdout.log certGen_stderr.log	Standard output and standard error logs for the generation of the root certificate for the Agent Manager certificate authority.
datastore.out	Log of the Data Definition Language (DDL) script that creates and initializes the registry database.
datastore_stdout.log datastore_stderr.log	Standard output and standard error logs for creating and initializing the tables in the registry database.
ds_install.log	An ISMP log for installing the files necessary to create the registry database.
db_stdout.log db_stderr.log	Standard output and standard error logs for the creation of the registry database.
encrypt_stdout.logencrypt_ stderr.log	Standard output and standard error logs for the EncryptAMProps program.

Table 6. Agent Manager installation log files (continued)

Log File	Description
guid_install.logguid_stdout. logguid_stderr.log	Standard output and standard error logs for installing the Tivoli® globally unique globally unique identifier (GUID).
msg_EPM_Install.log trace_EPM_Install.log	Messages and trace information generated during the installation and configuration of the Agent Manager applications in WebSphere Application Server.
serverStatus_out.log serverStatus_err.log	Standard output and standard error logs for starting the application server for the Agent Manager.
startserver_stdout.log startserver_stderr.log	Standard output and standard error logs for starting the Agent Manager server under WebSphere.
jacl/amApp_out.log jacl/amApp_err.log	Standard output and standard error logs generated while installing the AgentManager and AgentRecoveryService applications and `.WAR` files. These logs are generated by the `EPMInstallApp.jacl` configuration script.
jacl/appServer_out.log jacl/appServer_err.log	Standard output and standard error logs generated while installing the application server for the Agent Manager. These logs are generated by the `EPMAppServer.jacl` script.
jacl/checkcell_out.log jacl/checkcell_err.log	Standard output and standard error logs for verifying the cell for the WebSphere configuration. These logs are generated by the `EPMValidate.jacl` script.
jacl/checknode_out.log jacl/checknode_err.log	Standard output and standard error logs for verifying the node for the WebSphere configuration. These logs are generated by the `EPMValidate.jacl` script.
jacl/jdbc_out.log jacl/jdbc_err.log	Standard output and standard error logs for the configuration of the WebSphere JDBC provider, data source, and J2C Authentication Data Entry. These logs are generated by the `EPMJdbcProvider.jacl` script.
jacl/ssl_out.log jacl/ssl_err.log	Standard output and standard error logs for the SSL configuration. These logs are generated by the `EPMSSLConfig.jacl` script.
jacl/virHost_out.log jacl/virHost_err.log	Standard output and standard error logs for creating the WebSphere virtual host. These logs are generated by the `EPMVirtualHost.jacl` script.

Agent Manager uninstallation log files

Log files are created when you uninstall Agent Manager.

The log files generated when you uninstall the Agent Manager are located in the `<Agent_Manager_install_dir>\logs` directory. The following table lists the logs that are created.

Table 7. Agent Manager uninstallation log files

Log File	Description
uninstall.log	InstallShield MultiPlatform (ISMP) log for uninstalling the Agent Manager.
AMUninstallReturn Values.log	Summary of return values of the steps of the Agent Manager uninstallation.
msg_EPM_Install.logtrace_ EPM_Install.log	Messages and trace information that is generated when uninstalling the Agent Manager applications in WebSphere Application Server.

Agent Manager run-time log files

The run-time logs for the Agent Manager are located in the `<Agent_Manager_install_dir>\embedded\logs\<app_server_name>` directory, where `<app_server_name>` is the name of the application server. By default, this is **Agent Manager**.

Run-time Agent Manager-specific WebSphere log files

The run-time logs for the Agent Manager-specific WebSphere logs are located in the <Agent_Manager_install_dir>\embedded\logs\AgentManager directory.

Packaging the Agent Manager log files

You can use the **LogCollector** tool to collect logs and other information needed to debug problems with the Agent Manager. If you contact IBM Customer Support, you will need to provide this package.

If you encounter a problem that you cannot resolve immediately, you can use the **LogCollector** tool to package the Agent Manager log files. This preserves the information that you need to perform detailed problem determination and prevents you from having to scan back through messages and trace events that were recorded after the problem occurred. For more information about the **LogCollector** tool, see the readme file in the <Agent_Manager_install_directory>/toolkit directory, where <Agent_Manager_install_directory> is the directory where Agent Manager is installed. The readme file provides information on how to run this tool and what parameters you can specify.

1. Change to the following directory:

   <Agent_Manager_install_directory>/toolkit/bin

   where <Agent_Manager_install_directory> is the directory where Agent Manager is installed.
2. Run one of the following commands:

On Linux or UNIX	LogCollector.sh
On Windows	LogCollector.bat

   The **LogCollector.zip** file is created.

Determining the version of Agent Manager

You can run the **GetAMInfo** command to determine which version of Agent Manager is installed.

To display the version of the Agent Manager, complete the following steps:

1. Change to one of the following directories:

For AIX, Linux, and UNIX	<Agent_Manager_install_directory>/bin
For Windows	<Agent_Manager_install_directory>\bin

   where <Agent_Manager_install_directory> is the directory where Agent Manager is installed.
2. To display the version for the Agent Manager, run the following command:

For AIX, Linux, and UNIX	GetAMInfo.sh AgentManager
For Windows	GetAMInfo.bat AgentManager

3. To get the version for the Agent Recovery Service, run the following command:

For AIX, Linux, and UNIX	GetAMInfo.sh AgentRecoveryService
For Windows	GetAMInfo.bat AgentRecoveryService

Registration failures

This section lists the common reasons for an Agent Manager registration failure.

Common registration failures are:

- A short host name was used instead of a fully-qualified host name when the Agent Manager was installed. For example, myserver was used instead of myserver.ibm.com. The agent cannot resolve the registration URL (which contains the short name) with the Agent Manager server.
- The wrong truststore file was used when installing the agent or resource manager.
- The wrong agent registration password was used when installing the agent.
- The wrong resource manager registration user ID or password was used when installing the resource manager.
- The Agent Manager configuration files, such as AgentManager.properties or Authorization.xml, are missing or have been altered.
- The security credentials are invalid, expired, or revoked.

Chapter 4. Troubleshooting IBM Tivoli Storage Productivity Center problems

This section provides information to help you troubleshoot IBM Tivoli Storage Productivity Center problems.

Configuration files

This topic provides default file locations for IBM Tivoli Storage Productivity Center configuration files.

The default file locations for the configuration files are shown in the following table.

Table 8. Default file locations for Tivoli Storage Productivity Center configuration files

Component	Default file location

Tivoli Storage Productivity Center	For Windows:
	`c:\Program Files\IBM\TPC\config`
	For UNIX or Linux:
	`/opt/IBM/TPC/config`

Data Server	For Windows:
	`c:\Program Files\IBM\TPC\data\config`
	For UNIX or Linux:
	`/opt/IBM/TPC/data/config`

Device server	For Windows:
	`c:\Program Files\IBM\TPC\device\conf`
	For UNIX or Linux:
	`/opt/IBM/TPC/device/conf`

Common agent	For Windows:
	`c:\Program Files\IBM\TPC\ca\config`
	For UNIX or Linux:
	`/opt/IBM/TPC/ca/config`

Data agent	For Windows:
	`c:\Program Files\IBM\TPC\ca\subagents\TPC\Data\config`
	For UNIX or Linux:
	`/opt/IBM/TPC/ca/subagents/TPC/Data/config`

Fabric agent	For Windows:
	`c:\Program Files\IBM\TPC\ca\subagents\TPC\Fabric\conf`
	For UNIX or Linux:
	`/opt/IBM/TPC/ca/subagents/TPC/Fabric/conf`

server.config file

This topic lists the parameters that are set in the `server.config` file. These parameters include Controller, Logging, Repository, and Service.

Table 9. Parameters for server.config file

Parameter	Description
Controller parameters:	
name	Data Manager Server Name (host computer name)
port	Port on which the server listens for requests (9549)
maxConnections	Maximum number of concurrent sockets that the server will open (500)
routerThreads	Number of threads redirect incoming requests to the appropriate service provider (1)
serviceThreads	Number of threads to allocate for the Server internal service provider (2)
agentErrorLimit	Number of consecutive attempts to reach an agent before the agent is displayed as DOWN. When an agent is in this state, no attempts to connect are made until either the agent contacts the Server or the agent status is manually changed to UP (3)
adminGroup	Name of the group a user must be a member of in order to perform administrative functions from the Graphic User Interface (adm)
commEncrypted	Switch that secures communication between the Server/Agent and the Server/GUI by encrypting the data stream. • `0` = Off. Do not encrypt the data stream. • `1` = On. Encrypt the data steam.
hostAlias	This parameter appears if the HOST_ALIAS is not specific and represents the name of the server. This value for this parameter is used when multiple computers have the same name or the name cannot be determined.
Logging parameters:	
logsKept	Number of server logs to keep (5)
messagesPerLog	Maximum number of messages in a log, when this number is reached the log is closed and a new log is created. (100,000)
Repository parameters:	
driver	Name of the JDBC driver to use, normally: • Oracle: `oracle.jdbc.driver.OracleDriver` • MS SQL: `com.inet.tds.TdsDriver` • Sybase: `com.sybase.jdbc2.jdbc.SybDriver` • UDB/DB2: `COM.ibm.db2.jdbc.app.DB2Driver` • Cloudscape®: `com.ibm.db2j.jdbc.DB2jDriver`

Table 9. Parameters for server.config file (continued)

Parameter	Description
url	The URL used to connect to the database, normally: • Oracle: `jdbc:oracle:thin:@<host_name>:<port>:<SID>` • MS SQL: `jdbc:inetdae:<host_name>` • Sybase: `jdbc:sybase:Tds:host_name:port` • UDB: `jdbc:db2:database_name`
user	User name that IBM Tivoli Storage Productivity Center uses to connect to the repository
connectionPool	Number of database connections in a pool of reusable open connections (10)
Service parameters	
name	Repeating section that indicates the service providers to start. REQUIRED: • `TStorm.server.svp.GuiSvp` • `TStorm.server.svp.AgentSvp` • `scheduler.Scheduler`

scheduler.config file

This topic lists the parameters that are set in the scheduler.config file. These parameters include Concurrency parameters and Jobs parameters.

Table 10. Parameters for scheduler.config file

Parameter	Description
Concurrency parameters:	
maxSubmitThreads	Number of threads to create that handle the submission of jobs (3)
maxCompleteThreads	Maximum Number of threads to create to handle job completions. Initially will create a pool of 1/2 the number that can grow to the maximum (3)
Jobs parameters:	
minutesAdvanced	Number of minutes in advance of scheduled time to begin the scheduling process. This allows for the overhead time involved in scheduling a job so that the job will actually start close to the scheduled time (1)

TPCD.config file

This topic lists parameters that are set in the `TPCD.config` file. These include Server parameters and GUI parameters.

Parameter	Description
Server parameters:	

Parameter	Description
threadPoolSize	Number of initial threads to create for handling requests (3)
abbreviatedProbe	Only Small Computer Systems Interface (SCSI) commands are sent to disk drives for inquiry and disk capacity information (1).
maxThreads	Maximum number of threads allowed for handling requests (8)
pingReceiveTimeout	Number of seconds to wait before indicating that a ping has failed (10)
skipAutoFS=1	Set to 1 for discovery on the Solaris Data agent to skip the automounts process. By default, discovery always processes automounts on the Solaris Data agent.
GUI parameters:	
threadPoolSize	Number of initial threads to create for handling user interface requests (3)
maxThreads	Maximum number of threads allowed for handling user interface requests (10)
reportRowLimit	Maximum number of rows that will be sent at a time to the user interface. If this number is exceeded, a **More** button will be displayed above the table, along with a warning message (5000)
keepCachedReport	Number of minutes to retain incomplete reports in the server's *tmp* directory (120)

agent.config file

The **agent.config** file contains configuration parameters for the Data agent. These parameters are set when the Data agent is installed; they can also be changed manually by modifying the file.

The following table contains the parameters for the **agent.config** file. If the Data agent is installed in the default location, this file is located at either `/opt/IBM/TPC/ca/subagents/TPC/Data/config` or `C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\config`.

Table 11. Parameters for the **agent.config** *file*

Parameter	Description
agentPort	Port on which the Data agent listens for requests. By default, this is set to 9510.
serverHost	Fully qualified host name of the system on which the Data server is installed.
serverPort	Port on which the Data server listens for requests. By default, this is set to 9549.
logFilesKept	Maximum number of Data agent logs that are retained. When this number is reached, the oldest log file is overwritten. By default, this is set to five.
messagesPerLog	Maximum number of messages in a Data agent log file. When this number is reached, the a new log file is created. By default, this is set to 100,000.

Table 11. Parameters for the **agent.config** *file (continued)*

Parameter	Description
maxBacklog	Maximum number of uncompleted jobs that are permitted. When this number is reached, if additional job requests are made, any error is generated. By default, this is set to 500.
sendFailWait	Number of seconds to wait before the Data agent attempts to resend a message to the Data server. By default, this is set to 30.
maxIdleThreads	Maximum number of idle threads to retain for use by future jobs. By default, this is set to 10.
uptimePoll	How often (in seconds) should agent check to ensure it is up (20).
hostAlias	This parameter appears if the `HOST_ALIAS` is not specific and represents the name of the server. This value for this parameter is used when multiple computers have the same name or the name cannot be determined.
honorSentScripts	If this parameter is set to '1', 't', 'T', 'y', or 'Y', the Data agent can run scripts sent from the Data server. Otherwise, only scripts that are exist in the **scripts** directory on the system where the Data agent is installed can be run.
TPCInstallLocation	Directory where the Data agent is installed.

Log files

There are several product logs files to check when you have a problem.

Installation log files

This section provides information about the installation log files for IBM Tivoli Storage Productivity Center IBM Tivoli Storage Productivity Center for Replication, and IBM Tivoli Integrated Portal.

Installation log files for IBM Tivoli Storage Productivity Center for Replication

If you have an installation problem for IBM Tivoli Storage Productivity Center for Replication, collect the log files as indicated in Table 12. These log files will help you and IBM Customer Support identify where the problems are.

For other troubleshooting information for IBM Tivoli Storage Productivity Center for Replication, see the Information Center. Click **Tivoli Storage Productivity Center for Replication > Troubleshooting**.

Table 12. Installation log files for IBM Tivoli Storage Productivity Center for Replication

Log file	Directory location	Description

TPCRMInstall.log	UNIX:
	`/opt/IBM/replication`
	Windows:
	`C:\Program Files\IBM\replication`
		Main IBM Tivoli Storage Productivity Center for Replication log file for installation

Table 12. Installation log files for IBM Tivoli Storage Productivity Center for Replication (continued)

Log file	Directory location	Description
TPCRInstallIS.log	UNIX: /opt/IBM/TPC/log/tpcr/install Windows: C:Program Files\IBM\TPC\log\tpcr\install	Main IBM Tivoli Storage Productivity Center log file for IBM Tivoli Storage Productivity Center for Replication installation
WebSphere server logs: • startServer.log • SystemErr.log • SysteOut.log • stopServer.log	UNIX: /opt/IBM/replication/eWAS/profiles/CSM/logs/server1 Windows: C:\Program Files\IBM\replication\eWAS\profiles\CSM \logs\server1	You will see these logs if there are WebSphere start errors. Copy these logs at the first sign of an installation failure before the installation finishes.

Installation log files for IBM Tivoli Integrated Portal

Table 13. Installation log files for IBM Tivoli Integrated Portal

Log file	Directory location	Description
IA-TIPInstall-00.log	AIX: / Linux: /root Windows: C:\Documents and Settings\Administrator	Main IBM Tivoli Integrated Portal log file for installation

Installation log files for IBM Tivoli Storage Productivity Center

Table 14. Installation log files for IBM Tivoli Storage Productivity Center

Log file	Directory location	Description
<component> Install.log or <component> InstallIS.log	UNIX: /opt/IBM/TPC/log/<component>/install Windows: C:\Program Files\IBM\TPC\log\<component>\install	Main IBM Tivoli Storage Productivity Center log files for installation. <component> can be: • ca • cli • data • dbSchema • device • gui • subagents/TPC/<Data or Fabric>
TIPInstallIS.log	UNIX: /opt/IBM/TPC/log/tip/install Windows: C:\Program Files\IBM\TPC\log\tip\install	Installation log for IBM Tivoli Integrated Portal

Table 14. Installation log files for IBM Tivoli Storage Productivity Center (continued)

Log file	Directory location	Description

TPCRInstallIS.log	UNIX:
	`/opt/IBM/TPC/log/tpcr/install`
	Windows:
	`C:\Program Files\IBM\TPC\log\tpcr\install`
		Installation log for IBM Tivoli Storage Productivity Center for Replication

Default log file locations

Check the following default log file locations when you have a problem.

Table 15. Default log file locations for IBM Tivoli Storage Productivity Center components

Component	Log file location

Data Server	For Windows:
	`c:\Program Files\IBM\TPC\data\log`
	For UNIX or Linux:
	`/opt/IBM/TPC/data/log`

Device server	For Windows:
	`c:\Program Files\IBM\TPC\device\log`
	For UNIX or Linux:
	`/opt/IBM/TPC/device/log`

Common agent	For Windows:
	`c:\Program Files\IBM\TPC\ca\logs`
	For UNIX or Linux:
	`/opt/IBM/TPC/ca/logs`
	For agent on Virtual I/O Server (see note):
	`/home/padmin/agentInstall.log`

Data agent	For Windows:
	`c:\Program Files\IBM\TPC\ca\subagents\TPC\Data\log`
	For UNIX or Linux:
	`/opt/IBM/TPC/ca/subagents/TPC/Data/log`
	For agent on Virtual I/O Server (see note):
	`/home/padmin/agentInstall.log`

Fabric agent	For Windows:
	`c:\Program Files\IBM\TPC\ca\subagents\TPC\Fabric\log`
	For UNIX or Linux:
	`/opt/IBM/TPC/ca/subagents/TPC/Fabric/log`
	For agent on Virtual I/O Server (see note):
	`/home/padmin/agentInstall.log`

GUI	For Windows:
	`c:\Program Files\IBM\TPC\gui\log`
	For UNIX or Linux:
	`/opt/IBM/TPC/gui/log`

Table 15. Default log file locations for IBM Tivoli Storage Productivity Center components (continued)

Component	Log file location

Database schema	For Windows:
	`c:\Program Files\IBM\TPC\dbschema\log`
	For UNIX or Linux:
	`/opt/IBM/TPC/dbschema/log`

Audit logs

IBM Tivoli Storage Productivity Center has two audit logs. The audit logs provide an audit log for all administrator activities including identification of users initiating actions.

The audit logs are:

- Data server audit log, which is initiated by the GUI.
- Device server audit log, which is initiated by the API or CLI.

Audit logging is performed at the ″point of entry″:

- GUI commands are logged in the Data server audit log.
- API and CLI commands are logged in the Device server audit log.

Communication between the services are not logged:

- Data server to Device server
- Data server to agents
- Device server to agents

The Data server audit log is in the following directory:

`<TPC_install_directory>\IBM\TPC\data\log\AuditTrace.log`

The Device server audit log is in the following directory:

`<TPC_install_directory>\IBM\TPC\device\log\auditTPCDeviceServer.log`

The Device server audit log includes the following information:

- Timestamp
- User ID
- IP address of client
- Service and operation performed
- Key input and output parameters

Diagnosing IBM Tivoli Storage Productivity Center problems

This section provides information about the log files, audit files, configuration files, and what steps to take to diagnose a IBM Tivoli Storage Productivity Center problem.

When a problem occurs with IBM Tivoli Storage Productivity Center, you should collect information to help diagnose the problem. The following sections help you gather the information you need.

For specific troubleshooting topics, see Chapter 7. Troubleshooting in the *Installation and Configuration Guide*.

General information

When a problem occurs with IBM Tivoli Storage Productivity Center, you should gather the following information:

- An exact description of the problem.
- The function being used.
- The sequence of steps that resulted in the problem.
- The expected results from the failing step.
- Any error messages that you see.
- The date and time when the problem occurred.
- The log files collected with the service utility.
- ESS microcode level or SAN Volume Controller version, if known. The microcode level can be obtained by looking at the properties of the discovered device.
- ESS, SVC, DS4000, DS5000, DS6000, or DS8000 CIM Agent version, if known.
- The last time that inventory collection was performed (this indicates that the repository is in synchronization with the real configuration).
- Whether the error is repeatable or it occurs intermittently.
- The answers to the following connectivity questions help you determine whether there is a communication issue between IBM Tivoli Storage Productivity Center and the CIM Agent server:
  - Is there any firewall enabled interfering with the communication between the IBM Tivoli Storage Productivity Center components and the CIM agents?
  - Is it possible to ping the CIM agents?
  - Is it possible to telnet to the CIM agents?
  - Is it possible to contact the CIM agents with a CIM browser?
  - Are the CIM agents locate in the local subnet with or is a DA in another subnet used for discovery?
  - Is the DA on the remote subnet registered with IBM Tivoli Storage Productivity Center?
  - Was the CIM agents SLP registration successful? Use the **slptool findsrvs** command.

When you schedule a job (for example, discovery, probe, scan, and so forth), the actual start time of the job will start one minute earlier (which is the default) than the scheduled time. This means, that if you schedule a job to run at 1:30 PM, the job will actually start at 1:29 PM. The 1:29 PM timestamp will show up in the GUI for that job and in the job log files also. This allows for the overhead time involved in scheduling a job so that the job will actually start close to the scheduled time.

Discovery

For discovery problems, there are several service level logs to check:

- msgTPCDeviceServer.log
- traceTPCDeviceServer.log
- dmSvcTrace.log
- tracePerfMgr.log

- TPCZoneControl.log

These logs are located in this directory:

```
<TPC_install_directory>\IBM\TPC\device\log
```

You can get a core dump file in this directory:

```
<TPC_install_directory>\device\apps\was
```

For Windows, follow these steps:

1. Go to this directory:

   ```
   C:\Program Files\IBM\TPC\device\apps\was\bin
   ```

2. Enter the following:

   ```
   wsadmin set jvm [$AdminControl completeObjectName
   type=JVM,process=server1,*] $AdminControl invoke $jvm dumpThreads
   ```

For AIX or Linux, follow these steps:

1. Enter this command:

   ```
   ps -ef | grep "IBM/TPC" | grep "apps/was" | awk '{print $2}' | xargs kill -3
   ```

To discover the job activity, issue the following command:

```
srmcp -u <user_ID> -p <password> DiscoverService list jobs
```

To determine the DB2 activity, issue the following DB2 commands:

```
db2 list applications show detail > dbListApplicationDetails.out
db2 update monitor switches using statement on lock on table
   on uow on
db2 connect to tpcdb
db2 get snapshot for all on tpcdb > dbsnap.out
```

Common user errors

A discovery or probe operation does not complete. This can be caused by the following:

- There is a lock contention when IBM Tivoli Storage Productivity Center accesses the database tables.
- DB2 does not return information from a query.
- The CIM agent starts returning information and then stops.

To work around this problem, stop and then restart the Data server and Device server. Also collect log information using the service tool. For information about the service tool, see “Service tool: collecting information” on page 9.

Monitoring service

The monitoring service provides a notification of events like the CIM indication and SNMP trap and fabric in-band events from the underlying devices to other IBM Tivoli Storage Productivity Center components.

IBM Tivoli Storage Productivity Center provides an event or alter publishing service for the discovery and control components when you add a new device, the device status changes, or the device is not detected. IBM Tivoli Storage Productivity Center provides an alert generation on behalf of the performance monitoring service when the threshold changes and there is a collection failure. The database table snapshot for a certain device type is based on the performance monitoring service’s request.

For general diagnosis of the problem, check these items:

- Make sure the alert trigger condition is set up correctly.
- Make sure that the changes did happen.
- Check the alerts displayed in the GUI alert panel.

For specific diagnosis of the problem, check these items:

- Locate the Device Server trace log and search for possible exceptions generated from the monitor service.
- If no errors and exceptions are found for the monitor service, it means that the monitor service is running OK. Next, check to see if some other component is causing the problem such as the Change Detection component.

Common usability problems

These are some common user errors:

Why do I not get the right alert?
Check to see if you have configured the correct trigger condition for the right device or device type.

Why is overall performance slow?
A large number of database activities to form SNMP and TEC events might affect overall performance if too many alerts get created. You can configure only the needed alerts to control the number of alerts.

Why am I losing an event or alert?
An event or alert might be lost because the required data is not available in the database or there are network communication problems. Once the alert creation fails, the alert for this changed device will not be recreated.

Logs and traces

When a problem occurs, you can set the trace level parameter `san.eventFactoryTrace.level = INFO` using the CLI command to get detailed trace information. Make sure that the created FabricAlert object is sent to the Data Server successfully.

The detailed trace information is saved in the following directory:

```
<TPC_install_directory>\data\log\traceTPCDeviceServer.log
```

Use the **service** tool and **repocopy** tool to collect trace information for IBM Tivoli Storage Productivity Center and the database.

When an expected alert is not presented, do the following:

- Check the trace log to make sure the device changes is reported to the monitor service. For example, search the log for the corresponding API call.
- Check to see if the right alert populator is used and any other exception is encountered during the process.
- Make sure that the alert object is sent to the Device Server.

Performance monitoring

This section provides information about performance monitoring issues.

Tuning switch performance for collecting data

When you are using the switch performance monitor and you run into timeout problems, you can change the attributes that affect the performance monitor. The performance monitor uses a couple of algorithms to collect performance information. The association algorithm is optimized for environments with larger numbers of switches managed by a CIM Agent and the enumeration algorithm is optimized for environments with fewer switches managed by a CIM Agent. IBM Tivoli Storage Productivity Center attempts to determine which algorithm to use based on the fabric configuration.

You can change the following attributes in the `pm.conf` file:

MinSwitchPortRatio
This attribute defines the minimum port ratio (for example, the ratio of monitored ports to total ports for a switch CIMOM). This value is used to determine which algorithm to use for querying the performance statistics data for the ports through the CIM Agent. If not set, the default cutoff value for the enumeration algorithm is 40%, so that if less than 40% of a CIMOM's ports are targeted for performance data collection, the association algorithm is used instead, provided that the **MaxSwitchPortLimit** is not exceeded (see **MaxSwitchPortLimit** below). This does not apply to Cisco CIMOMs.

MaxSwitchPortLimit
This attribute defines the maximum port limit (for example, the maximum number of ports to be monitored using the association algorithm for switch CIMOMs). This value is used to determine which algorithm to use for querying the performance statistics data for the ports through the CIM Agent. If not set, the default cutoff value for the association algorithm is 256 ports, so that performance data will not be collected for more than 256 ports simultaneously using the association algorithm. Instead, the enumeration algorithm will be used for any CIMOMs after the limit has been reached.

To change these attributes in the `pm.conf` file, go to the following directory:

```
<TPC_install_dir>/device/conf/pm.conf
```

Remove the pound sign (#) from these attributes and modify your setting:

```
#com.ibm.tpc.perf.MinSwitchPortRatio = 0.4
#com.ibm.tpc.perf.MaxSwitchPortLimit = 256
```

Save the file. The Device server needs to be restarted for these changes to take effect.

Topology viewer

This section provides information about how to troubleshoot the topology viewer.

Using debug mode for the topology viewer lets you see the following:

- The entire data model sent by the Data Server to the GUI.
- Any exceptions that occurred.
- A special debug ID column in the tabular view. Each entry contains:
  - The prefix for the entity (″server″ or ″node″, and so forth).
  - The ID of that entity. You can use this to easily look up that entity in the database tables.

The data model lets you see exactly what will be rendered on the screen. If something looks wrong in the GUI, check the data model first:

- If the data model contains the correct information, it is a GUI defect.
- If the data model contains the wrong information, you need to check the database. Either the underlying data is incorrect, or the Data Server portion of the topology viewer is not handling the data correctly.

To enable debug mode, a flag needs to be passed to Java when launching the GUI:

```
-DTopologyViewerDebug=5
```

This will cause the topology viewer to provide debug information to `stdout` (and show an extra debug ID column in the tabular view). Note that if IBM Tivoli Storage Productivity Center is launched using `javaw.exe` (as used by the default in the Windows shortcut), you will not see any output because `javaw.exe` hides the console window. In this case, you need to change the shortcut to use `java.exe`.

An easier way to route the debug information to a file is to use the following:

```
-DTopologyViewerDebugOutputFile="C:\topo.log"
```

You can then send the file to IBM support.

Here is an example of how to edit the target portion of the Windows shortcut for IBM Tivoli Storage Productivity Center. The following example shows this.

```
"C:\Program Files\IBM\TPC\jre\bin\javaw.exe"
-Xmx512m -classpath
"C:\Program Files\IBM\TPC\gui\TSRMgui.jar"
com.tivoli.itsrm.gui.GUiMain
wanda.tpc.storage.sanjose.ibm.com:9549
```

To generate the log file `C:\topo.txt`, change the ″target″ field in the shortcut properties to this:

```
"C:\Program Files\IBM\TPC\jre\bin\javaw.exe"
-DTopologyViewerDebug=5
-DTopologyViewerDebugOutputFile="C:\topo.txt"
-Xmx512m -classpath
"C:\Program Files\IBM\TPC\gui\TSRMgui.jar"
com.tivoli.itsrm.gui.GUiMain
wanda.tpc.storage.sanjose.ibm.com:9549
```

For UNIX or Linux, you need to edit the `TPCD.sh` script in the following directory:

```
/<usr or opt>/IBM/TPC/gui/
```

Add the same flags. The `TPCD.sh` script also contains some characters at the end that pipe the output elsewhere. At the end of the script, remove the sequence of characters that look like this:

```
>/dev/null 2>&1
```

Common usability problems

Note the following when linking the topology viewer with IBM Tivoli Storage Productivity Centerreports and alert logs:

- The user will not be offered a link to asset reports if they have selected more than one entity (for example, computer, and storage subsystem).
- Links to alerts show all alerts, not just any that happen to be selected in the topology viewer.

- When linking to ″manage storage subsystems″ or ″manage zone control″, if the user has selected storage subsystems or fabrics, only one selected storage subsystem or fabric will be highlighted in the panel.

Single sign-on

Single sign-on is an authentication process that enables you to enter one user ID and password to access multiple applications. For example, you can access 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 in context feature to enable you to move smoothly from one application to a functionally-related location in a second application. This section provides information on how to troubleshoot issues with single sign-on.

Log files

When you have issues with single sign-on, make sure that tracing is turned on for single sign-on. You will see these parameters in the DataStore.properties file in directory <TPC_install_directory>/device/conf:

- san.SingleSignonServiceTraceLogger.level=ALL
- san.SingleSignonServiceTraceLogger.logging=true
- san.SingleSignonServiceTraceLogger.listenerNames=file.trace
- san.SingleSignonServiceTraceLogger.LoggerType=TraceLogger

Traces will be in the following file:

```
<TPC_install_directory>/device/log/traceTPCDeviceServer.log
```

In addition, you should collect the Device server WebSphere log files and first-failure data capture (FFDC) files. Look in this directory:

```
<TPC_install_directory>/device/apps/was/profiles/deviceServer/logs
```

If you have an issue with synchronization of the products, browse through the trace log files, and make sure that loginID and loginPassword properties are set correctly in the soap.client.props file. For IBM Tivoli Integrated Portal, this file is located in the following directory:

```
<TIP_install_directory>/profiles/TIPProfile/properties
```

Storage Optimizer

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. The Storage Optimizer helps you develop storage migration or storage consolidation plans, and helps you plan for the growth of your storage infrastructure.

Note:

- The Storage Optimizer does not actually perform any migrations or make any modifications to subsystem configurations. Its primary purpose is to provide you with a performance analysis and optimization recommendations that you can choose to implement at your discretion.
- Non-IBM subsystems, including SAN Volume Controllers that use non-IBM back-end subsystems, are not supported by Storage Optimizer. Solid state drives (SSDs) are not supported at this time. Subsystems that contain solid state drives

may produce unexpected results in the Storage Optimizer heat maps. Storage Optimizer analysis of subsystems containing solid state drives should be avoided until support is available.

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

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

The following IBM storage subsystems or applications are supported:

- DS8000
- DS6000
- DS4000
- Tivoli Storage Enterprise Storage Server
- SAN Volume Controller

Before running Storage Optimizer, you must set up performance monitors and collect performance monitoring data for all storage subsystems that you want Storage Optimizer to analyze. You must also collect performance monitoring data for a SAN Volume Controller's back-end subsystems in order to produce the most accurate Storage Optimizer analysis.

Storage Optimizer analyzes performance monitoring data for the time interval you specify. The analysis will be less accurate if there are any gaps in the data collection for the time interval being analyzed. It is recommended that you collect at least one week's worth of performance monitoring data before using Storage Optimizer. Providing a longer time interval for data collection will increase the accuracy of the Storage Optimizer analysis and recommendations. To create a subsystem performance monitor job, navigate to **Disk Manager > Monitoring > Subsystem Performance Monitors**. Right-click **Subsystem Performance Monitors** and choose **Create Subsystem Performance Monitor**.

You must provide the Storage Optimizer with the following input:

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

The Storage Optimizer provides the following output:

- An analysis report that displays performance heat maps and tables that graphically illustrate the performance utilization of the storage subsystems and pools 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 performance monitoring data over a longer time interval before using the Storage Optimizer.

Because 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.

Log files

The log files for the Storage Optimizer are found in these files:

AnalyzerJobList.log
This file is found in the following directory:

```
<TPC_install_directory>\device\log\msg.control.<nnn>.AnalyzerJobList.log
```

where *nnn* is a generated number.

OptimizerJobList.log
This log file can provide information about why no recommendations were made. This file is found in the following directory:

```
<TPC_install_directory>\device\log\msg.control.<nnn>.OptimizerJobList.log
```

where *nnn* is a generated number.

Device server trace logs

The DataStore.properties file contains these default values for the Storage Optimizer:

```
san.OptimizerMsgLogger.listenerName=file.message
san.OptimizerMsgLogger.logging=true
san.OptimizerMsgLogger.LoggerType=MessageLogger

san.OptimizerTraceLogger.level=WARN
san.OptimizerTraceLogger.listenerNames=file.trace
san.OptimizerTraceLogger.logging=true
san.OptimizerTraceLogger.LoggerType=TraceLogger
```

The DataStore.properties file found in the following directory:

```
<TPC_install_directory>\device\conf
```

Storage Optimizer

Use this section to troubleshoot and resolve Storage Optimizer problems.

Optimization report contains no recommendations:

The Storage Optimizer optimization report contains no recommendations.

Problem

The Storage Optimizer optimization report contains no recommendations. Some typical reasons for this are the following:

- the performance threshold value used as input to the optimization report was set too low or too high
- there is not enough disk space in the target pools, so none of the source pools can be migrated
- the source and target storage subsystems are incompatible

Furthermore, if "Space" is the only utilization metric that exceeds the specified performance threshold, Storage Optimizer will not provide recommendations based

on space alone, since it is not designed to be used as a space planner.

Action

Storage Optimizer attempts to provide optimization recommendations based on the inputs provided, but it may not be able to create recommendations in all cases. Keep in mind that generating an optimization report is an iterative process. Try one or more of the following actions to increase the likelihood that recommendations will be created:

- specify different values for the performance threshold
- specify different combinations of source and target entities

Performance heat map cells are all white:

The heat map in a Storage Optimizer analysis report consists only of white cells, which means that no data was returned.

Problem

Each cell in a Storage Optimizer heat map represents a storage pool. If all the cells in an analysis report heat map are white, then some possible reasons are:

- no data was returned by the storage subsystem performance monitor used as input in the analysis
- data values of zero were returned by the subsystem performance monitor for the time interval chosen for the Storage Optimizer analysis

Action

To determine if subsystem performance data exists for the time interval chosen for the Storage Optimizer analysis, generate storage subsystem performance reports, as follows:

1. Navigate to **Disk Manager > Reporting > Storage Subsystem Performance**.
2. Using daily summation, generate the following storage subsystem performance reports:
   - By Storage Subsystem
   - By Array
   - By Volume
3. Include the following columns in each report:
   - Read I/O rate (normal)
   - Read I/O rate (sequential)
   - Read I/O rate (overall)
   - Write I/O rate (normal)
   - Write I/O rate (sequential)
   - Write I/O rate (overall)
   - Read Cache Hits Percentage (overall)
   - Write Cache Hits Percentage (overall)
   - Cache to Disk Transfer Rate
4. Use the information provided in the reports to determine if you need to adjust the time interval used for the Storage Optimizer analysis, make other adjustments to the Storage Optimizer analysis settings, or make adjustments to the settings used for subsystem performance monitors.

5. Run the Storage Optimizer analysis again.

Analysis report includes additional subsystems:

The analysis report for a SAN Volume Controller includes additional subsystems.

Problem

In the analysis report for a SAN Volume Controller, additional subsystems are listed. These additional subsystems are included in the analysis report because they are the back-end storage to the SAN Volume Controller.

If no performance data was collected for these back-end subsystems, or if the back-end subsystems are non-IBM devices, the utilization percentages for these subsystems will be listed as ″N/A″ (not applicable) in the analysis report.

Space utilization is listed as 0%:

Use this section to understand why space utilization is listed as 0%.

Problem

When space utilization is listed as 0%, some reasons for this are the following:

- There is no managed capacity (no possibility to use the volumes as a target)
- There is managed capacity, but no volumes (there are no volumes to use as a target)

No recommendations will be made based on space utilization alone.

Reports with target volume columns that are empty or not empty:

In the reports for the Storage Optimizer, there might be target volume columns that are empty or not empty. This section provides an explanation for these reports.

Problem

When the Target Volume column is empty for a storage system, this indicates that you need to create a new volume in the Target Pool.

When the Target Volume column is not empty, this indicates that one of the listed volumes is unassigned and has the same characteristics as the Source Volume.

Rollup reports

This section provides information about what to look for when diagnosing problems with rollup reports.

Rollup reports let you view the storage information collected by the master IBM Tivoli Storage Productivity Center server from its associated subordinate servers. Rollup reports enable you to have a network-wide perspective of storage usage in an environment where multiple IBM Tivoli Storage Productivity Center servers are deployed and collecting storage metrics.

Things to note

These are some things to note:

- The master server should not monitor more than 500 unique sources.
- Subordinate servers should not have more than 1200 unique data sources. This includes Data agents, Fabric agents (in-band and out-of-band), CIM agent, and VM agents.
- You must have IBM Tivoli Storage Productivity Center superuser or IBM Tivoli Storage Productivity Center administrator authority to perform administrative operations such as adding or deleting a subordinate server to or from the master server list.
- For rollup reports, a master server collects the following information from subordinate servers:
  - Asset information: detailed statistics about agents (IBM Tivoli Storage Productivity Center), computers, storage subsystems, disk and volume groups, disks, file systems, logical volumes, LUNs, and fabrics.
  - Database asset information: detailed statistics about the RDBMSs in your environment, including Oracle, SQL Server, Sybase, and DB2 UDB.
  - Capacity information: storage metrics related to the disk capacity, file system capacity, file system used space, and file system freespace of the storage entities in your environment.
  - Database capacity information: storage metrics related to the storage capacity of the RDBMSs in your environment, including Oracle, SQL Server, Sybase, and UDB.

  Additional storage information collected by a subordinate server and used within IBM Tivoli Storage Productivity Center functions such as the topology viewer, data path explorer, volume provisioning, volume performance, SAN planner, and so forth is available for that subordinate server **only**. That information is not collected by the master server and thus not rolled up into the reports available through the master server.

Diagnosing the problem

The first thing to look at to troubleshoot rollup reports is to look at the log files:

- Data Server log file on the master server:

  ```
  <TPC_install_dir>/data/log/server_xxxxxx.log
  <TPC_install_dir>/data/log/TPCD_xxxxxx.log
  ```

  You can look at these log files from the GUI:

  1. Open the IBM Tivoli Storage Productivity Center GUI and navigate to: **Administrative Services > Services > Data Server > Server**.
  2. For the server, right-click on **Server**. Click on **View Log**. The log will be displayed in the right pane.
  3. For the GUI, right-click on **GUI**. Click on **View Log**. the log will be displayed in the right pane.
- Data Server trace log file on the master server:

  ```
  <TPC_install_dir>/data/log/ServerTrace.log
  ```

  Turn on the trace function through the GUI. For information about how to turn on the trace function, see "Tracing the servers and agents" on page 11.
- Device Server log on the subordinate server:

  ```
  <TPC_install_dir>/device/log/msgTPCDeviceServer.log
  ```

- Device Server trace log on the subordinate server:
  `<TPC_install_dir>/device/log/traceTPCDeviceServer.log`

You will need to run **repocopy** in most cases to gather information about your environment. IBM Support Center will probably request that you run **repocopy**. You will need to run **repocopy** for both the master server and subordinate server.

SAN Planner

This section provides information about what to look for when diagnosing problems with the SAN Planner.

The SAN planner assists the user in end-to-end planning involving fabrics, hosts, storage controllers, storage pools, volumes, paths, ports, zones, and zone sets . Once a plan is made, the user can select to have the plan implemented by the SAN planner.

Dependencies

For the Volume Planner, the dependencies are:

- Requires an ESS, DS6000, or DS8000 storage subsystem
- Storage subsystems must previously be discovered and probed
- By default, a subsystem performance monitor must have been previously run for the storage device or devices you are creating a plan for. You can create a plan without performance data by specifying the ″Space Only″ workload profile.

For the Path Planner, the dependencies are:

- Requires a Linux, AIX, HP-UX, Solaris, or Windows host with a Data agent installed on the host, and the host has been probed.
- SDD must be installed on the host:
  - On UNIX or Linux hosts, any version of SDD is supported (no SDDPCM support)
  - On Windows, SDD 1.6.1 or higher is required
  - The host needs at least two channel ports if you are using the Path Planner.

For the Zone Planner, the dependencies are:

- The out-of-band Fabric agent or CIM agent must be managing the switch or fabric
- In-band Fabric agent is required on the host
- A fabric probe must have occurred before Zone Planner use

Diagnosing the problem

To turn tracing on for the SAN Planner in the Device server, follow these steps:

1. Run the **srmcp** tracing command.

   Windows

   To run the **srmcp** tracing command on Windows, follow these steps:

   a. Run the following command: **setenv**.
   b. Go to this default directory on Windows:
      `C:\Program Files\IBM\TPC\device\bin\w32-ix86`
   c. Run this command:

```
srmcp -u <user_id> -p <password> log set
san.PlannerTraceLogger -filterkey INFO
```

The filter key can be: INFO, WARN, or ERROR.

UNIX

To run the **srmcp** tracing command on UNIX, follow these steps:

a. Go to this default directory on UNIX:

```
/<usr or opt>/IBM/TPC/device/bin/aix (for AIX)
/<usr or opt>/IBM/TPC/device/bin/linux (for Linux)
```

b. Run this command:

```
./setenv.sh
```

c. Run this command:

```
./srmcp.sh -u <user_id> -p <password> log set
san.PlannerTraceLogger -filerkey INFO
```

The filter key can be: INFO, WARN, or ERROR.

The output will be logged in `traceTPCDeviceServer.log`.

If you receive an "Unexpected error occurred." GUI message, look for the exception in the following log:

`<TPC_install_dir>/data/log/server_0000xx.log`

For IBM support to recreate a problem, a **repocopy** of the database is probably required.

Configuration History and Analysis

This section provides information about what to look for when diagnosing problems with Configuration History and Analysis.

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.

When you set up the Configuration Analysis tool for policy checking, you define the scope to perform the checking:

- All Fabrics
- One Fabric
- One Zoneset

The Configuration Analysis tool can check for these policies (the policies supported depend on the scope you select).

The Configuration History views can show an entity two or more times. This may happen when an entity gets removed manually from the database or by the removed resource retention function, and IBM Tivoli Storage Productivity Center rediscovers that entity again. IBM Tivoli Storage Productivity Center assigns the entity to a different ID when it is rediscovered.

Common usability problems

These are some common usability problems:

- Policy violations are not correctly sorted in the Alert Overlay within the topology viewer.
- When a missing fabric is selected as a scope in a configuration analysis job, then no violations will be reported for policies that need zoning information (policies 3, 4, 5, 6, 12 and 13). The zone information gets removed from the IBM Tivoli Storage Productivity Center database when a fabric goes missing.
- By default the limit of 50 policy violations are stored per analysis job. This is done to prevent overloading the topology viewer with a lot of violations. You can change this policy violation limit by using the **tpctool** command. For example, if you wanted to change the policy violations limit to 200, run this command:

```
tpctool setdscfg -user <user_ID> -pwd <password>
-url localhost:9550 -property MaxPolicyViolationsStored 200
```

Dependencies

Configuration analysis requires the following:

- A mixture of in-band Fabric agents, out-of-band Fabric agents, and CIM agents for the fabric to collect all fabric topology information (switches, interconnections, and zoning). This information is needed for the evaluation of all configuration analysis policies.
- CIM agents for storage subsystems to collect the storage subsystem information (subsystem details, storage volumes, storage ports, and storage volumes to host mappings). The information is needed for the evaluation of configuration analysis policies 2, 3, 5, 7, and 11.
- Data agents to collect the host information (operating system type and version). This information is needed for the evaluation of configuration analysis policies 3 and 9.
- In-band Fabric agents to collect the attribute information from the host information (HBA and host port details). This information is needed for the evaluation of configuration analysis policies 2, 3, 4, 8, and 9.
- CIM agents for tape to collect the tape information. This is needed for the evaluation of configuration analysis policy 2.

Limitations

The following items are not supported by Configuration History:

- Alert and performance overlays
- Data Path Explorer view
- Pin and refresh settings actions

Diagnosing problems

Check the following when you encounter a problem with Configuration History and Analysis:

- Configuration Checker job logs are located here:

  `<TPC_install_dir>/data/log/*config_analysis*.log`
- Message log for the Device Server. The log file is at this location:

  `<TPC_install_dir>/device/log/msgTPCDeviceServer.log`

  The message prefix is ″HWNCC.″
- Device Server trace log.

To turn tracing on for the SAN Planner in the Device server, follow these steps:

1. Run the **srmcp** tracing command.

   Windows

   To run the **srmcp** tracing command on Windows, follow these steps:

   a. Run the following command: **setenv**.
   b. Go to this default directory on Windows:

      ```
      C:\Program Files\IBM\TPC\device\bin\w32-ix86
      ```

   c. Run these commands:

      ```
      srmcp -u <user_id> -p <password> log set
      san.ConfigCheckerTraceLogger -filterkey INFO

      srmcp -u <user_id> -p <password> log set
      san.ConfigRoverTraceLogger -filterkey INFO
      ```

      The filter key can be: INFO, WARN, or ERROR.

   UNIX

   To run the **srmcp** tracing command on UNIX, follow these steps:

   a. Go to this default directory on UNIX:

      ```
      /<usr or opt>/IBM/TPC/device/bin/aix (for AIX)
      /<usr or opt>/IBM/TPC/device/bin/linux (for Linux)
      ```

   b. Run this command:

      ```
      ./setenv.sh
      ```

   c. Run this command:

      ```
      ./srmcp.sh -u <user_id> -p <password> log set
      san.ConfigCheckerTraceLogger -filerkey INFO

      ./srmcp.sh -u <user_id> -p <password> log set
      san.ConfigRoverTraceLogger -filerkey INFO
      ```

      The filter key can be: INFO, WARN, or ERROR.

   The output will be logged in `traceTPCDeviceServer.log`.

- Data Server trace log. For information about enabling tracing for the Data Server, see "Tracing the servers and agents" on page 11. Select DEBUG_MAX when enabling tracing. The log file is at this location:

  ```
  <TPC_install_dir>/data/log/ServerTrace*.log
  ```

Data Path Explorer

This section provides information about what to look for when diagnosing problems with Data Path Explorer.

The Data Path Explorer is a new type of view in the topology viewer. Data path explorer combines several of the usual topology views into a visualization that allows users to follow connectivity from one end point in the topology view to another. Storage administrators can use this view to debug connectivity and performance problems in the storage environment. Data Path Explorer will display the different data paths (SAN access paths or I/O paths) for each host disk. This allows you to use the Data Path Explorer to check the health and performance statistics of the storage components along the I/O path.

Things to note

When you have the problem of data paths not appearing in the topology viewer, check these items:

- Are the storage volumes visible as disks on the initiator?
- Are the agents configured correctly for the host, fabric, and storage subsystem?
- Have probes been run on these agents?
- Using the L2:Computers view and L2:Subsystem view, check which switches (and fabric) the host and subsystem are connected to.
- Does the L2:Fabric view show the interconnection between the switches shown above?

Common usability problems

The entities and connections involved in a data path segment are not highlighted in the Data Path Explorer graphical view when a data path segment is selected in the Data Path tabular view.

Diagnosing problems

Check these log files:

- Message log. The message log can be accessed through the Tivoli Storage Productivity Center GUI. The log file is located here:

  ```
  <TPC_install_dir>/data/log/server_*.log
  ```

  To look at the log file, follow these steps:

  1. Open the Tivoli Storage Productivity Center GUI and navigate to: **Administrative Services > Services > Data Server > Server**.
  2. For the server, right-click on **Server**. Click on **View Log**. The log will be displayed in the right pane.
  3. For the GUI, right-click on **GUI**. Click on **View Log**. the log will be displayed in the right pane.
- Server trace log. For information about how to enable tracing, see "Tracing the servers and agents" on page 11. When enabling tracing, specify DEBUG_MAX. The log file is located here:

  ```
  <TPC_install_dir>/data/log/ServerTrace*.log
  ```
- GUI trace log.

  Here is an example of how to edit the target portion of the Windows shortcut for Tivoli Storage Productivity Center. The following example shows this.

  ```
  "C:\Program Files\IBM\TPC\jre\bin\javaw.exe"
  -Xmx512m -classpath
  "C:\Program Files\IBM\TPC\gui\TSRMgui.jar"
  com.tivoli.itsrm.gui.GUiMain
  wanda.tpc.storage.sanjose.ibm.com:9549
  ```

  To generate the log file `C:\topo.txt`, change the ″target″ field in the shortcut properties to this:

  ```
  "C:\Program Files\IBM\TPC\jre\bin\javaw.exe"
  -DTopologyViewerDebug=5
  -DTopologyViewerDebugOutputFile="C:\topo.txt"
  -Xmx512m -classpath
  "C:\Program Files\IBM\TPC\gui\TSRMgui.jar"
  com.tivoli.itsrm.gui.GUiMain
  wanda.tpc.storage.sanjose.ibm.com:9549
  ```

SMI-S fabric probe

This section provides information about what to look for when diagnosing problems with the SMI-S fabric probes.

IBM Tivoli Storage Productivity Center supports the SMI-S Agent (CIM agent) for the fabric. The SMI-S Agent collects basic fabric, switch, and port information used for performance monitoring. For Brocade and McDATA fabrics, full topology information and zoning information is also collected during the fabric probe through SMI-S. For other switch vendors, monitoring of the fabric topology requires the IBM Tivoli Storage Productivity Center in-band Fabric agent or out-of-band Fabric agent.

For a list of switches and directors supported by IBM Tivoli Storage Productivity Center, see http://www.ibm.com/software/sysmgmt/products/support/IBMTotalStorageProductivityCenterStandardEdition.html. Search for **4.1 switch & director**.

This section provides information about what to look for when diagnosing problems with the SMI-S fabric probes.

Common usability problems

Some of the common usability problems are:

Switch Vendor and Model is blank or unknown
Some switch vendors do not provide a distinguishable vendor name or model name through some of the fabric agents.

Switch performance monitoring job fails immediately
Check to see if a fabric probe has been run.

Switch performance monitoring job fails
A switch performance monitoring job may fail for switches that are known to the McDATA or Cisco CIMOM but are not managed by the CIMOM.

For McDATA SMI-S Agent Direct Mode
To get detailed information for McDATA, make sure that users have the SMI-S Agent configured to manage each switch in the fabric. IBM Tivoli Storage Productivity Center will gather switch information from the SMI-S Agent.

For McDATA SMI-S Agent Proxy Mode
To get detailed information for McDATA, make sure that users have EFCM configured to manage each switch in the fabric. EFCM communicates with the SMI-S Agent and IBM Tivoli Storage Productivity Center will gather switch information from the SMI-S Agent.

For Cisco
To get detailed information for Cisco switches, make sure that users have IBM Tivoli Storage Productivity Center configured to manage the SMI-S agent for each switch in the fabric. This will allow you to collect information about FC ports. Cisco switches have embedded SMI-S agents.

Cisco logical switches and VSANs are missing
Check to see if the out-of-band Fabric agents for the Cisco switches are configured.

"Switch port", "Index", "Slot", or "Enabled State" is blank for the Brocade switch
Check to see if a n SMI-S agent has been configured for this fabric, and check to see that a fabric probe has been run.

CIMOM discovery failed
Check to see if the switch CIMOMs are at the supported levels.

Brocade SMI Agent

For information about the Brocade SMI Agent for FOS switches, see *Brocade SMI Agent User's Guide Supporting SMI Agent 120.8.0* at the Brocade Web site. For information about the Brocade SMI agent for EOS switches (also known as McDATA), see *Brocade SMI Agent for EOS Products User Guide 2.6* at the Brocade Web site.

Things to note

- When a fabric probe is performed, much topology information is collected redundantly from the in-band Fabric agent and the out-of-band Fabric agent. However, the following information is collected only through the CIMOM:
  - Brocade blades
  - Brocade Switch FCPort EnabledState and Index
- For Cisco switches, VSAN and logical switches are not collected through the CIMOM. You must use the out-of-band Fabric agent to collect this information.
- When a CIMOM discovery is performed, IBM Tivoli Storage Productivity Center finds fabrics and switches through the switch CIMOMs. IBM Tivoli Storage Productivity Center finds basic information for switches that are not registered with the CIMOM, if they are in a fabric with a switch that is registered with the CIMOM. IBM Tivoli Storage Productivity Center also finds and persists fabrics and the fabric-to-switch relationships.
- When a fabric probe is performed, the following occurs:
  - Out-of-band Fabric agent and in-band Fabric agent scans are performed.
  - Runs the SMIS probe for Brocade and McDATA switches:
    - Re-runs the CIMOM discovery algorithm for fabrics to get the latest information on fabric and switches.
    - Performs requests to get switch FCPorts, blades, port connections, and zoning information.
  - Runs the SMIS probe for Cisco and QLogic switches:
    - Re-runs the CIMOM discovery algorithm for fabrics to get the latest information on fabric and switches.
    - Performs requests to get switch FCPorts.
  - For Brocade and McDATA CIMOMs, it is best to group fabrics for a fabric probe by the CIMOM managing the fabric.

VMware ESX

This section provides information about what to look for when diagnosing problems with VMWare ESX servers.

IBM IBM Tivoli Storage Productivity Center now supports the VMWare Virtual Infrastructure which consists of the ESX Server and VMWare VirtualCenter. 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 a data center's ESX Servers. You would install a Data agent on each virtual machine that you want to monitor.

IBM Tivoli Storage Productivity Center supports the following:

- The Virtual Center discovery is the only successful criteria.
- Alerts are for probe only (no event, traps, or alarms from the Virtual Infrastructure).
- Reporting only, no active management of the Virtual Infrastructure through IBM Tivoli Storage Productivity Center.
- The information from VMWare is equivalent of a read-only Data agent (no scripts, no FS extension).

To get complete information for space and capacity reports, you must do the following:

- Probe all the ESX servers.
- Install Data agents on all the virtual machines and probe them. Otherwise, you get a partial picture with incomplete capacity data.

Common usability problems

These are some common user errors:

- Did the user import the SSL certificates from VMWare for secure communication? When the SSL certificate is not properly defined, the "Add VMWare VI Data Source" fails.
- There are new totals in the dashboard and reports.
- Note the indications for the hypervisor status.
- Differentiate information from a Data agent for a virtual machine versus the ESX Server.

Dependencies

These are some of the dependencies:

- VMWare ESX 3.0.1 and Virtual Center 2.0.1
- The hypervisor and Data agent probes are required for complete reporting.
- Hypervisors and VirtualCenters have https communication enabled and http communication disabled by default.
- If you are using the https protocol, make sure that you import the VMWare certificate into the trust store before selecting the https protocol:
  - The trust store needs to be created under `<TPC_install_dir>/device/conf/vmware.jks`.
  - The VMWare hypervisor's certificates can be found on the hypervisor under the location `/etc/vmware/ssl/rui.crt`.
  - The VMWare VirtualCenter's certificate can be found in the following location:

    ```
    \Documents and Settings\All Users\Application Data\
    VMware\VMware VirtualCenter\SSL\rui.crt
    ```

  - To create and import the trust store, use **keytool** located in the following directory:

    ```
    <TPC_install_dir>/jre/bin/keytool
    ```

- The command to import the trust store is:

```
keytool -import -file rui.crt -alias <alias>
-keystore vmware.jks
```

Diagnosing the problem

When troubleshooting a VMWare problem, note the following:

- Look at the job logs for discovery and probe create jobs which may refer to tracing information.
- Look for trace messages in the general trace file:
  `<TPC_install_dir>/device/log/traceTPCDeviceServer.log`
- Set the trace level.

  To turn tracing on for the Device server, follow these steps:

  1. Run the **srmcp** tracing command.

     Windows

     To run the **srmcp** tracing command on Windows, follow these steps:

     a. Run the following command: **setenv**.
     b. Go to this default directory on Windows:
        `C:\Program Files\IBM\TPC\device\bin\w32-ix86`
     c. Run this command:
        ```
        srmcp -u <user_id> -p <password> log set
        san.VMManagerTraceLogger -filterkey INFO
        ```
        The filter key can be: INFO, WARN, or ERROR.

     UNIX

     To run the **srmcp** tracing command on UNIX, follow these steps:

     a. Go to this default directory on UNIX:
        ```
        /<usr or opt>/IBM/TPC/device/bin/aix (for AIX)
        /<usr or opt>/IBM/TPC/device/bin/linux (for Linux)
        ```
     b. Run this command:
        `./setenv.sh`
     c. Run this command:
        ```
        ./srmcp.sh -u <user_id> -p <password> log set
        san.VMManagerTraceLogger -filerkey INFO
        ```
        The filter key can be: INFO, WARN, or ERROR.

     The output will be logged in `traceTPCDeviceServer.log`.

Reporting groups

This section provides information about what to look for when diagnosing problems with reporting groups.

A reporting group is a user-defined set of resources and groups upon which you want to report. Reporting groups are available in Data Manager and Data Manager for Databases only. Use reporting groups to view storage information about:

- Resources that appear in different monitoring groups. Because a resource can appear in only one monitoring group at a time, you can use reporting groups to include resources from different monitoring groups and view reports on them.

- A subset of resources within a monitoring group. For example, you can create a reporting group for a specific subset of computers with the Payroll monitoring group and generate a report containing information about those computers.

Diagnosing problems

To diagnose problems for reporting groups, you can get an exception stack trace dump in the command window:

- For AIX or Linux, it will be the shell window
- For Windows, it will be the DOS window

You need to run the GUI manually using the following command:

```
"<TPC_install_dir>\jre\bin\java.exe" -Xmx256M
-classpath "<TPC_install_dir>\gui\TSRMgui.jar"
com.tivoli.itsrm.gui.GuiMain localhost:9549
```

You should also do the following:

- Turn tracing on in the GUI. For information about turning tracing on, see "Tracing the servers and agents" on page 11.
- Run **repocopy** of the IBM Tivoli Storage Productivity Center database. For information on how to run **repocopy**, see "Repository Copy tool" on page 6.
- Run the **service** tool for both the servers and common agents. For information on how to run **service** tool, see "Service tool: collecting information" on page 9.
- Get DB2 logs (`db2diag.log`).
- Get output of **db2 list applications** command. To get output, run these commands:

  ```
  db2 connect to tpcdb
  db2 list applications
  ```

Monitored Computer Storage Space reports

This section provides information about what to look for when diagnosing problems with Monitored Computer Storage Space reports.

Disk Storage On Storage Subsystems reports

The reports for Disk Storage On Storage Subsystems display the storage subsystem name from where storage being reported on originates. The storage subsystem name is one of the following:

Storage Subsystem Serial Number
The storage subsystem serial number is retrieved by a probe of a host. Depending on the storage subsystem from where the storage originates, this might be a series of characters and numbers unique to the storage subsystem. For some storage subsystems, this value is a hexadecimal value and requires conversion to match a decimal number. This value should only be displayed when IBM Tivoli Storage Productivity Center has not probed the storage subsystem from where the storage originates. If IBM Tivoli Storage Productivity Center has probed the storage subsystem, this is a problem with the way the storage subsystem is retrieved by the probe of the host, and requires an investigation of the SCSI data returned by the probe of the host and the probe results from the storage subsystem.

Probed Storage Subsystem Display Name
When a storage subsystem is probed, the display name of the storage

subsystem returned from the probe represents the native name stored in the SMI-S provider of the storage subsystem.

User Defined Storage Subsystem Display Name
A storage subsystem discovered through a CIMOM discovery might have the storage subsystem name altered by the user of IBM Tivoli Storage Productivity Center. This is a user-defined storage subsystem name to represent the storage subsystem that was probed.

Unknown Storage Subsystem
If IBM Tivoli Storage Productivity Center is unable to identify the storage subsystem, the storage subsystem name is set to unknown. This is a problem and requires an investigation of the SCSI data returned by the probe of the host.

All Disk Storage and Disk Storage On Storage Subsystem reports

The reports specific to computer disks and file systems or logical volumes have a **Correlated** column when displayed in the All Disk Storage reports or the Disk Storage On Storage Subsystem reports. This column indicates that IBM Tivoli Storage Productivity Center has been is able to relate the host storage directly to the storage subsystem storage if set to **Yes**. Therefore, the Disk Manager reports for Storage Subsystem Computer Views and Storage Subsystem Views are able to generate data for these rows. These rows that indicate the storage has been correlated can then be related back to the actual storage subsystem storage for a more detailed view of how the host storage is allocated.

If a computer disk, file system, or logical volume is reported in the Disk Storage On Storage Subsystem reports, but is not correlated, then it might be a result of one of the following reasons:

- The storage subsystem has not been probed.
- The storage subsystem storage has been unassigned from the host storage, but the host has not been probed. This would require a re-probe of the host.
- The storage subsystem storage was not correctly identified by the host probe. This is a problem and requires an investigation of the SCSI data returned by the probe of the host and the probe results from the storage subsystem.

Computer disk reports

The reports specific to computer disks have an **Overallocated** column. This column indicates an error with the probe of the host. IBM Tivoli Storage Productivity Center was not able to uniquely identify the disk storage across more than one host. This is a problem and requires an investigation of the SCSI data returned by the probe of the host.

Disks Without Serial Numbers reports

The reports for Disks Without Serial Numbers indicate that IBM Tivoli Storage Productivity Center was unable to query the device to extract a serial number that would uniquely identify the disk. This is know to occur on internal disks within a hypervisor and virtual machines that have storage created from internal disks within a hypervisor. If this problem occurs for other types of storage, then the probe log of the host needs to be investigated for messages relating to SCSI command errors.

Non-Disk Storage reports

The reports generated for Non-Disk Storage indicate that IBM Tivoli Storage Productivity Center was not able to relate the file system or logical volume to disk storage. This occurs for Network Attached Storage, remote file system mounts, and temporary file systems on the processor cache. These reports should only report on these types of file systems. If there are logical volumes or file systems that are not the expected types, then this indicates that the probe of the host was unable to relate the logical volumes back to disk storage. Investigate the host probe logs for messages relating to file system or logical volume errors.

Debugging the problem

If the reports indicate that there are probe problems, gather as much information as you can. Before you collect the service logs on the host system, turn tracing on and set the debug level to DEBUG_MID. Once you have configured tracing, the probe on the host system should be run again. If there is a problem with the correlation of the storage subsystem volumes or the storage subsystem name, then the IBM Tivoli Storage Productivity Center server logs should also be collected.

If there are problems with serial numbers that are not being reported, then the following messages provide guidance on investigating these problems:

```
STA0035W: SCSI command failure occurred on device <device_path>,
manufacturer <manufacturer>, model <model> Please check the message
manual for more information on the SCSI command failure.

STA0036I: Sense key: 0xsense keyASC:0xadditional data ASCQ: 0xqualifier.
```

For more information about these messages, see the messages in the Information Center.

FlashCopy

IBM Tivoli Storage Productivity Center FlashCopy support labels volumes as having a FlashCopy property of ″source″, ″target″, or ″none″ as appropriate. This ensures that the ″consumable volume space″ value displayed in the Data Manager system-wide asset reports by storage subsystem does not include capacity used by the FlashCopy target volumes.

This section provides information about what to look for when diagnosing problems with FlashCopy support.

IBM Tivoli Storage Productivity Center does not provide the following information:

- Additional information about FlashCopy relationships
- Allow the IBM Tivoli Storage Productivity Center user to create or manipulate FlashCopy relationships
- Provide a report (or any information) that links together the source and target volumes of a particular FlashCopy relationship

Diagnosing problems

Most issues arise from one of two causes:

Missing data
This occurs when the subsystem does not report the FlashCopy data in a

manner that IBM Tivoli Storage Productivity Center expects. This is most likely to occur in the event of the release of new microcode or CIM agents for a subsystem. Check the following:

- That the CIMOM knows of the FlashCopy relationships. Use the CIMOM browser.
- There is very little logging available for IBM Tivoli Storage Productivity Center for a such a situation. The best way to determine what the problem is: recreate the problem using breakpoints at the start of the subsystem probe (to verify that the information is being discovered by IBM Tivoli Storage Productivity Center) and at the start of the Storage Volume mapper (to verify that it is being persisted properly).

Improperly created FlashCopy relationships

This occurs when the user did not create a persistent FlashCopy relationship. If the relationship is temporary, chances are it will be gone by the time IBM Tivoli Storage Productivity Center's probe requests the information. You should verify that the subsystem's management software Web site still reports the FlashCopy relationship.

Element Management

The Element Management is integrated with the IBM Tivoli Storage Productivity Center GUI. This allows the user to manage multiple DS Element Managers within the IBM Tivoli Storage Productivity Center GUI.

Common usability problem

By default, the configuration for the DS8000 CIM agent reports the DS8000 GUI using the https protocol. If you want to use the https version, do not manually add the http version for this CIM agent. Manually adding an http version for the CIM agent and then running a CIMOM discovery results in two different Element Managers being listed for one HMC.

Logs and traces

When you encounter a problem in this area, turn tracing on using the maximum debug level. For information about turning tracing on for the Data Server, see "Tracing the servers and agents" on page 11.

For information about turning tracing on for the Device Server, see "Tracing the Device server " on page 12. Run the following command:

```
srmcp -u <user_id> -p <password> log set
san.ElementManagerMgmtTrcLogger -filterkey INFO
```

The filer key can be: INFO, WARN, or ERROR.

Restart the Device Server and collect the Data Server and Device Server logs by running the **service.bat** or **service.sh** command.

The most useful log is:

```
<TPC_install_dir>\device\log\traceTPCDeviceServer.log
```

For element manager issues, look for entries that include the related Java classes:

```
com.ibm.tpc.dmc.***
```

JavaServer Pages (JSP)

To turn on the JSP debug trace, edit this file:

```
<TPC_install_dir>/device/apps/was/profiles/deviceServer/installedApps/
DefaultNode/DeviceServer.ear/DeviceServer.war/*.jsp
```

Change the following line:

```
var djConfig=isDebug: false
to
var djConfig=isDebug: true
```

No server reboot is required. You will then be able to view the JSP log entries in the GUI panel.

IBM Tivoli Storage Productivity Center universal agent

IBM Tivoli Storage Productivity Center provides a universal agent which collects data from IBM Tivoli Storage Productivity Center for use with IBM Tivoli Monitoring.

IBM Tivoli Monitoring provides enterprise-wide reporting on the health of the data center. Tivoli Monitoring is comprised of these components:

- Tivoli Enterprise Monitoring Server
- Tivoli Enterprise Portal Server
- Tivoli Enterprise Portal GUI client
- DB2 which stores the warehoused data for trending and historical analysis
- Monitoring agents which collect detailed information and metrics on the operating systems, networks, and applications

The IBM Tivoli Storage Productivity Center universal agent is found on the disk1 image of Tivoli Storage Productivity Center. The agent is in the `/tools` directory.

Note: Do not collect data more frequently than once every 10 minutes. This could have a detrimental performance impact on IBM Tivoli Storage Productivity Center.

Troubleshooting

The `TPCUA.log` is the main source of debug and trace information. Things to note about this log:

- This log contains anything that is displayed on the console for the user.
- Problems contacting a monitored IBM Tivoli Storage Productivity Center server is also logged in this file.
- The log is rolled over once it hits 200 KB and the last five copies will be kept.
- Logging is already set at the lowest level by default.
- INFO and ERROR entries are displayed both in the log and on the console.
- DEBUG entries are only entered in the log.

Collect this information for the IBM service representative:

```
<ITM_install_dir>\tmaitm6\scripts\tpcua\log\*.*
<ITM_install_dir>\tmaitm6\scripts\tpcua\config\*.*
<ITM_install_dir>\tmaitm6\metafiles\TPC_Network.mdl
```

Where `<ITM_install_dir>` is where IBM Tivoli Monitor is installed. By default, `<ITM_install_dir>` is:

```
For Windows:
C:\Program Files\IBM\ITM

For UNIX or Linux:
/opt/IBM/ITM
```

The IBM service representative may also need **repocopy** output.

McDATA Intrepid 10000

This section provides information about what to look for when diagnosing problems with the McDATA Intrepid 10000.

IBM Tivoli Storage Productivity Center supports the McDATA Intrepid 10000 Director with full support for topology discovery, zone control, and performance monitoring.

Things to note

Before TotalStorage Productivity Center 3.3, switch performance data collection was done using association calls to the CIMOM, which resulted in a huge number of jobs in the Device Server and a larger number of calls to the CIMOM in a relatively short period. In IBM Tivoli Storage Productivity Center, the enumeration approach is used when the monitored ports is over 40% of the total ports by default. This change applies to all switch performance monitors; not only for the McDATA Intrepid 10000.

Use best practices guidelines in handling multiple switch performance monitors:

- Have a few switch monitor jobs for all switches managed under one Enterprise Fabric Connectivity Manager (EFCM) to reduce the chance for a time out.
- Increase the collection interval time from the default of 15 minutes to 30 minutes, for example.
- Use a data collection frequency greater than the sample interval to reduce the amount of data stored in the IBM Tivoli Storage Productivity Center database.

Dependencies

The dependencies are:

- Requires McDATA 10000 EOSn 9.2.4.
- Hardware consideration for McDATA CIMOM scalability: the Java heap size needs to be increased to 1 GB when monitoring more than 300 ports.

Increasing the Java heap size for McDATA CIMOM

It is recommended that you increase the heap size of the JVM to 1 GB when managing medium or large fabrics (when the number of ports is greater than 300). By default the JVM is set to 512 MB. The maximum heap size can be set using the **-Xmx** JVM option.

To configure the JVM settings when running from the command line:

1. Stop the McDATA CIMOM.
2. Go to the following directory:
   ```
   wbemservices/cimom/bin
   ```
3. Edit the following file:

start_cimom.bat (for Windows)
start_cimom.sh (for UNIX)

4. Update the file and set the **-Xmx1024** JVM option.
5. Save the file.
6. Restart the McDATA CIMOM.

To configure the JVM settings when running as a service:

1. Open a command prompt or shell window.
2. Go to the following directory:

   wbemservices/cimom/bin

3. Edit the following file:

   start_cimom.bat (for Windows)
   start_cimom.sh (for UNIX)

4. Update the file and set the **-Xmx1024** JVM option.
5. Save the file.
6. Uninstall the server running as a service by using the following command:

   uninstall_service.bat (for Windows)
   ./uninstall_service.sh (for UNIX)

7. Reinstall the server using the following command:

   install_service.bat (for Windows)
   ./install_service.sh (for UNIX)

8. Restart the McDATA CIMOM.

Diagnosing problems

The McDATA CIMOM ships with a CIM browser (CIM Workshop) that can be used to check the CIMOM status quickly. The CIM Workshop can be started using the following:

C:\Program Files\SMI-S_Interface_Provider
wbemservices\bin\cimworkshop.bat

If you want to check which switches are configured to be managed by this CIMOM, check the **Mcdata_ManagedSwitch** instance in the /root/Mcdata namespace.

If you want to check which virtual switches for the McDATA 10000 will show up in IBM Tivoli Storage Productivity Center, check the **Mcdata_Switch** instance in the /root/Mcdata namespace.

To check the current version of the CIMOM, check the **Mcdata_SoftwareIdentity** instance in the /interop namespace.

Check these log files in the following directory:

<TPC_install_dir>/device/log/tracePerfMgr.log
<TPC_install_dir>/device/log/dmSvcTrace.log

To check to see if IBM Tivoli Storage Productivity Center is using the ″enumeration″ or ″association″ method, check the tracePerfMgr.log. An example of the output is shown below:

2007-04-0119:15:00,109-07:00ColThread8 PmFabricCache
isEnumerationPreferred
INFO Selected Enumeration; CIMOM https://xxx.xxx.xxx.xxx:5989;
TimeSlot 3; 267 / 298.0 = 0.89593295880014981; MinPortRatio 0.4,
MaxPortLimit 256

```
...
2007-04-0119:15:00,156-07:00ColThread8  CounterDataServiceSwitch
collectFCPortAndRateStats
INFO  Using Enumeration approach
...
```

tpctool

The **tpctool** is a standalone Java client and connects to the Device Server only. The **tpctool** connects through TCP, HTTP, and SOAP to the Web service APIs. The commands provide query, control and reporting capabilities only. The commands do not initiate discovery, probes, or configuration and control of agents.

The **tpctool** is installed in the following default directories:

```
C:\Program Files\IBM\TPC\cli (for Windows)
/<usr or opt>/IBM/TPC/cli (for UNIX or Linux)
```

The Windows command to run the tool is **tpctool**. The UNIX or Linux command to run the tool is **tpctool.sh**.

There are two distinct kinds of authentication:

- User authentication
- Super user or host-based authentication

The user authentication requires a user ID and password authenticated in the Device Server's authentication domain. Role-based authorization is enforced on a per command basis.

The super user or host authentication is for the user ID **tpc_superuser**. The password for this user ID bypasses role-based authorization. This is the required authentication method for AIX-based Device Servers.

Some of the control commands run a long time. An example is the **mkvol** command. There is no way to determine the intermediate status of the command. If the **tpctool** client stops (for example the user presses **Ctrl-C** or the node crashes), all connections with the job are lost. There is no way for reconnecting to the host for checking on the status.

The Device Server logs and trace files (assuming that tracing is on) are shown below. For the Disk commands, see this log:

```
DiskManagerService: dmSvcTrace.log
```

For the Fabric commands, see these logs for the Fabric Manager Service:

```
TPCZoneControl.log
msgTPCDeviceServer.log
traceTPCDeviceServer.log
```

For the reporting commands, see this log:

```
PerformanceService: tracePerfMgr.log
```

For fabric reports, the commands pass through the Fabric service for authorization. For subsystem reports, the commands pass through the Disk service for authorization.

For configuration commands, see this log:

```
ConfigService: traceTPCDeviceServer.log
```

For all other commands, see this log:

```
msgTPCDeviceServer.log
```

Some notes about passwords:

- You can use the GUI to change the Device Server host authentication password.
- If the GUI is used to change the password, the password will be updated in the database and the Device Server will be notified.
- The configuration file for all the Fabric agents need to be manually changed.

Common usability problems

Some of the common user problems are:

- Do not know where the installation directory is. Is the installation directory in a nonstandard location?
- Did not set the PATH or **chdir** to the `<TPC_install_dir>/cli` directory.
- Using **tpctool** on UNIX or Linux instead of **tpctool.sh**.
- Did not provide the following parameters in the command:
  - **-url**
  - **-user**
  - **-pwd**

This is a common error message for **tpctool**:

```
AAJ000009E Error communicating to the App server.
```

This same error message indicates an invalid port, invalid host, or the Device Server is unreachable. To distinguish what the problem is, you can set a debug flag which prints the Java stack if an exception occurs. Most errors are propagated to **tpctool** as exceptions. To set this flag:

- For Windows: **set TPCCLIDBG=1**.
- For UNIX or Linux: **export TPCCLIDBG=1**.

For an invalid host, the host name is embedded in the message. For example:

```
[SOAPException: faultCode=SOAP-ENV:Client: msg=Error opening socket:
java.net.UnknownHostException:badhost:...]
```

For a valid host but an invalid port or when the Device Server is down, an example of a message is as follows:

```
[SOAPException: faultCode=SOAP-ENV:Client: msg=Error opening socket:
java.net.ConnectException: Connection refused:...]
```

An invalid port is the most common problem.

If the Device Server node is unreachable, an example of the message is as follows:

```
[SOAPException: faultCode=SOAP-ENV:Client: msg=Error opening socket:
java.net.SocketException: Operation timed out:
connect:could be due to invalid address:...]
```

srmcp commands

Here are some simple **srmcp** commands for SANEventCorrelatorFactory for fabric:

- To list the configured filters:

```
cd <TPC_install_dir>\device\bin\w32-ix86
srmcp -u <user_ID> -p <password> SANEventCorrelatorFactory
  list
```

- To add a filter with the specified values (any or all can be specified). You can specify just the enterprise, the OIDs, or a combination.

```
cd <TPC_install_dir>\device\bin\w32-ix86
srmcp -u <user_ID> -p <password> SANEventCorrelatorFactory
  add filter [enterprise=<enterprise>][genericTrap=<number>]
  [specificTrap=<number>][<OID>=<value>]
```

- To specify an IP address that will be filtered for the specified filter:

```
cd <TPC_install_dir>\device\bin\w32-ix86
srmcp -u <user_ID> -p <password> SANEventCorrelatorFactory
  add address <ID> <IP_address>
```

- To remove a specified filter:

```
cd <TPC_install_dir>\device\bin\w32-ix86
srmcp -u <user_ID> -p <password> SANEventCorrelatorFactory
  remove filter <ID>
```

- To remove the IP address from the filter list:

```
cd <TPC_install_dir>\device\bin\w32-ix86
srmcp -u <user_ID> -p <password> SANEventCorrelatorFactory
  remove address <ID> <IP_address>
```

Here is the **srmcp** command to change the host authentication password:

```
cd <TPC_install_dir>\device\bin\w32-ix86
srmcp -u <user_ID> -p <password> ConfigService setAuthenticationPw
  <new_host_password>
```

Here is the **srmcp** command to change the DB2 password (not the db2admin password) that the server uses:

```
cd <TPC_install_dir>\device\bin\w32-ix86
srmcp -u <user_ID> -p <password> ConfigService setPw
  <new_host_password>
```

Fabric-specific problems

This section provides information for troubleshooting fabric-specific problems.

Fabric configuration (zone control)

Check these things for fabric configuration problems:

- Look at the GUI or CLI error code. Refer to *IBM Tivoli Storage Productivity Center Messages* for the error codes. The error codes have two parts: command failure code and native failure code.
- Check if the appropriate agent is available for zone control. The capabilities available depend on the agent.
- Check if the agent is connected to the SAN and if the fabric is configured appropriately for zoning.

The logs to check are:

- GUI job log.
- `AuditTrace.log` in the Data Server log directory.
- `TPCZoneControl.log` in the Device Server log directory.

- Check the traceTPCDeviceServer.log, traceTPCDeviceAgent.log and traceNative.log files for the Fabric agent. Check the Device Server log directory for native code and the out-of-band agent and the Fabric agent logs for the in-band agent.

How to adjust log and trace levels

You can adjust the following trace levels in the DataStore.properties and nativelog.properties files (both on the Device Server and Fabric agent).

Note: You should make a backup copy of the DataStore.properties file before modifying this file. A corrupt DataStore.properties file can cause Tivoli Storage Productivity Center to not start.

```
san.ControlTraceLogger.level=ALL
san.SanZoneControlTraceLogger.level=ALL
san.SanZoneControlAgentTraceLogger.level=ALL
native.trace.BrocadeZoneControl.level=DEBUG_MAX
native.trace.GS3ZoneControl.level=DEBUG_MAX
native.trace.brocadeScanner.level=DEBUG_MAX
native.trace.topologyScanner.level=DEBUG_MAX
```

Commands and tools for troubleshooting

Check the following:

- Look for the processes running (BrocadeApiScanner and BrocadeZoneControl).
- Look for the switch management tools:
  - Firmware
  - Topology
  - Current® zoning configuration
  - Transaction lock (cfgTransShow and cfgTransAbort).
- Look for IBM Tivoli Storage Productivity Center GUI or CLI instances that can be holding a transaction lock. It might be difficult to abandon the transaction lock on older switch models without rebooting the switch.
- Look at the IBM Tivoli Storage Productivity Center database for the agent and scanner status, and the token table.

Fabric agents

The logs to check for Fabric agent problems are as follows:

```
msgTPCDeviceServer.log
traceTPCDeviceServer.log
```

The logs are in this directory:

```
<TPC_install_dir>/device/log/
```

The Fabric agent is now supported on Solaris 10. The Fabric agent runs only in the global zone. The Fabric agent is **not** supported on Solaris 10 running in the local zone.

If you upgrade from Solaris 9 to Solaris 10, you must reinstall all agents.

Fabric discovery

The types of fabric discovery problems you can have are as follows:

- An error in the logs. Typically the first error you see is the most important.
- Information in the GUI that is incorrect or unexpected.

For information about error messages, see the *IBM Tivoli Storage Productivity Center Messages* publication and try to follow the explanations, user responses, and administrator responses.

For fabric discovery problems, follow these steps:

1. Run the **service** tool. See "Service tool: collecting information" on page 9.
2. Run **repocopy** to capture the IBM Tivoli Storage Productivity Center database. Follow these steps:
   a. Go to the following directory:
      ```
      C:\Program Files\IBM\TPC\data\server\tools (for Windows)
      /opt/IBM/TPC/data/server/tools (for UNIX or Linux)
      ```
   b. Run **repocopy.bat** (for Windows) or **repocopy.sh** (for UNIX or Linux).
   c. Select **Export data from repository tables**.
   d. Select the directory for the export file location.
   e. Use the default values for delimiter and quotes.
   f. Take the default values for the connection properties.
   g. Zip all the files in the output directory and send to IBM Support.
3. Get additional trace information.
   a. Go to the following directory:
      ```
      cd C:\Program Files\IBM\TPC\device\bin\w32-ix86
      ```
   b. Run **setenv**.
   c. Run the following command:
      ```
      srmcp -u <user_ID> -p <password> log set -filterkey INFO
      ```
      To turn logging off:
      ```
      srmcp -u <user_ID> -p <password> log set -defaults
      ```
4. Correlate the time stamps in the job log.
5. Check the health of the Device Server.
6. Check these files:
   ```
   C:\Program Files\IBM\TPC\device\log\traceTPCDeviceServer.log
   C:\Program Files\IBM\TPC\device\log\msgTPCDeviceServer.log
   ```
7. Correlate the time stamps in the job log.

To check on incorrect information, check these items:

- Did IBM Tivoli Storage Productivity Center receive an event? Look for the following:
  - Events in the alerts.
  - The `T_Alert_Log` table in the database MSG row.
  - ALR4100I: Received an SNMP trap notification from source.
  - ALR4101I Received an inband notification from source.
- The health of the probe:
  - Look in the log file:
    ```
    C:\Program Files\IBM\TPC\device\log\msg.probeFabricAgents.x.x.log
    ```

  - Look for exceptions or error messages.
- If the device reported incorrect information:
  - Check the BM files or scanner output from the service output.
  - Check to see if the XML file is readable with the text editor.
- If populating the results into DB2 failed:
  - Check the following log file:

    ```
    C:\Program Files\IBM\TPC\device\log\msg.probeFabricAgents.x.x.log
    ```
  - Look for exceptions or error messages.

Some known problems are:

- Some unsupported devices like multi-protocol routers could cause job failures during fabric discovery.
- Data migration might show duplicate peripheral entities as missing.
- AIX HBA information is limited; only provides name, manufacturer, firmware, and model.
- HBA hardware and BIOS versions are not shown in the IBM Tivoli Storage Productivity Center GUI.

Fabric Removed Resource Retention

The logs to check for fabric Removed Resource Retention problems are:

```
server_xxxxx.log
TPCD_xxxxx.log
```

These logs are in the following directory:

```
<TPC_install_dir>/data/log/
```

Zone control

For zone control, an agent capable of zoning must be installed. For non-Brocade environments, an in-band agent that had a probe run successfully on it is required. For Brocade switches, an out-of-band agent with the correct user ID and password and a successful probe is required.

For some non-standard switches, the zone and zone members are non-standard. This can prevent IBM Tivoli Storage Productivity Center from performing zoning operations.

The logs to check are as follows:

```
TPCZoneControl.log
msgTPCDeviceServer.log
traceTPCDeviceServer.log
```

These logs are in this directory:

```
<TPC_install_directory>\device\log\
```

Fabric alerts

The logs to check are as follows:

```
server_xxxx.log
TPCD_xxxx.log
```

These logs are in this directory:

```
<TPC_install_dir>/data/log/
```

Chapter 5. Troubleshooting IBM Tivoli Integrated Portal

IBM Tivoli Storage Productivity Center includes IBM Tivoli Integrated Portal, which contains an authentication service that interfaces with an LDAP-compliant repository and supports application single sign-on. This authentication service allows non-WebSphere-based applications (such as the DS8000 GUI) to participate in the IBM Tivoli Storage Productivity Center single sign-on environment using LTPA tokens. This section provides information on troubleshooting issues that arise in the authentication service.

Turn tracing on for IBM Tivoli Integrated Portal

There are two types of errors related to the authentication service:

Authentication client errors
These types of errors are logged by the application that is using the authentication client (such as the DS8000 GUI). These errors result when the application tries to communicate with the IBM Tivoli Integrated Portal authentication service to authenticate its users against IBM Tivoli Storage Productivity Center's LDAP-based repository. These types of errors include configuration, runtime, network, and authentication.

Authentication service errors
These types of errors are logged to IBM Tivoli Integrated Portal's SystemOut.log file.

To start tracing for IBM Tivoli Integrated Portal, follow these steps:

1. Open the IBM Tivoli Integrated Portal console.
2. Click **Troubleshooting → Logs and Trace**.
3. In the right pane, select a server (**server1**). Click on the server and click **Change Log Detail Levels**.
4. On the **Configuration** tab, set the **com.ibm.security.ess.*** log level to **All Messages and Traces**. Click **OK** and save the setting.
5. On the **Runtime** tab, set the **com.ibm.security.ess.*** log level to **All Messages and Traces**. Click **OK** and save the setting.
6. Restart the IBM Tivoli Integrated Portal server.

Turn tracing on for the token service infrastructure

The authentication service username provider and token service infrastructure logging is done through the **log4j** utility. The **log4j** utility is a Java-based logging utility. You should contact IBM Customer Support before running this utility to help you identify what package to debug.

To turn tracing on for the token service infrastructure, follow these steps:

1. Modify the log4j.properties file in the following directory:
   ```
   <TIP_home_directory>/profiles/TIPProfile/InstalledApps/TIPCell/
   authnsvc_ctges.ear/com.ibm.security.ess.war/WEB-INF/classes
   ```
2. In the log4j.properties file, change the log level for the desired package to debug.
3. Restart the IBM Tivoli Integrated Portal server.

Chapter 6. Troubleshooting DB2 and the database

IBM Tivoli Storage Productivity Center stores information in a DB2 database. This section provides information for troubleshooting the DB2 database and for daily operations for DB2.

DB2's health monitor

DB2 has a health monitor which is a server-side tool that adds a management-by-exception capability by constantly monitoring the health of an instance and active databases. The health monitor also has the capability to alert a database administrator (DBA) of potential system health issues. The health monitor proactively detects issues that might lead to hardware failures, or to unacceptable system performance or capability. The proactive nature of the health monitor enables users to address an issue before it becomes a problem that affects system performance.

The health monitor checks the state of your system using health indicators to determine if an alert should be issued. Preconfigured actions can be taken in response to alerts. The health monitor can also log alerts in the administration management-by-exception model to free up valuable DBA resources by generating alerts to potential system health issues without requiring active monitoring.

The health monitor gathers information about the health of the system using interfaces that do not impose a performance penalty. It does not turn on any snapshot monitor switches to collect information.

It is recommended that you monitor DB2's database performance for table spaces for IBM Tivoli Storage Productivity Center. You can do this through DB2's Health Center. A health indicator measures the health of some aspect of a particular class of database objects, such as table spaces. Criteria are applied to the measurement to determine healthiness. The criteria applied depends on the type of health indicator. A determination of unhealthiness based on the criteria generates an alert.

Threshold-based indicators are measurements that represent a statistic (on a continuous range of values) of the behavior of the object. Warning and alarm threshold values define boundaries or zones for normal, warning, and alarm ranges.

Note: An SMS table space is considered full if there is no more space on any of the file systems for which containers are defined, although it may still have some space left so that other data containers or their file systems.

To get to DB2's Health Center on Windows, go to **Start > All Programs > IBM DB2 > Monitoring Tools > Health Center**. To get to DB2's Health Center on UNIX, log on as the DB2 user like db2inst1 and run this command: **db2cc**.

On the Health Center window, click **Health Center > Configure > Health Indicator Settings**. On the Health Indicator Configuration Launchpad window, click **Instance Settings**. On the Instance Health Indicator Configuration window, enter **DB2** for instance. These defaults should be set for the Monitor Heap Utilization:

- Default is Yes

- Evaluate is Yes (indicates that evaluation is enabled)
- Warning is 85% (indicates that the warning threshold level is set)
- Alarm is 95% (indicates that the alarm threshold level is set)

On the Health Center window, click **Health Center > Configure > Health Indicator Settings**. On the Health Indicator Configuration Launchpad window, click **Global Settings** On the Global Health Indicator Configuration window, enter DB2 for the Instance and Database for Object Type. Check **Application Currency**, Deadlock Rate. The defaults should be set as follows:

- Default is Yes
- Evaluate is Yes (indicates that evaluation is enabled)
- Warning is 5 (indicates that the warning threshold level is set)
- Alarm is 10 Deadlocks per hour (you get warning if this threshold is reached).

Other parameters you might want to check for global health monitor are:

- Logging - Log filesystem Utilization
- Database - Automatic Storage Utilization

You can check the health of the database by going to the DB2 Health Center and follow the directions for the Health Center.

IBM Tivoli Storage Productivity Center uses the default configuration of IBM DB2's health monitor. This health monitor configuration should be sufficient in most cases. If you wish to get earlier warnings or include more parameters to monitor, you can do so.

For more information about DB2's Health Center, see *IBM DB2 Universal Database System Monitor Guide and Reference*.

Troubleshooting the database

For database problems, follow these steps:

1. Run the Service Tool. See "Service tool: collecting information" on page 9.
2. Run **repocopy** to capture the IBM Tivoli Storage Productivity Center database. For more information about **repocopy**, see "Repository Copy tool" on page 6. Follow these steps:
   a. Go to the following directory:
      ```
      C:\Program Files\IBM\TPC\data\server\tools (for Windows)
      /<usr or opt>/IBM/TPC/data/server/tools (for UNIX or Linux)
      ```
   b. Run **repocopy.bat** (for Windows) or **repocopy.sh** (for UNIX or Linux).
   c. Select **Export data from repository tables**.
   d. Select the directory for the export file location.
   e. Use the default values for delimiter and quotes.
   f. Take the default values for the connection properties.
   g. Zip all the files in the output directory and send to IBM Support.

Accessing the DB2 Control Center

To access the DB2 Control Center:

Windows

To start the DB2 Control Center on Windows, click the following:

```
Start > Programs > IBM DB2 > General Administration >
Tools > Control Center
```

Linux and AIX

To start the DB2 Control Center on Linux or AIX, follow these steps:

1. In a command prompt window, issue the following command:

   ```
   su - db2inst1
   ```

   Where **db2inst1** is the db2user.
2. Issue the following command to start the DB2 Control Center:

   ```
   db2cc
   ```

Accessing the DB2 command-line processor

To access the DB2 command-line processor:

Windows

To access the DB2 command-line processor on Windows, click the following:

```
Start > Programs > IBM DB2 > Command Line >
Tools > Command Line Processor
```

Linux and AIX

To access the DB2 command-line processor on Linux and AIX, follow these steps:

1. In a command prompt window, issue the following command:

   ```
   su - db2inst1
   ```

   Where **db2inst1** is the db2user.
2. Issue the following command to access the DB2 command-line processor:

   ```
   db2
   ```

   This opens the DB2 command-line processor window.
3. To exit the DB2 command-line processor window, issue the following command:

   ```
   quit
   ```

Starting DB2 when there is an exhausted buffer space error

When a memory leak causes the DB2 `MON_HEAP_SZ` parameter to run out of space, DB2 stops and does not start again until this parameter is increased in size. A temporary solution is to increase the `MON_HEAP_SZ` value.

Follow these steps to increase the `MON_HEAP_SZ` value:

1. Stop the Device Server and Data Server.
2. From the DB2 Control Center, click on the plus sign (+) until **DB2** appears under **Instances**.
3. Right-click **Configuration Parameters**. When a Database Manager configuration window appears, scroll down to `MON_HEAP_SZ`.
4. Change the setting for this parameter from 66 to 132.

Determining when a DB2 table is accessible

To determine whether the DB2 table is accessible, follow these steps:

1. Ensure that all tablespace states are 0. For UNIX or Linux, use the **db2 list tablespaces | grep State | sort | uniq** command.
2. Use the **db2 select * from <table_name>** command to determine whether you can read one record of the table. Fetch only the first row.

Looking up the table sample contents in the DB2 Control Center

To look up the table sample contents, follow these steps:

1. From the DB2 Control Center, click on the following:

   `System ► Instances ► DB2 ► Databases`
2. Select a database and navigate to **Tables**.
3. Select a table from the right pane. (The right pane contains all the tables under the database that you chose.)
4. Right-click the table that you selected, and click **Sample Contents** from the menu that appears.

Using DB2 commands

Here are some useful DB2 commands:

Determining the DB2 release and version

From the DB2 command-line processor, enter **quit** to exit the DB2 prompt, and then enter **db2level**.

Looking up the DB2 message code

From the DB2 command-line processor (CLI) DB2 prompt, enter **? ErrorCode**. For example:

```
db2=> ? sql0289
```

Connecting to a database

From the DB2 command-line processor, enter the following command:

```
db2=> connect to <database> user <user_ID> using <password>
```

You must supply the database name, user ID, and password.

Changing the password for a user ID

From the DB2 command-line processor, enter the following command:

```
db2=> connect to <database> user <user_ID> change password
```

Verifying the existence of a database

From the DB2 command-line processor, enter the following command:

```
db2=> list db directory
```

You can also enter:

```
db2=> list db directory show details
```

Looking at the database manager configuration file

From the DB2 command-line processor, enter the following command:

```
db2=> get db cfg
```

Looking at the database configuration file

From the DB2 command-line processor, enter the following command:

```
db2=> get db cfg for <database>
```

Looking at the DB2 registry

From the DB2 command-line processor, enter the following command:

```
db2=> quit
```

Then enter this command (you must be out of command-line mode to issue this command):

```
db2set -all
```

Listing active DB2 applications

Enter this command to list active DB2 applications:

```
db2 list applications show detail
```

Displaying the status change time

Enter this command to see the status change time:

```
db2 update monitor switches using uow on
```

Chapter 7. Troubleshooting performance and memory problems

This section provides information on how and when to tune IBM Tivoli Storage Productivity Center. This section also provides information on troubleshooting and common user errors.

This information applies to the Data Manager.

From a server perspective, you might not have to tune IBM Tivoli Storage Productivity Center if you meet these criteria:

- You are using less than 500 agents.
- You are using all the product defaults for schedules and history retention.

Common user errors

Here are some common user errors:

Why do certain devices, such as switches or subsystems, not show up in the selection list when defining a performance monitor?

The device will only show up in the selection list if it supports performance data collection. This situation can occur if the CIM Agent is not at the proper version, or an incorrect namespace was used for discovery. Also check to see if the CIM Agent discovery has completed successfully.

Why was I not successful in getting performance data?

For a successful performance data collection to occur, the device must have been discovered and probed successfully. For switches, an out-of-band fabric scan must also have been run.

Why does my performance monitor job fail immediately after it starts?

This usually occurs when there is a network connectivity problem with the CIM Agent.

Why do I get the error message ″performance data files could not be correlated″?

For SVC subsystems, this error message can occur during performance data collection. This is usually a problem with mismatching timestamps for the I/O statistics dump files on the SVC itself. This problem can be resolved by setting the timezone properly on the cluster.

Why do my CIM Agent operations time out?

You get the error message:

```
HWNPM4103E CIM/OM operation time out (30 seconds) expired.
```

Check whether the CIM Agent is up and running. You can do this by clicking: **Administrative Services ► Agents ► CIMOM**. Select **Test CIMOM Connection**.

Why does my CIM Agent run slowly?

Try stopping and restarting the IBM Tivoli Storage Productivity Center services. If this does not help, increase the value of com.ibm.tpc.perf.ConteConnectTimeOut in the following file:

```
<TPC_install_dir>\device\conf\pm.conf
```

Why does the performance correlation step take a long time to complete?
The performance correlation step can take a long time to complete if there are a lot of volumes. This causes a delay in collecting the first performance sample. To work around this problem, place the CIM Agent and IBM Tivoli Storage Productivity Center servers on faster machines. You can also use multiple CIM Agents to monitor a set of storage subsystems.

Why can I not define a performance monitor for a switch?
If you cannot define a performance monitor for a switch because it does not appear in the list of storage subsystems that can be monitored, do the following:

- Check to see that the SMI-S version supported by the switch vendor is 1.1 or higher.
- Check to see if the CIM Agent supports the Fabric sub-profile.
- Check to see if the CIM Agent discovery completed successfully.
- Check to see if an out-of-band fabric discovery was run after the CIM Agent discovery completed successfully. This is needed even if you are not using out-of-band Fabric agents if you want to use the switch performance monitor.

Disk performance reports

When trying to troubleshoot disk performance reports, ask these questions or perform these actions:

- Does the GUI hang?
- Are error messages viewed in the GUI?
- Is information in the GUI correctly populated?
- Is the Data Server running? Go to **Administrative Services ▸ Services ▸ Data Server ▸ <node>**.
- Is the Device Server running? Go to **Administrative Services ▸ Services ▸ Device Server ▸ <node>**.
- Check the Data Server and Device Server logs. Are there Java exceptions in the logs? Are the messages in the logs for information, warning, or error?

To view the monitor status, go to: **Disk Manager > Monitoring > Subsystem Performance Manager > <monitor> > <run>**. You will be able to access the logs from the monitor run status panel.

You can also run the GUI from a command prompt window:

```
<TPC_install_directory>\jre\bin\java.exe -Xmx512m
-classpath <TPC_install_directory>\com.tivoli.itsrm.gui.GuiMain
<host>:9549
```

You might see exceptions when you run this command.

Logs and traces

The logs to check are as follows:

```
<TPC_install_directory>\data\log\server_<xxx>.log
<TPC_install_directory>\device\log\tracePerfMgr<xxx>.log
```

There are no GUI logs, so you need to use the Service Tool to collect trace data. For information about the Service Tool, see “Service tool: collecting information” on page 9. You can gather trace data for review by running the Service Tool:

<TPC_install_directory>\service\service.<bat or sh>

Troubleshooting the performance data query API

Follow these steps for a general diagnosis procedure:

- Look in the tracePerfMgr.log for calls to PerfReportingService during the time the report was being generated. The parameters to the call can be seen in the log entry.
- Look for related calls to lower-level services. These show more detailed information. For example, you will be able to see the SQL queries that were made to the database.
- If the exception was reported by the GUI or other client, then it most likely originated from a lower-level exception that will be listed in the tracePerfMgr.log.

Server and agent memory settings

You might run into a situation where you are seeing a Java out–of–memory exception in either the server log or agent log. You should investigate the cause of this before increasing the memory allocation for either the server or the agent.

Note: If you need to substantially increase agent memory for one agent or several, then you will need to increase the memory on the server at **least** 25% greater than that of the largest agent memory setting. This is so the server will be able to handle the entire agent result set in memory.

The causes of an agent out–of–memory situation are almost always linked to a scan job. The possible reasons that a scan job can do this are as follows:

- Many directories are being scanned.
- Profiles associated with scans are configured to bring back large file lists (up to 32767 in any profile entrée box).
- You are using jobs to scan many local file systems or remote file systems.
- Many thousands of users own files on file systems being scanned.
- Several jobs are running simultaneously.
- Any combination of the above.

Keep in mind that storage resources like users, operating system groups, file systems, and directories are all buckets of information that the agent must create in memory. The more there are the more buckets the agent must create and maintain until the scan job is complete. Forethought and understanding in configuring the various jobs can usually avoid agent out–of–memory issues. However, there are situations where you must increase the agent's memory.

The causes for a server out–of–memory error are as follows:

- The history aggregation job in combination with the server receiving job results from agents.
- The agent result set for a scan is larger than the server available memory.

For information about how to increase the memory allocation, see *Installation and Configuration Guide*, Chapter 7. Administering IBM TotalStorage Productivity Center, "Increasing memory allocation".

Tuning switch performance for collecting data

When you are using the switch performance monitor and you run into timeout problems, you can change the attributes that affect the performance monitor. The performance monitor uses a couple of algorithms to collect performance information. The association algorithm is optimized for environments with larger numbers of switches managed by a CIM Agent and the enumeration algorithm is optimized for environments with fewer switches managed by a CIM Agent. TotalStorage Productivity Center attempts to determine which algorithm to use based on the fabric configuration.

You can change the following attributes in the `pm.conf` file:

MinSwitchPortRatio
: This attribute defines the minimum port ratio (for example, the ratio of monitored ports to total ports for a switch CIMOM). This value is used to determine which algorithm to use for querying the performance statistics data for the ports through the CIM Agent. If not set, the default cutoff value for the enumeration algorithm is 40%, so that if less than 40% of a CIMOM's ports are targeted for performance data collection, the association algorithm is used instead, provided that the **MaxSwitchPortLimit** is not exceeded (see **MaxSwitchPortLimit** below). This does not apply to Cisco CIMOMs.

MaxSwitchPortLimit
: This attribute defines the maximum port limit (for example, the maximum number of ports to be monitored using the association algorithm for switch CIMOMs). This value is used to determine which algorithm to use for querying the performance statistics data for the ports through the CIM Agent. If not set, the default cutoff value for the association algorithm is 256 ports, so that performance data will not be collected for more than 256 ports simultaneously using the association algorithm. Instead, the enumeration algorithm will be used for any CIMOMs after the limit has been reached.

To change these attributes in the `pm.conf` file, go to the following directory:

```
<TPC_install_dir>/device/conf/pm.conf
```

Remove the pound sign (#) from these attributes and modify your setting:

```
#com.ibm.tpc.perf.MinSwitchPortRatio = 0.4
#com.ibm.tpc.perf.MaxSwitchPortLimit = 256
```

Save the file. The Device server needs to be restarted for these changes to take effect.

Some commonly used parameters

AbbreviatedProbe parameter

When IBM Tivoli Storage Productivity Center probes a machine's hardware, it attempts to gather the same set of information for all disk devices. This generates nonfatal errors in the probe log for certain disk devices that do not support some of the probe's data gathering commands (like getting a list of disk defects). This can also cause errors to be generated in the machine's own logs and trigger alerts on other systems management consoles, which may be undesirable. To alleviate this problem, the **abbreviatedProbe** parameter may be set on the server. It is a global setting and will affect all agents.

To set this parameter, follow this procedure:

1. Stop the Data server.
2. Edit the TPCD.config file on the server, in <TPC_Data_Server_home>/config. In the server section, add the **abbreviatedProbe** parameter:

```
[server]
threadPoolSize=3
maxThreads=8
pingReceiveTimeout=10
abbreviatedProbe=1    //add this line and save the file

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

   Save the file.
3. Restart the Data server.
4. Run a probe on the Data agent or agents.

The only SCSI commands that will be sent to ANY disk drives ANYWHERE are:
- INQUIRY
- READ CAPACITY

This means that the following information will not be available for any SCSI or FC drives:
- medium-defect counts
- error counts
- I/O counts
- failure predicted
- cylinder count
- head count

saveNonRoot parameter

This server parameter affects the discovery of a NAS device's file systems. When the product discovers a NAS device with SNMP and the NAS device is then licensed, another discovery job must be run to discover the NAS device's exported file systems. The default behavior of this discovery is to discard export paths that are not at the root of the file system. For CIFS (Windows shares), the default behavior is to not have the administrative bit set. This is done so that a situation cannot occur where the product double counts information on a file system because it is treating two export jobs to the same file system as separate file systems. If for some reason the above discovery criteria cannot be met, then it can be overridden by setting the **saveNonRoot** parameter. This will then discover ALL exports (shares) of a given NAS device as file systems. Insert a parameter **SaveNonRoot=1** in the Server section of the TPCD.config file and restart the server to take effect.

Server parameters

server.config file

name="myserver"
This is the server's instance name. On installing the IBM Tivoli Storage Productivity Center server for the first time, the instance name is automatically set to the server's host name. If you move IBM Tivoli Storage Productivity Center

to a new host and copy the repository to a new DBMS using a backup/restore or database dump, this parameter will need to reflect the original instance name (which will always be the first server's host name). Otherwise, you will not see any of the old data on the new server. Once the new server is brought up, it will detect that its host name has changed and will broadcast this change out to the agent machines. The agents will then change this entry in their agent configuration files.

maxConnections=500
This is the maximum connections that the server will handle simultaneously. Once this number of connections is reached, the rejected agent connections will wait for 4 hours and retry. This number can be increased. On Solaris, this will require increasing the number of file descriptors per 32–bit process. By default, on current versions of Solaris, this number is 1024. The file descriptor is created per connection to IBM Tivoli Storage Productivity Center. Other operating systems may require tuning as well.

Note: GUI connections will override this setting.

routerThreads=1
All incoming connections are handled by the server's routing thread before they are rejected or handed to another service provider (agent, GUI, or scheduler). Connections to this process are not logged or limited by IBM Tivoli Storage Productivity Center. This can have the affect of allowing many more connections to the server than the **maxConnection** parameter would seem to allow in large environments. This is because the routing thread allows all connections to be queued before it hands the connection off to another service provider or tells the connecting agent that the maximum connection limit has been reached. Setting this parameter to its maximum setting of 3 threads will help by not allowing many *unclassified* connections to queue up.

scheduler.config file

maxSubmitThreads=3
This controls the number of threads that the scheduler service provider can spawn to submit new jobs to agents. The scheduler has a hard coded 10 minute window in which to submit the job to all agents required. If it goes over the 10 minute window, you will see ″wait time exceeded″ errors. This can be seen usually in large environments with agents that are on the other side of a slow WAN link from the server. Each submit thread will wait for a maximum of 3 minutes for an agent to respond to a request to send it a job. If each of the threads are tied up with slow agents or agents that do not respond, it is possible for the scheduler to run out of time to submit the job to all of the required agents. The first solution would be to break up the job so that it did not need to submit the

request to as many agents. The second solution would be to set this parameter to a higher setting, with a maximum setting of **8**.

nas.config file

The `nas.config` file is used by the NAS discovery process to discover possible NAS devices. It uses the NAS vendor's SNMP Enterprise designation as assigned in IANA:

`http://www.iana.org/assignments/enterprise-numbers`

You should not set any of the other parameters unless told to do so by support. Setting some of these could cause database deadlocks to occur.

You can edit the configuration files for the Data Manager to customize the operation of the server and agent components within your environment. These files reside in the `<TPC_install_directory>/data/config` directory, where `<TPC_install_directory>/data/config` represents the directory where you installed the product. When you make changes to the configuration file for the server, you must stop and restart the server before those changes will take effect.

Chapter 8. Troubleshooting the tape library

This section provides information on how to troubleshoot tape library problems.

Tivoli Storage Productivity Center supports the following tape libraries:

- Tivoli Storage Productivity Center 3494 Tape Library - The 3494 tape library can be used for data consolidation to help achieve higher performance and reduced requirements for tape drives and cartridges, environmental controls and personnel. The 3494 tape library supports WORM and standard rewritable media, providing further opportunity for consolidation.
- IBM System Storage TS3500 Tape Library (formerly Tivoli Storage Productivity Center 3584 Tape Library) - The TS3500 Tape Library is designed to provide a highly scalable, automated tape library for mainframe and open systems backup and archive in midrange to enterprise environments.
- IBM System Storage TS3310 Tape Library is a modular, scalable tape library designed to grow as your needs grow.

Tivoli Storage Productivity Center supports IBM and non-IBM tape libraries through the SMI-S (CIM) agent. This 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 really logical partitions of the given library, and not physical accessors.
- Depending on the size of the tape library and network latency between the SMI-S agent host and the tape library, probes of the tape library might take a long time and might fail because of time-outs.
- The SMI-S agent does not display an exception to the SMI-S client such as Tivoli Storage Productivity Center when communication between the SMI-S agent and the tape library fails. Instead, empty result sets are returned. This limits Tivoli Storage Productivity Center's capabilities in detecting connection failures relating to tape libraries.
- In this release, support for the IBM 3494 Tape Libraries is limited to discovery and in-place launching of the ETL Specialist. This assumes that the SMI-S agent has been configured accordingly.
- 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 because the increased load on the Agent would increase the likelihood of time outs. 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 of the tape cartridges are returned to the CIM client, in this case, Tivoli Storage Productivity Center. To resolve this situation, restart the IBM SMI-S Agent for Tape; refer to the documentation for the SMI-S Agent for instructions on how to do this.

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

Chapter 9. Troubleshooting Network problems

This section provides information on how to troubleshoot network problems.

There are some commands that are helpful in identifying network connectivity and problems:

- **netstat**
- **nslookup**
- **ping**
- **telnet**
- **tracert**

These commands are described in this section.

netstat command

The **netstat** command is used to query TCP/IP about the network status of the local host. The syntax of this command is implementation–dependent. This is a useful command for debugging purposes.

This command provides information on:

- Active TCP connections on the local host.
- State of all TCP/IP servers on the local host and the sockets used by them.
- Devices and links used by TCP/IP.
- The IP routing tables (gateway tables) in use by this local host.

Syntax

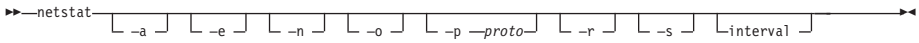

Parameters

–a Displays all connections and listening ports.

–e Displays Ethernet statistics. This may be combined with the **–s** option.

–n Displays addresses and port numbers in numerical form.

–o Displays the owning process ID associated with each connection.

–p ***proto***
Shows connections for the protocol specified by ***proto***. ***Proto*** can be: TCP, UDP, TCPv6, or UDPv6. If this option is used with the **–s** option to display per–protocol statistics, ***proto*** may be IP, IPv6, ICMP, ICMPv6, TCP, TCPv6, UDP, or UDPv6.

–r Displays the routing table.

–s Displays per–protocol statistics. By default, the statistics are shown for IP, IPv6, ICMP, ICMPv6, TCP, TCPv6, UDP, and UDPv6. The **–p** option may be used to specify a subset of the default.

interval
Displays selected statistics, pausing interval seconds between each

display. Click **CTRL+C** to stop redisplaying statistics. If omitted, **netstat** will print the current configuration information once.

nslookup command

The **nslookup** command is a common DNS application utility for querying name servers. **Nslookup** allows you to locate information about network nodes, examine the contents of a name server database and establish the accessibility of name servers. This command can help identify network connectivity problems due to misidentification of the host name (fully–qualified host name or short name) or the IP address (or both).

Syntax

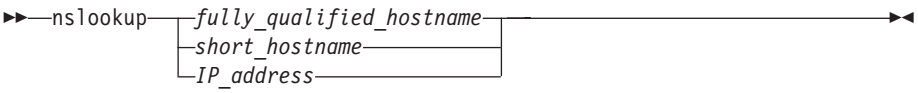

Parameters

fully_qualified_hostname, short_hostname, IP_address
Specify a fully–qualified host name, short host name, or IP address.

ping command

The **ping** command sends ICMP Echo Request (ECHO_REQUEST) packets to the host once every second. Each packet that is echoed back through an ICMP Echo Response packet is written to the standard output, including round–trip time. This command is intended for use in testing, managing, and measuring network performance. It should be used primarily to isolate network failures.

Note: Most UNIX systems provide a similar utility called **traceroute**.

Syntax

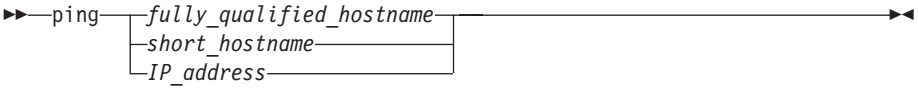

Parameters

fully_qualified_hostname, short_hostname, IP_address
Specifies the target host.

telnet command

IBM Tivoli Storage Productivity Center uses **telnet** to contact a specific port to perform its scheduling activities and to establish communication with the Server and Agent. If **telnet** is invoked without arguments, it enters command mode as indicated by its prompt (**telnet>**). In this mode, **telnet** accepts and executes the commands. If **telnet** is invoked with arguments, it performs an open command with those arguments.

In the event of an error, or if the **telnet** connection is closed by the remote host, **telnet** returns a value of 1. Otherwise, **telnet** returns a value of zero (0).

Syntax

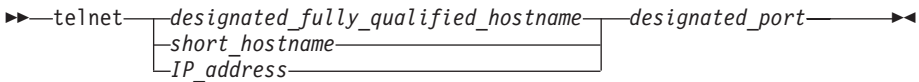

Parameters

designated_fully_qualified_hostname, short_hostname, IP_address
Specifies the target host.

designated_port
Specifies the port.

Error messages

Some of the following error messages may be displayed by **telnet**:

Error! Could not retrieve authentication type.
The type of authentication mechanism is obtained from a system file which is updated by **inetsvcs_sec**. If the system file on either the local host or the remote host does not contain known authentication types, this error message is displayed.

telnet/tcp: Unknown service
telnet was unable to find the TELNET service entry in the services (4) database.

hostname: Unknown host
telnet was unable to map the host name to an Internet address. Your next step should be to contact the system administrator to check whether there is an entry for the remote host in the hosts database.

?Invalid command
An invalid command was typed in **telnet** command mode.

tracert command (Windows only)

The **tracert** command displays the network route to a specified host and identify faulty gateways. A packet might be redirected to large Internet Centers and then rerouted to the smaller pipelines, before eventually finding its way to the IP address specified.

This command can provide valuable information about each location the packet is sent to before it is received.

Syntax

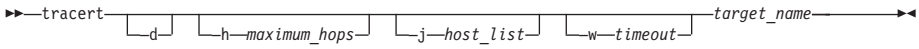

Parameters

–d Do not resolve the addresses to host names.

–h ***maximum_hops***
Specifies the maximum number of hops to search for a target.

–j ***host_list***
Specifies the source route along a host list.

–w ***timeout***
Specifies the timeout in milliseconds for each reply.

target_name
Specifies the address or target that you want to trace.

Make sure that you are connected to a network. Your prompt should be set to:
`C:\windows\>`

Chapter 10. Troubleshooting SAN FS problems

This section provides information on how to troubleshoot SAN FS problems.

Problem descriptions

Note: This section applies only to IBM Tivoli Storage Productivity Center for Data.

Problem: loss of connectivity.

There is a loss of connectivity to the SAN.FS LUNs.

Solution

Check the physical connectivity to the SAN.FS LUNs. LUNs need to be zoned and assigned to the agent machine. Depending on the fibre channel HBA used, a program can be used to verify that the SAN.FS LUNs are visible to the agent machine. When using a QLogic HBA, the SANBlade Control FX program can be used to verify that the LUNs are connected to the agent machine.

Problem: metadata is not available.

The metadata is not available to the agent machine.

Solution

Log onto the SAN.FS console and ensure that the agent machine is defined and has been granted root privileges. Also, ensure that all SAN.FS LUNs are available.

Problem: unable to read or write to SAN FS.

You are unable to read or write to SAN.FS from the agent machine.

Solution

Ensure that the SAN.FS client code has been activated and check that the LUNs have been labeled correctly. For example, run the following command on a Sun Solaris system:

```
dd if=/dev/dsk/c3t6d12s2 count=1 skip=8
```

If the **SDISK...** text is displayed, the LUN is correctly labeled. Otherwise, the **labellun** command should be run against the raw disk representation of the LUN. For example:

```
/usr/tank/client/bin/labellun -d /dev/rdsk/c2t6d12s12
```

The SAN.FS disks should be specifically specified as something like **c3t6dXXs2** where XX is the number of the SAN.FS LUN. When setting up the SAN.FS client, the device should be specified something like this:

```
device=/dev/dsk/c3t6d*s2
```

Check to see that the SAN.FS LUNs can still be listed by running the following command:

```
/usr/tank/client/bin/sanfs_ctl disk list -instance
<sanfs_mount_point>
```

Where the `<sanfs_mount_point>` is something like this: `/mnt/sanfs`. If that fails, then run the following command:

```
sanfs_ctl disk add -instance <sanfs_mount_point> -disk
<disk_device_path, ...>
```

The disk device path should be in this format:

```
</dev/dsk/c3t6d0s2,/dev/dsk/c3t6d1s2,...,/dev/dsk/c3t6d12s2>
```

The exact disk names will be different.

Problem: unable to report capacity values.

Unable to report capacity values of the SAN.FS.

Solution

This is working as designed. The values will show up as N/A after a successful probe of a SAN.FS. The reason for this is because the native functions used to collect the capacity values return incorrect hard–coded values. If a future release is planned, work may be done to report those values correctly.

Chapter 11. Configuring and troubleshooting the Novell NetWare servers and agents

This section provides hints and tips and troubleshooting information for Novell NetWare server and agents.

Hints and tips for NAS

1. You should understand the flow of the IBM Tivoli Storage Productivity Center. For example, a scan on a new volume cannot be performed before a probe has been run to discover the new volume. For NAS, never do a POC for NAS or Netware if you have not configured this elsewhere before. Typically a lack of preparation on the System Engineer's part becomes an emergency to Support in this case. You cannot install a NAS (or Netware) without prior experience or guidance.
2. For a CIFS NAS discovery; it will not discover NAS devices in a domain across a "trusted" relationship. You must have an agent machine in the same domain as the NAS device. This goes for AD transitive trusts as well.
3. For NAS and Netware scanning, make sure that SNMP ports are opened between the proxy scanning agent and the NAS/Netware machines.
4. You do not install agents to the NAS filers themselves–rather, you install them to Windows or UNIX machines (or both) that have access to those NAS filers. For example:
   - Install agents on Windows machines within the same domain as NAS filers.
   - Install agents on UNIX machines that have imports for the file systems within the NAS filers.

Configuring and discovering NetWare servers

To discover the NetWare servers and volumes within your environment, follow this procedure:

1. Access the NetWare Tree Logins window by clicking **Administrative Services ▸ Configuration ▸ NetWare Tree Logins** of the function tree.
2. Enter a valid, fully–qualified user login ID and password for each listed NDS trees found by the agent. The user logins must have permission to enumerate the servers and volumes in every NDS tree against which you want to run a discovery. See the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for more information.
3. Once you have configured IBM Tivoli Storage Productivity Center to access your NDS trees, you can run a discovery job on all the agents for which you have specified a valid ID and password on the NetWare Tree Logins window. You might still need to configure agents to discover, probe, and scan the NDS boxes.

To run the NetWare discovery job, right–click the **Administrative Services ▸ Monitoring –> Discovery ▸ NetWare Filer** in the function tree. Select **Run Now** from the pop–up menu. The discovery job will discover all the NetWare servers in the NDS trees, and log in to those servers to gather volume and disk information. The user ID and password supplied on the NetWare Tree logins window is used to log in.

When you set up the NDS logins, ensure that all fields listed below (from the NetWare Tree Login window and the Tree Login Editor window) have values:

Tree Name
The name of the NDS tree IBM Tivoli Storage Productivity Center discovered during installation.

Discovery Agent
The name of the agent that will perform the discovery. The Agent will be the machine name that has an Data agent installed and can access the NDS tree.

Loginid
The login ID the Agent will use when logging into the NDS tree.

Agent name, ID and password
Verify the Agent name, user ID, and password by clicking on the **Edit** button to bring up the Tree Login Editor.

Troubleshooting NAS and Novell NetWare problems

Problem
Why does the NetApp Filer not show up when you have a Windows domain? The server, agent and NetApp filer are not in a domain or in the same domain. For IBM Tivoli Storage Productivity Center to see a NetApp Filer, the server, agent (used to access the filer) and NetApp filer must all be in the same domain.

Solution
Create a domain and add the IBM Tivoli Storage Productivity Center server, agent and NetApp Filer. You should be able to see the NetApp Filer.

Problem
How do I obtain normal quotas from a Qtree using NetApp Filers?

Solution

To obtain normal quotas from a Qtree, follow this procedure:

1. Configure and run a NetApp quota job. This will include a setting for triggering an alert based on the percentage of quota utilized. The first time this is run, it will not trigger any alerts; it will only populate the `tivoli.default_qtree` directory group.
2. Configure and run a scan on the Filer which includes applying the `tivoli.default_qtree` group to the scan. This will gather the statistics for each qtree (which is really just a directory).
3. Rerun the configured NetApp quota job. This will trigger the alerts.

Problem
Why am I unable to obtain NetApp Filer user and group quotas under a Qtree?

Solution
The product will not gather user and group quotas when they exist under a Qtree. This is a limitation of the product.

Appendix. Accessibility features for IBM 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:

- IBM Tivoli Storage Productivity Center functions are available using the keyboard for navigation instead of the mouse. You can use keys or key combinations to perform operations that can also be done using a mouse. However, you must use the mouse to navigate the Topology Viewer and report graphs. Standard operating system keystrokes are used for standard operating system operations.
- You can use screen readers to read the user interface.
- The user interface communicates all information independently of color.
- The *IBM Tivoli Storage Productivity Center Information Center*, and its related publications are accessibility-enabled and include the following accessibility features:
  - The information center is provided in XHTML 1.0 format, which is viewable in most Web browsers. XHTML allows you to view documentation according to the display preferences set in your browser. It also allows you to use screen readers and other assistive technologies.
  - All documentation is available in PDF format.
  - All images are provided with alternative text, so that users with vision impairments can understand the contents of the images.

The accessibility features of the information center are described at http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=/com.ibm.sspc_v132.doc/fqz0_c_accessibility_overview.html.

Keyboard navigation

This product uses standard Microsoft Windows navigation keys.

Interface information

Use the options available in the **Preferences** > **Look and Feel** menu to select how to display the IBM Tivoli Storage Productivity Center user interface. To do this, complete the following steps:

1. Start the IBM Tivoli Storage Productivity Center user interface.
2. Select one of the following options from the **Preferences** > **Look and Feel** menu to change the visual appearance of the user interface to best suit your visual needs:
   - Windows Classic
   - Windows
   - CDE/Motif

- Metal

Related accessibility information

You can view the publications for IBM Tivoli Storage Productivity Center in Adobe Portable Document Format (PDF) using the Adobe Acrobat Reader. You can access the PDFs from the Printable PDFs topic in the information center at http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp.

IBM and accessibility

See the IBM Human Ability and Accessibility Center website for more information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Software Interoperability Coordinator, Department YBWA
3605 Highway 52 N
11400 Burnet Road
Rochester, MN 55901
U.S.A

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml.

Adobe and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Red Hat and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Glossary

This glossary includes terms and definitions for IBM Tivoli Storage Productivity Center.

The following cross-references are used in this glossary:

- *See* refers the reader from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- *See also* refers the reader to a related or contrasting term.

To view glossaries for other IBM products, go to: http://www.ibm.com/software/globalization/terminology/.

agent An entity that represents one or more managed objects by sending notifications regarding the objects and by handling requests from servers for management operations to modify or query the objects.

Basic HyperSwap™
In System z, a replication feature that performs the following actions:

- Monitoring for events that indicate a storage device has failed
- Determining whether the failing storage device is part of a Peer-to-Peer Remote Copy (PPRC) pair
- Determining from policy, the action to be taken
- Ensuring that data consistency is not violated
- Swapping the I/O between the primary logical devices in the consistency group with the secondary logical devices in the consistency group.
- Allowing only CKD volumes to be added to the HyperSwap session.

CIM See Common Information Model.

CIM agent
The code that consists of common building blocks that can be used instead of proprietary software or device-specific programming interfaces to manage CIM-compliant devices.

CIM object manager (CIMOM)
The common conceptual framework for data management that receives, validates, and authenticates the CIM requests from the client application. It then directs the requests to the appropriate component or service provider.

CIMOM
See CIM object manager

CKD Count key data

cluster
1. In SAN Volume Controller, a pair of nodes that provides a single configuration and service interface. 2. In IBM System Storage DS8000, a partition capable of performing all DS8000 series functions. With two clusters in the DS8000 storage unit, any operational cluster can take over the processing of a failing cluster.

Common Information Model (CIM)
An implementation-neutral, object-oriented schema for describing network management information. The Distributed Management Task Force (DMTF) develops and maintains CIM specifications.

comma-separated value file
A text file, created in a spreadsheet program such as Microsoft Excel. A CSV file includes each of the copy sets you want to add to the session separated by a comma.

community name
The part of an SNMP message that represents a password-like name and that is used to authenticate the SNMP message.

consistency group
The set of target volumes in a session that have been updated to preserve write ordering and are therefore, recoverable.

copy set
A set of volumes that contain copies of the same data. All the volumes in a copy set are the same format (count key data [CKD] or fixed block) and size.

CSV See comma-separated value file

data collection
See *discovery*.

data exposure
The time between the point at which the data is written to primary storage, and when it is replicated to secondary storage. Data exposure includes factors such as:
- Requested consistency-group interval time
- Type of storage systems
- Physical distance between the storage systems
- Available bandwidth of the data link
- I/O load on the storage systems

discovery
The process of finding resources within an enterprise, including finding the new location of monitored resources that were moved. Discovery includes the detection of changes in network topology, such as new and deleted nodes or new and deleted interfaces. See also *discovery interval*.

discovery interval
The frequency at which topology and attribute information is gathered. The discovery interval is set by a schedule to occur either periodically or at specific times. Discovery can also occur at other times, such as when triggered by an event from a SAN switch.

discovery job
A job that enables you to find new Windows machines that have been introduced into your environment, identify the servers and volumes within NetWare trees (NDS trees), discover the file systems within NAS filers, and discover the CIMOMs in your environment and the storage subsystems managed by those CIM/OMs.

enterprise repository
A component of the Data server that records and stores all information about the monitored computers' storage assets and their usage over time. The repository is organized into relational database tables and is accessed by Data server using Java Database Connectivity (JDBC).

event Any significant change in the state of a system resource, network resource, or network application. An event can be generated for a problem, for the resolution of a problem, or for the successful completion of a task. Examples of events are: the normal starting and stopping of a process, the abnormal termination of a process, or the malfunctioning of a server.

fabric A complex network using hubs, switches, and gateways. Fibre channel uses a fabric to connect devices.

failover and failback
The implementation of a complex local or remote disaster-recovery solution with the capability of a two-way site switch.

fibre channel
A technology for transmitting data between computer devices. It is especially suited for attaching computer servers to shared storage devices and for interconnecting storage controllers and drives.

FlashCopy
An optional feature of the DS8000 series that can make an instant copy of data; that is, a point-in-time copy of a volume.

global copy
An optional capability of the DS8000 remote mirror and copy feature that maintains a fuzzy copy of a logical volume on the same DS8000 storage unit or on another DS8000 storage unit. In other words, all modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume at a later point in time. The original order of update is not strictly maintained. See also *remote mirror* and *copy* and *metro mirror*.

global mirror
An optional capability of the remote mirror and copy feature that provides a 2-site extended distance remote copy. Data that is written by the host to the storage unit at the local site is automatically maintained at the remote site. See also Metro Mirror and Remote Mirror and Copy.

globally unique identifier (GUID)
An algorithmically determined number that uniquely identifies an entity within a system.

heat map
A color-coded data chart where colors are used to differentiate values in a data set.

host A computer that is connected to a network (such as the Internet or a SAN) and provides a point of access to that network. Also, depending on the environment, the host can provide centralized control of the network. The host can be a client, a server, both a client and a server, a manager, or a managed host.

host volume
A volume that represents the volume functional role from an application point of view. The host volume can be connected to a host or server, and receives read, write, and update application I/Os, depending on the site that the application is writing to.

in-band discovery
The process of discovering information about the SAN, including topology and attribute data, through the fibre-channel data paths. Contrast with *out-of-band discovery*.

intermediate volume
The target of the remote copy relationship, and the source of a FlashCopy relationship in which the target of the FlashCopy is the H2 volume.

job scheduler
A component of the Data server that deploys all monitoring activities. The job scheduler controls when monitoring jobs are run by agents.

journal volume
A volume that functions like a journal and holds the required data to reconstruct consistent data at the Global Mirror remote site. When a session must be recovered at the remote site, the journal volume is used to restore data to the last consistency point.

logical unit number (LUN)
An identifier used on a SCSI bus to distinguish among devices (logical units) with the same SCSI ID. For a SCSI bus, a LUN represents a storage volume.

LUN See *logical unit number*.

managed disk (MDisk)
A SCSI logical unit that a Redundant Array of Independent Disks (RAID) controller provides and a cluster manages. The MDisk is not visible to host systems on the SAN.

managed host
A host that is managed by Tivoli Storage Productivity Center and one or more active in-band fabric agents. Install in-band fabric agents on host systems with host bus adapters (HBAs) that are connected to the SAN fabrics that you want to manage.

Management Servers
Increased availability of the replication management software with the implementation of a high-availability configuration such that one management workstation runs as standby, ready to take over in case of a failure of the active workstation.

Note: The takeover is not automatic and requires you to issue a takeover command.

metro mirror
A function of a storage server that maintains a consistent copy of a logical volume on the same storage server or on another storage server. All modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume. See also *Remote Mirror* and *Copy* and *Global Copy*.

Metro Global Mirror
The three-site remote mirroring solution.

out-of-band discovery
The process of discovering SAN information, including topology and device data, without using the fibre-channel data paths. A common mechanism for out-of-band discovery is the use of SNMP MIB queries, which are invoked over a TCP/IP network. Contrast with *in-band discovery*.

ping job
A job that tracks the availability of assets and that is performed by an agent. Several ping jobs can be used to monitor the availability of any computer or subset of computers in the network.

pool A named set of storage volumes that is the destination for storing client data.

primordial pool
Unallocated storage capacity on a storage device. Storage capacity can be allocated from primordial pools to create storage pools.

probe job
A job that itemizes and creates an inventory of assets, such as computers, controllers, disk drives, file systems, and logical units, and that is performed by an agent. Several probe jobs can be used on any computer or subset of computer

RAID See *Redundant Array of Independent Disks*.

Recovery point objective (RPO)
The maximum amount of data that you can tolerate losing in the case of a disaster.

remote console
A console that is installed on a machine other than the one on which the server is installed. A remote console lets you access Tivoli Storage Productivity Center from any location.

remote mirror and copy
A feature of a storage server that constantly updates a secondary copy of a logical volume to match changes made to a primary logical volume. The primary and secondary volumes can be on the same storage server or on separate storage servers.

role A function that a volume assumes is the copy set. The role is is composed of the intended use and, for Global Mirror and Metro Mirror, the volume's site location. Every volume in a copy set is assigned a role. A role can assume the functions of a host volume, journal volume, or target volume. For example, a host volume at the primary site might have the role of `Host1`, while a journal volume at the secondary site has the role of `Journal2`.

role pair
The association of two roles in a session that take part in a copy relationship. For example, in a metro mirror session, the role pair could be the association between the volume roles of `Host1` and `Host2`. In another example, a `Host1` volume could be a host volume on the primary site, and a `Host2` volume could be a host volume on the secondary site.

SAN See *storage area network*.

scan job
A job that monitors the usage and consumption of your storage and the constraints and that is performed by an agent. Several scan jobs can be used to monitor the file systems on any computer or subset of computers.

SCSI See *Small Computer Systems Interface*.

session
A collection of multiple copy sets that comprise a consistency group.

site switching
See also *failover and failback*.

SMI-S See *Storage Management Initiative - Specification*.

SMI-S agent
See *CIM Object Manager (CIMOM)*. See also *Storage Management Initiative - Specification (SMI-S)*.

SNIA See *Storage Networking Industry Association*.

source The site where production applications run while in normal operation. The meaning is extended to the disk subsystem that holds the data as well as to its components: volumes and LSS.

storage area network
A dedicated storage network tailored to a specific environment, combining servers, storage products, networking products, software, and services.

storage group
A collection of storage units that jointly contain all the data for a specified set of storage units, such as volumes. The storage units in a group must be from storage devices of the same type.

Storage Management Initiative - Specification (SMI-S)
The standard that defines the protocol used for communication with SMI-S agents.

Storage Networking Industry Association (SNIA)
An alliance of computer vendors and

universities that focus on developing and promoting industry standards for storage networks.

storage pool
An aggregation of storage resources on a SAN that have been set aside for a particular purpose.

System z® Global Mirror
See also *Global Mirror*.

target The site to where the data is replicated, the copy of the application data. The meaning is extended to the disk subsystem that holds the data as well as to its components: volumes and logical subsystem (LSS).

target volume
A volume that receives data from a host volume or another intermediate volume. It is used only in FlashCopy sessions.

topology
The physical and logical arrangement of devices in a SAN. Topology can be displayed graphically, showing devices and their interconnections.

VDisk See *virtual disk*.

virtual disk (VDisk)
A device that host systems attached to the storage area network (SAN) recognize as a Small Computer System Interface (SCSI) disk.

virtualization
A concept in which a pool of storage is created that contains several disk subsystems. The subsystems can be from various vendors. The pool can be split into virtual disks that are visible to the host systems that use them.

virtual storage area network (VSAN)
A Cisco technology that allows independent logical fabrics to be defined from a set of one or more physical switches. A given switch port is assigned to only one VSAN. Each VSAN is completely isolated from the other VSANs and functions as a separate and independent fabric with its own set of fabric services (for example, Name Services, zoning, routing, and so on).

volume
The basic entity of data storage as defined by the SCSI protocol. A volume is a logical address space, having its data content stored on the systems disk drives.

VSAN See *virtual storage area network*.

zone A segment of a SAN fabric composed of selected storage devices nodes and server nodes. Only the members of a zone have access to one another.

zone set
A group of zones that function together on the fabric. Each zone set can accommodate up to 256 zones. All devices in a zone see only devices assigned to that zone, but any device in that zone can be a member of other zones in the zone set.

Index

X

Program Number: 5608-WB1, 5608-WC0, 5608-WC3, 5608-WC4

Printed in USA

GC27-2342-00