

Tivoli Storage Productivity Center
Version 4.2.1

Problem Determination Guide

IBM

Tivoli Storage Productivity Center
Version 4.2.1

Problem Determination Guide

IBM

Note:

Before using this information and the product it supports, read the information in "Notices" on page 95.

| This edition applies to version 4, release 2, modification 1 of IBM Tivoli Storage Productivity Center (product
| numbers 5608-WB1, 5608-WC0, 5608-WC3, 5608-WC4, 5608-E14, 5608-WB2, and 5608-WB3) and to all subsequent
| releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 2001, 2010.

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

Contents

Preface	v	SAN Planner	39
Who should read this guide	v	Configuration History and Analysis	45
Publications	v	Data Path Explorer	47
Contacting IBM Support Center	v	SMI-S fabric probe	48
Reporting a problem	vi	NAS and NetApp	50
Conventions used in this guide	vi	VMware ESX	51
		Reporting groups	53
		Monitored Computer Storage Space reports	53
Chapter 1. Introduction to problem determination	1	FlashCopy	55
Tivoli Storage Productivity Center problem determination resources	3	Element Management	56
		McData Intrepid 10000	57
		tpctool	59
		Fabric-specific problems	61
		CIM agents	67
Chapter 2. Tivoli Storage Productivity Center tools	5	Chapter 4. Troubleshooting Tivoli Integrated Portal	69
Repository Copy tool	5		
Exporting repository data	5	Chapter 5. Troubleshooting DB2 and the database	71
Importing repository data	6		
Service tool: collecting information	7	Chapter 6. Troubleshooting performance and memory problems	75
Running the Service tool for the servers	8		
Running the Service tool for the agents	9	Chapter 7. Troubleshooting the tape library	83
Tracing the servers and agents	10		
		Chapter 8. Network connectivity	85
Chapter 3. Troubleshooting IBM Tivoli Storage Productivity Center	19	Appendix A. Accessibility features for Tivoli Storage Productivity Center	89
Configuration files	19		
server.config file	19	Appendix B. Accessibility features for IBM Tivoli Storage Productivity Center for Replication	91
scheduler.config file	21		
TPCD.config file	21	Notices	95
agent.config file	22	Trademarks	97
Log files	22	Glossary	99
Installation log files	22		
Default log file locations	23	Index	107
Audit logs	24		
Diagnosing IBM Tivoli Storage Productivity Center problems	25		
Discovery	26		
Configure Devices wizard	28		
Monitoring service	29		
Database monitoring	30		
Performance monitoring	31		
File system scans	32		
Topology viewer	34		
Single sign-on	35		
Storage Optimizer	36		
Rollup reports	38		

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

Publications

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

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

Contacting IBM Support Center

This topic provides information about 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 website at http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition.

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

- Customers in the United States can call 1-800-IBM-SERV (1-800-426-7378).
- For international customers, go to the Tivoli Storage Productivity Center technical support website for customer support telephone numbers.

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

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

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

Reporting a problem

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

Have the following information ready when you report a problem:

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

Conventions used in this guide

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

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

The following typeface conventions are used in this publication:

Bold

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

Italic

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

monospace

- Commands and command options in examples

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

For syntax notation, these conventions are used:

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

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. You can use this information to troubleshoot problems, including how to view and package log files, use network connectivity tools, and identify known problems and resolutions.

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 to 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. A good description of the problem 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 question 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 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?

- Is the database 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. You can most easily do get a detailed timeline of events by working backwards. 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 can look only as far as the first suspicious event that you find in any diagnostic log. However, the first event 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 these questions help 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 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 have 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 the Tivoli Storage Productivity Center library, see <http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp>.

Tivoli Storage Productivity Center Technical Support website

See the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Technical Support website 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 Tivoli Storage Productivity Center for Replication 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 website: http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition.

Chapter 2. Tivoli Storage Productivity Center tools

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

Repository Copy tool

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

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

Note:

- You must stop the Device server and Data server before running **repocopy**.
- You cannot import **repocopy** data into the IBM Tivoli Storage Productivity Center database if the migration tool has not been run against the database. When you upgrade the Device server, you must run the migration tool before importing **repocopy** data.

Exporting repository data

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

To export repository data, follow these steps:

1. 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
2. Run the **repocopy.bat** or **repocopy.sh** command. A window is displayed prompting you for what you would like to do.
3. Select **Export data from repository tables**, and then click **Next**. The Options for Import/Export window is displayed.
4. In the Options for Import/Export window, enter information in the following fields:

Directory for Export

Enter the directory where the comma delimited file is saved.

Delimiter

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

Quote Enter the symbol that contains string data (double quotation marks is the default).

IBM Tivoli Storage Productivity Center exports the data into the comma delimited file you specify, and place it in a file named *tablename.txt*. Click **Next**.

5. The Choices for DB2[®] Export window are displayed. You have the following choices:

- Export Using DB2 native format
- Export Using text files

Make a choice and click **Next**.

- 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**.

- 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. Manually enter the database information. Click **Finish**.

- Tivoli Storage Productivity Center then connects to the database and displays the database and connection information. Click **Run** to begin the export process.
- A window is displayed containing the export progress log. As you progress through the export process, messages are written to this progress log. You can keep track of the steps through the progress log.

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 unavailable for loading `libTSRMinsudb.so`. You can ignore this message. 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:

- Stop the Device server and Data server.
- 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`
- Run the **repocopy.bat** or **repocopy.sh** command. A window is displayed prompting you for what you would like to do
- Select **Import data into repository tables**, and then click **Next**. The Options for Import/Export window is displayed.
- 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 quotation marks is the default). You must use the same double quotation marks that were used in the exported file.

Delete before inserting

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

IBM Tivoli Storage Productivity Center imports 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 then connects to the database and displays 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. You can use the progress log to track 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 unavailable for loading `libTSRMinsudb.so`. You can ignore this message. Click **OK** and continue.

Service tool: collecting information

This topic 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 compressed 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 class path
- 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 about 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 (information from the **db2level** command)
- Information from the **ipconfig /all** command (Windows)
- Information from the **ipconfig -a** command (Linux or UNIX)
- Information 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).
- Information from the scanners which are run manually from the command line.
- Everything in the log directory (including subdirectories).
- Everything in the conf directory
- Directory listing of the lib directory
- Information from the **ipconfig /all** command (for Windows).
- Information from the **ifconfig -a** command (for UNIX or Linux).
- Information 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 must 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, follow these steps:

1. Log on to 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. A compressed file is created in the directory where you ran the **Service** tool. The following compressed file is created:
 - `TPCServiceFiles.zip`

Running the Service tool for the agents

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

To run the **Service** tool for the Data agent, Fabric agent, or Storage Resource agent, perform the following steps:

1. Log on to the system. You must have administrator authority on Windows or root authority on UNIX or Linux.
2. For the Storage Resource agent, follow these steps:
 - a. In the navigation tree, expand **Administrative Services** → **Data Sources** → **Data/Storage Resource Agents**. In the content pane, select the agent. From the Select Action menu, click **Collect Service Data**.
 - b. A compressed file is created in the directory where you ran the Service tool. The following compressed files are created:
 - For Windows: `c:\Program Files\IBM\TPC\data\log\SRA_traces\SRA_computer_name\TPCServiceInfo.zip`
 - For UNIX or Linux: `/opt/IBM/TPC/data/log/SRA_traces/SRA_computer_name/TPCServiceInfo.zip`

where *SRA_computer_name* represents the name of the computer on which the Storage Resource agent is located.

- c. If the collection of service data fails, Tivoli Storage Productivity Center displays a message that indicates that the job was unsuccessful. See the server log file or the services script information for an indication of why the data collection failed. These files are located in the following default locations:
 - Server log file (on the computer where the Tivoli Storage Productivity Center server is installed):
 - For Windows: `c:\Program Files\IBM\TPC\data\log`
 - For UNIX or Linux: `/opt/IBM/TPC/data/log`
 - Services script information file (on the computer where the Storage Resource agent is installed):
 - For Windows: `c:\Program Files\IBM\TPC\SRA_computer_name\services\TPCServiceInfo.log`
 - For UNIX or Linux: `/opt/IBM/TPC/SRA_computer_name/services/TPCServiceInfo.log`

where *SRA_computer_name* represents the name of the computer on which the Storage Resource agent is located.

3. For the Data agent or Fabric agent, follow these steps:
 - a. 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: `/opt/IBM/TPC/ca/subagents/TPC/service`.
 - b. Run the following program:
 - For Windows: `service.bat`.

- For Linux or UNIX: `service.sh`.
- c. A compressed file is created in the directory where you ran the Service tool. The following compressed files are created:
 - TPCDATAservice.zip (for the Data agent)
 - TPCDEVCAservice.zip (for the Fabric agent)

Tracing the servers and agents

This topic provides information about the tracing tool. 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
- Agent
- Scheduler
- Device
- GUI

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

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 on one of these components:
 - Server
 - Agent
 - Scheduler
 - Device
 - GUI

Click **Configure Tracing**.

3. The tracing dialog window is displayed. Check the box for **Enable Trace** and enter a value for the following options.

Note: For the Agent, Scheduler, Device, or GUI, only the check box for **Enable Trace** and level are displayed.

Level The level can be:

DEBUG_MIN

The minimum debugging level. This option is the default and is always turned on. This option 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 option 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 option includes entry and exit information for every method in addition to all DEBUG_MID and DEBUG_MIN statements. This option can be a performance hit.

Maximum Number of Files

The maximum number of files used before the Data server starts reusing the tracing files. The default is five files.

Maximum File Size (kilobytes)

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 default trace level for an agent is modified through the Agent Tracing Configuration panel, those settings are only effective while the agent is running. If that agent is restarted, the defaults specified in the `nativelog.config` file are reread and used.

Traces are written by default to `TPC_installation_directory\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 content pane, select a Data agent or Storage Resource 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 option is the default and is always turned on. This option 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 option 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 option includes entry and exit information for every method in addition to all DEBUG_MID and DEBUG_MIN statements. This option can be a performance hit.

Maximum Number of Files

The maximum number of files used before the Data server starts reusing the tracing files. The default is five files.

Maximum File Size (kilobytes)

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 you can set for tracing a specific area 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 topic provides information about how to change the tracing levels for these areas. When collecting information for the investigation of problems, IBM support would provide information about 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 option contains large amounts of data logged in the trace files and is used generally when investigating specific issues with the product.

WARN

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

ERROR

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

ALL This option contains all warning and error messages.

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

To turn on tracing for the Device server for the Event Correlator, 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:

```
/opt/IBM/TPC/device/bin/aix (for AIX)  
/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 information is logged in `traceTPCDeviceServer.log`.

Tracing the Fabric agent

To turn on tracing 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:
`/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 information is placed in `C:\Program Files\IBM\TPC\ca\subagents\TPC\Fabric\log` for Windows and `/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. You can also see the traps that have been caught by the filters after they have been added. The traps are placed in the trace log file with a message indicating which filter caught the trap.

To turn on tracing 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:

```
/opt/IBM/TPC/device/bin/aix (for AIX)
/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 information is placed in `traceTPCDeviceServer.log`.

Trace log file locations

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

Table 1. Java trace log files

Class	Java trace log file
Storage Resource agent	<code>TPC_installation_directory/agent/log</code> <code>TPC_installation_directory/agent/log/</code> <code>name_of_server_SRA_communicates_with</code>
Data agent	<code>TPC_installation_directory/data/ca/subagents/TPC/data/log/</code> <code>hostname/AgentTrace.log</code>
Server	<code>TPC_installation_directory/data/log/ServerTrace.log</code>

Table 2. Native trace log files

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

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_installation_directory/data/ca/subagents/TPC/data/config` is configured

for Agent tracing. Alternatively, if you wanted to perform ImportExport tracing, make sure that `TPC_installation_directory/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 log files. The standard error log is redirected to a temporary file. The temporary message log file is copied to the appropriate log directory, and the standard error log file is redirected to reflect the new message log file location.

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

Table 3. Trace files for executable routines

Executable	Platform	Trace log file location, file name, and description
udbexec.exe	Windows	<p><code>\hostname\traceTSRMNativeUdbexec_PID.log</code> <code>TPC_installation_directory\data\ca\ traceTSRMNativeUdbexec_PID_yyyymmddhhmss.log</code></p> <p>This executable routine is run whenever a DB2 probe or scan occurs. A DB2 probe occurs upon registration of a DB2 instance.</p>
ExecSvc.exe	Windows	<p><code>temp\traceTSRMNativeExecSvc.log</code></p> <p>This executable routine runs only on the remote computer during a Windows push installation.</p>
udbexec.exe (Windows) or udbexec (UNIX or Linux)	Windows, UNIX, and Linux	<p><code>TPC_installation_directory/data/ca/ traceTSRMNativeudbexec_pid_yyyymmddhhmss.log</code></p> <p>This executable routine is run whenever a DB2 probe or scan occurs. A DB2 probe occurs upon registration of a DB2 instance.</p>
DbMonitor.exe (Windows) or DbMonitor (UNIX or Linux)	Windows, UNIX, and Linux	<p>This executable routine is run on Windows for all database monitoring activities. This executable routine generates trace files that are named base on the name of the job being run.</p> <p>For example, the default database scan generates a trace file similar to the following file: <code>TPCUser.Default Db Scan.0001_dbmonitor.00009163.trace</code></p>

Configuration files for tracing

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

Table 4. Trace configuration files

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

Table 4. Trace configuration files (continued)

File	Directory
Java tracing configuration file (for agent)	<i>TPC_installation_directory</i> /data/ca/subagents/TPC/data/Config/ AgentTraceLog.config

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

nativeLog.config file

During installation or a maintenance upgrade of Tivoli Storage Productivity Center, a file named `nativeLog.config` is created in the *TPC_installation_directory*/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 do not have the `#StopAgent` section that UNIX platforms have. The reason for the difference is because the agent service is started 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 about these settings and their default values, see “Tracing the Data server ” on page 10.

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 run 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 run 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 does 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 can be changed to reflect this condition. After this change is made, the server must 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 wanted trace level is applied.

Alternatively, trace settings can be changed during run time through the GUI. For more configuration information, see "Tracing the Data server " on page 10.

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 produces trace information.

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 produces trace information.

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 IBM Tivoli Storage Productivity Center

This topic 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 5. 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
Storage Resource agent	For Windows: c:\Program Files\IBM\TPC\agent\config\agent.config For UNIX or Linux: /opt/IBM/TPC/agent/config/agent.config

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 6. 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 opens (500)

Table 6. Parameters for server.config file (continued)

Parameter	Description
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. <ul style="list-style-type: none"> • 0 = Off. Do not encrypt the data stream. • 1 = On. Encrypt the data stream.
FileSystemScanNFSTimeout	Determines the numbers of seconds that a Storage Resource agent waits for a status system call on a network file system (NFS) before it times out.
hostAlias	This parameter is displayed 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: <ul style="list-style-type: none"> • Oracle: oracle.jdbc.driver.OracleDriver • MS SQL: com.inet.tds.TdsDriver • Sybase: com.sybase.jdbc2.jdbc.SybDriver • DB2: COM.ibm.db2.jdbc.app.DB2Driver • Cloudscape: com.ibm.db2j.jdbc.DB2jDriver
url	The URL used to connect to the database, normally: <ul style="list-style-type: none"> • Oracle: jdbc:oracle:thin:@host_name:port:SID • MS SQL: jdbc:inetdae:host_name • Sybase: jdbc:sybase:Tds:host_name:port • DB2: jdbc:db2:database_name
user	User name that IBM Tivoli Storage Productivity Center uses to connect to the repository

Table 6. Parameters for server.config file (continued)

Parameter	Description
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: <ul style="list-style-type: none"> • 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 7. 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 creates 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. Use this option to allow for the overhead time involved in scheduling a job so that the job starts 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:	
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	Set the maximum number of threads for handling requests (8)
pingReceiveTimeout	Number of seconds to wait before indicating that a ping has failed (10)

Parameter	Description
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	Set the maximum number of threads for handling user interface requests (10)
reportRowLimit	Maximum number of rows that are sent at a time to the user interface. If this number is exceeded, a More button is displayed above the table, along with a warning message (5000)
keepCachedReport	Number of minutes to retain incomplete reports in the tmp directory for the server (120)

agent.config file

The **agent.config** file contains configuration parameters for the Storage Resource agent. These parameters are set when the Storage Resource agent is installed. The parameters can also be changed manually by editing the file.

The following table contains the parameters for the **agent.config** file.

Table 8. Parameters for the agent.config file

Parameter	Description
Servername	Fully qualified host name of the system on which the Data server is installed.
Portnumber	Port on which the Storage Resource agent listens for requests. By default, the port is set to 9549.
IPAddress	IP address of the Storage Resource agent.

Log files

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

Installation log files

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

Installation log files for Tivoli Integrated Portal

Table 9. Installation log files for Tivoli Integrated Portal

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

Installation log files for Tivoli Storage Productivity Center

Table 10. Installation log files for Tivoli Storage Productivity Center

Log file	Directory location	Description
<i>component</i> Install.log or <i>component</i> InstallIS.log	UNIX: /opt/IBM/TPC/log/ <i>component</i> / install Windows: C:\Program Files\IBM\TPC\log\ <i>component</i> \install	Main Tivoli Storage Productivity Center log files for installation. <i>component</i> can be: <ul style="list-style-type: none"> • ca • cli • data • dbSchema • device • gui • agent (Storage Resource agent) • subagents/TPC/ (Data agent)
TIPInstallIS.log	UNIX: /opt/IBM/TPC/log/tip/ install Windows: C:\Program Files\IBM\TPC\log\tip\ install	Installation log for Tivoli Integrated Portal
TPCRInstallIS.log	UNIX: /opt/IBM/TPC/log/tpcr/ install Windows: C:\Program Files\IBM\TPC\log\tpcr\ install	Installation log for Tivoli Storage Productivity Center for Replication

Default log file locations

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

Table 11. Default log file locations for IBM Tivoli Storage Productivity Center and other components

Component	Log file location
Data Server	For Windows: <i>TPC_installation_directory</i> \data\log For UNIX or Linux: <i>TPC_installation_directory</i> /data/log
Device server	For Windows: <i>TPC_installation_directory</i> \device\log For UNIX or Linux: <i>TPC_installation_directory</i> /device/log
Storage Resource agents	<i>TPC_installation_directory</i> /agent/log <i>TPC_installation_directory</i> /agent/log/ <i>name_of_server_SRA_communicates_with</i>

Table 11. Default log file locations for IBM Tivoli Storage Productivity Center and other components (continued)

Component	Log file location
GUI	<p>For Windows: <i>TPC_installation_directory</i>\gui\log</p> <p>For UNIX or Linux: <i>TPC_installation_directory</i>/gui/log</p>
Tivoli Integrated Portal	<p>For Windows: <i>TPC_installation_directory</i>\logs\tip_directory <i>TPC_installation_directory</i>\tip_directory IA-TIP*.log files in C:\Documents and Settings\Administrstor <i>TIP_install_directory</i>_uninst <i>TIP_install_directory</i>\logs <i>TIP_install_directory</i>\profiles\TIPProfile\logs acsi\logs in C:\Program Files\IBM\common\acsi\logs</p> <p>For UNIX or Linux: <i>TPC_installation_directory</i>/logs/tip_directory <i>TPC_installation_directory</i>/tip_directory IA-TIP*.log files in / directory on AIX, /root directory on Linux <i>/TIP_install_directory</i>/_uninst <i>/TIP_install_directory</i>/logs <i>/TIP_install_directory</i>/profiles/TIPProfile/logs /acsi/logs in /usr/ibm/common/acsi/logs directory on AIX and Linux</p>
Note:	
<ol style="list-style-type: none"> 1. For Windows, the default directory <i>TPC_installation_directory</i> is C:\Program Files\IBM\TPC. 2. For UNIX or Linux, the default directory <i>TPC_installation_directory</i> is /opt/IBM/TPC. 	

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 is 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_installation_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 topic 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, collect pertinent information to help diagnose the problem. The following topics help you gather the information you need.

For specific troubleshooting topics, see the IBM Tivoli Storage Productivity Center Information Center.

General information

When a problem occurs with IBM Tivoli Storage Productivity Center, 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, DS4000[®], DS5000, or DS6000[™], CIM Agent version, if known.
- The last time that inventory collection was performed (this action 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 located in the local subnet 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 on), the actual start time of the job starts 1 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 actually starts at 1:29 PM. The 1:29 PM timestamp is displayed in the GUI for that job and in the job log files. This action takes into account the overhead time involved in scheduling a job so that the job actually starts close to the scheduled time.

Discovery

This topic provides information about troubleshooting discovery problems.

Discovery - general information

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_installation_directory\IBM\TPC\device\log

You can get a core dump file in this directory:

TPC_installation_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 command:
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
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 condition can be caused by the following:

- There is a lock contention when 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 7.

Unable to discover filers in a Windows environment

If you are unable to discover filers in a Windows environment, check the following items:

- Verify that there is an agent installed on a system in the Windows domain and that the agent can contact the filer. Verify that the Storage Resource agent is up and running and can communicate with the Data server.
- Check the license panel. If the filer is present, the filer was probably discovered but its file systems cannot/d not be retrieved because of invalid login information.
- Check the discovery log file:
 - For initial discovery, check the probe_000001.log file.
 - For discovery jobs, check the associated discovery log file.

- If the log file contains the following information, check the authentication information entered in the Filer Login panel:

```
DIS0001I FULL command selected.  
STA0246I Discovery started  
STA0339I Discovery started for domain xxx  
NAA0041E Local login failed -- SRM\db2admin  
NAA0041E Local login failed -- SRM\db2admin  
STA0249I Sending results to server  
STA0250I Server saved results
```

This information indicates:

- Authentication entered for a single filer is validated.
- The default login and password (used for multiple filers) is not validated.
- Verify that the SNMP community specified in the discovery job is the SNMP community that is used by the filer.
- Check the TPCUser.Discovery.*.trace log file. If the trace log file does not exist, set the agent tracing to the maximum value and rerun the discovery job.
The trace file for the discovery job contains helpful information about what hosts were found on the domain, which host is being contacted using SNMP, the SNMP community name used, and any error codes returned from the Windows API calls.

Unable to discover filers in a UNIX environment

Check the following items:

- Ensure that the agent can communicate with the filer and can mount a file system from the filer.

- Verify that the SNMP community name specified in the discovery job is the SNMP community that is used by the filer.
- Ensure that the mounts on the agent system are current and not stale. (Ensure that the mounted file systems still exist on the filer and were not deleted.)
- Check the TPCUser.Discovery.*.trace log file. If the trace log file does not exist, set the agent tracing to the maximum value and rerun the discovery job.
The trace file for the discovery job contains helpful information about what SNMP community name is used to communicate with the filer.

Other scenarios

Other scenarios that might occur are:

- If the **View Log File** button is disabled in the Job Management GUI after submitting a discovery job, click **Refresh All** until the log file name is updated in the GUI.

- If you see message STA0155E in the Windows discovery log file:

```
DIS0001I FULL command selected.
STA0246I Discovery started
STA0339I Discovery started for domain xxx
STA0155E Discovery was unable to find the Windows computers in the domain
        or workgroup.
STA0249I Sending results to server
STA0250I Server saved results
STA0252I Discovery completed successfully.
```

If you receive these messages, turn tracing on to generate a discovery trace file. Check the trace log for error code 6118. If you see code 6118, try starting the Computer Browser service on the Windows system, if the service is not turned on.

- If a filer is added using the **Other NAS** option, and the filer is deleted and added again without specifying the **Other NAS** option, Tivoli Storage Productivity Center continues to recognize the filer as **Other NAS**.
- For UNIX agents: You see error message STA0096E when you manually enter a filer for a UNIX agent. If you see this message, check the mounts on the UNIX system (there is a stale NFS handle). Unmount the reported file to proceed.
- For Windows agents: For hosts that are on different subnets but have the same workgroup name, Tivoli Storage Productivity Center considers the workgroup is on the same subnet and excludes one subnet when performing a discovery job. To work around this issue, disable the Storage Resource agent on one subnet and then perform the discovery job again.

Configure Devices wizard

The Configure Devices wizard is designed to simplify the discovery and configuration steps for Tivoli Storage Productivity Center to manage devices.

Diagnosing a problem

To determine which component is causing an issue depends on the error and where the error occurs:

- If the error occurs within the user interface flow of the wizard, then the issue is with the wizard user interface code.
- If there is an error during the discovery of the devices, then the issues are the same as it would be with running a normal discovery of the devices outside of the wizard. For example, the error might occur with data source issues, firewalls, problems with the Device server, and so on.

- If there is an error after the Summary step, these errors are displayed in the Results step. The Configure Devices wizard might be trying to perform one of the following actions:
 - Adding a device to the monitoring group
 - Creating a probe, ping, or alert definition
 - Database connection issues

with the Data server request handler code for the wizard

Logs

The Configure Devices wizard uses the existing Tivoli Storage Productivity Center Data server infrastructure and, therefore, uses the same trace logs as the Data server processing.

Logging is turned on by default and any data is logged in *TPC_installation_directory/data/log/server_xxx.log* and the discovery logs.

The tracing log is in *TPC_installation_directory/data/log/ServerTrace.log*.

If an error occurs, check the following logs. Which log to check depends on the time the error occurs.

- Data server logs
- Device server logs
- Other logs that depend on the type of error that occurs, for example, during the discovery process, schedule definition creation, and so on.

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 Tivoli Storage Productivity Center components.

Tivoli Storage Productivity Center displays an event for the discovery and control components for the following actions:

- When you add a new device
- The device status changes
- The device is not detected

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 request made to the performance monitoring service.

For general diagnosis of the problem, check these items:

- Make sure that 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 items 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?

Many database activities to form SNMP and Tivoli Enterprise Console® 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 unavailable in the database or there are network communication problems. After the alert creation fails, the alert for this changed device is not recreated.

Why do I see unexpected alert definitions and unexpected alerts with the keyword "Pre-defined" in the name?

These alerts are displayed when another application has subscribed to receive events from Tivoli Storage Productivity Center through the Tivoli Storage Productivity Center API.

You can change the setting to hide these "Pre-defined" alert definitions by using the **tpctool** command:

```
tpctool setdscfg -user user_ID -pwd password -url localhost:9550
-property APIAlert.Visibility false
```

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 **repopcopy** tool to collect trace information for Tivoli Storage Productivity Center and the database.

When an expected alert is not presented, check the following items:

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

Database monitoring

This topic provides information about what to look for when diagnosing problems with database monitoring.

Agent side issues or failures

If you have an agent issue or failure related to database monitoring, the likely cause is usually related to a command or query that is sent to the database. The agent logs indicate that the job has failed.

If you have an agent issue or failure during database monitoring, follow these steps:

1. Turn the tracing option on and set the tracing option to `DEBUG_MAX` for the Storage Resource agent.
2. Run the specific job that caused the issue (for example, a probe or scan job).
3. Check the trace log for the specific job run. The trace log includes information about the specific commands run against the database and also captures any error messages from the database.

Server side issues or failures

If you have a server issue or failure related to database monitoring, the likely cause is usually related to the server not being able to capture or process output from a probe or scan job. The agent logs indicate that the probe or scan job completed successfully.

If you have a server issue or failure during database monitoring, follow these steps:

1. Enable the tracing option and set the tracing option to `DEBUG_MAX` for the Data server and agent.
2. Run the specific job that caused the issue (for example, a probe or scan job).
3. Check the `ServerTrace.log` file on the Data server. Look for error messages listed with the following information:
`com.tivoli.itsrm.Db.server.svp.DbScanResultsZHndlr` (for scan jobs)
`com.tivoli.itsrm.Db.server.svp.DbProbeResultsHndlr` (for probe jobs)

Performance monitoring

This topic 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 20%, so that if less than 20% of the

ports for a CIMOM are targeted for performance data collection, the association algorithm is used instead, if the **MaxSwitchPortLimit** is not exceeded (see **MaxSwitchPortLimit**). This parameter 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 is not 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 number 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 must be restarted for these changes to take effect.

File system scans

If you have a file system scan issue, check this information for probable causes of the problem.

Some of the probable causes of file system scan issues are as follows:

The specified file system is not found on the Storage Resource agent system

Follow these steps to check for the file system:

1. Check the Storage Resource agent system to see if the file system is still available.
2. Run the computer probe job again on the Storage Resource agent system.

The file system is full

Reasons why you see this issue:

- The file system scan job did not start because a scan definition file cannot be copied to the Storage Resource agent system.
- During a scan or probe job, the log and trace files cannot be written to the Storage Resource agent system because the file system is full.

If the file system is full, free file system space by removing files, directories, or components that are not needed.

File system scan job abruptly ended

The file system scan job abruptly ended because Tivoli Storage Productivity Center indicates that the file system is already being scanned. This condition occurs when there is no file system scan currently running on the Storage Resource agent system.

The reasons for this condition are as follows:

- The file system scan might have failed. Check the agent log and trace files for error messages.

- The file system name was not properly removed from file `sra_scan_res` after the file system was scanned.

If you have this issue, delete file `sra_scan_res` in the following directory:

`SRA_directory/agent/log`

Rerun the file system scan job again.

The file system scan job cannot log on to the NAS filer

The file system scan job cannot log on to the NAS filer to scan the NAS filer.

If you have this issue, follow these steps:

1. Check Tivoli Storage Productivity Center to see if the correct user ID and password were entered for this NAS filer.
2. Try to manually mount and access the filer volume on the Storage Resource agent system.

The NAS filer volume or file system was not found

If you have this issue, check the following item:

- If the Storage Resource agent system still has the volume mounted.
- If the network connection to the NAS filer is still available.
- If the NAS filer volume is still available.

Correct the problem and rerun the computer probe job again.

The file system scan job takes a long time

The file system scan job takes a long time for a small file system.

If you have this issue, follow these steps:

1. Check to see if the tracing option is set for this Storage Resource agent. If the tracing option is set to `DEBUG_MAX`, there might be numerous tracing information written.
2. If the tracing option is set with `DEBUG_MAX`, lower the setting to `DEBUG_MIN` and rerun the job.

The file system scan fails to scan the clustered file system

If you have this issue, check the following items:

- Make sure that all nodes have either daemon Storage Resource agents or legacy Data agents installed and running.

If you want to use non-daemon Storage Resource agents, see issue 52701 in the Tivoli Storage Productivity Center 4.2 flash.

- Probe all nodes in the cluster.
- Make sure that the cluster resource group is configured with an IP address. In MSCS, you must configure a Network Name resource with an IP address resource, not just an IP address resource.
- Make sure that the IP address for the cluster resource group can be contacted by the Tivoli Storage Productivity Center server. Make sure that Tivoli Storage Productivity Center has access through any firewall.
- Make sure that Tivoli Storage Productivity Center can find the IP address for the cluster resource group when the clustered nodes are probed. Tivoli Storage Productivity Center collects one IP address per Network Name resource. Check the computer entry for the cluster resource group in Data Manager. In the navigation tree, click **Data Manager** → **Reporting** → **Asset** → **By Cluster** → **cluster_name** → **Cluster Resource Groups**. If the IP address listed is not the one that can be

contacted by the Tivoli Storage Productivity Center, check the
T_RES_IP_ADDRESSES table using the following query:
db2 select network_name, IP_address from tpc.t_res_ip_addresses

Topology viewer

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

Using debugging mode for the topology viewer. You can use the topology viewer to see the following information:

- 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 on).
 - The ID of that entity. You can use this ID to easily look up that entity in the database tables.

You can see the data model exactly as you see what is 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 must 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 turn on debugging mode, a flag must be passed to Java when starting the GUI:

```
-DTopologyViewerDebug=5
```

This action causes the topology viewer to provide debugging information to stdout (and show an extra debug ID column in the tabular view). Note that if Tivoli Storage Productivity Center is started using javaw.exe, you do not see any output data because javaw.exe hides the console window. javaw.exe is the default selection in the Windows shortcut.

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

```
-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 Tivoli Storage Productivity Center.

```
"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 statement:

```
"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 must edit the `TPCD.sh` script in the following directory:
`/opt/IBM/TPC/gui/`

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

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

Common usability problems

Note the following actions when linking the topology viewer with Tivoli Storage Productivity Center reports and alert logs:

- The user is not 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 is 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 topic provides information about 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 can see these parameters in the `DataStore.properties` file in directory `TPC_installation_directory/device/conf`:

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

Traces are located in the following file:

```
TPC_installation_directory/device/log/traceTPCDeviceServer.log
```

In addition, you can 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 Tivoli Integrated Portal, this file is located in the following directory:

```
TIP_installation_directory/profiles/TIPProfile/properties
```

Storage Optimizer

The Storage Optimizer uses data in the Tivoli Storage Productivity Center database to analyze your storage subsystems to identify performance bottlenecks, and recommend changes to improve performance. 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 modifications to subsystem configurations. Its primary purpose is to provide you with a performance analysis and optimization recommendations that you can select 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 might produce unexpected results in the Storage Optimizer heat maps. Storage Optimizer analysis of subsystems containing solid-state drives must be avoided until support is available.

Important: Storage Optimizer does not take into account any established replication relationships or sessions. Migration recommendations must 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
- TotalStorage 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 is less accurate if there are any gaps in the data collection for the time interval being analyzed. It is recommended that you collect at least one week's worth of performance monitoring data before using Storage Optimizer. Providing a longer time interval for data collection will increase the accuracy of the Storage Optimizer analysis and recommendations. To create a subsystem performance monitor job, navigate to **Disk Manager > Monitoring > Subsystem Performance Monitors**. Right-click **Subsystem Performance Monitors** and click **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 is included in the analysis if it is available.

The Storage Optimizer provides the following output data:

- 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 data.
- 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 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_installation_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_installation_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_installation_directory\device\conf

Rollup reports

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

Through rollup reports, you can view storage information collected by the master Tivoli Storage Productivity Center server from its associated subordinate servers. Rollup reports provide you with a network-wide perspective of storage usage in an environment where multiple Tivoli Storage Productivity Center servers are deployed and collect storage metrics.

Things to note

Note these conditions:

- The master server must not monitor more than 500 unique sources.
- Subordinate servers must not have more than 1200 unique data sources. This condition includes Data agents, Fabric agents (in-band and out-of-band), CIM agent, Storage Resource agents, and VM agents.
- You must have Tivoli Storage Productivity Center superuser or 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 (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 Tivoli Storage Productivity Center functions such as the topology viewer, data path explorer, volume provisioning, volume performance, SAN planner, and so on, 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_installation_directory/data/log/server_XXXXXX.log
TPC_installation_directory/data/log/TPCD_XXXXXX.log

You can look at these log files from the GUI:

1. Open the Tivoli Storage Productivity Center GUI and click **Administrative Services** → **Services** → **Data Server** → **Server**.
2. For the server, right-click on **Server**. Click **View Log**. The log is displayed in the right pane.

3. For the GUI, right-click on **GUI**. Click **View Log**. the log is displayed in the right pane.

- Data Server trace log file on the master server:

TPC_installation_directory/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 10.

- Device Server log on the subordinate server:

TPC_installation_directory/device/log/msgTPCDeviceServer.log

- Device Server trace log on the subordinate server:

TPC_installation_directory/device/log/traceTPCDeviceServer.log

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

SAN Planner

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

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

Supported Subsystems

SAN Planner supports TotalStorage Enterprise Storage Server, IBM System Storage® DS6000, IBM System Storage DS8000, and IBM System Storage SAN Volume Controller. SAN Planner supports the **Space Only** workload profile option for any other storage system.

Dependencies

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

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

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

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

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

A storage controller performance monitor must be run to select the **Workload Profile** options. For SAN Volume Controller, it is required that the performance data is collected for all its backend subsystems as well. Additionally, the backend subsystems must be among DS8000, DS6000, DS4000, or ESS. If a performance monitor has not been run, the **Space Only** workload profile option is permitted for planning based on available capacity only.

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

Note:

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

Run the following command on the Solaris system:

```
cfgadm -al
```

If the system displays this message:

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

run this command:

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

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

You must reboot the Solaris system.

- The SAN Planner does not support Easy Tier™ (mixed pools) as input to a plan for automated storage provisioning. These pools are not displayed when you select resources for planning or request a list of recommended pools.

For replication planning, you must have an appropriate Tivoli Storage Productivity Center for Replication license installed and the appropriate subsystem device feature codes enabled. You must add all of the subsystems to be used in planning to Tivoli Storage Productivity Center for Replication. You must also ensure that connectivity is established between the subsystems.

Diagnosing the problem

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

1. Run the **srmp** tracing command.

Windows

To run the **srmp** 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:
/opt/IBM/TPC/device/bin/aix (for AIX)
/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 data is logged in traceTPCDeviceServer.log.

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

TPC_installation_directory/data/log/server_0000xx.log

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

Log files

SAN Planner involves two major operations:

- Plan recommendation or planning
- Plan execution

Planning

If you encounter problems with planning in SAN Planner, check the following logs:

TPC_installation_directory/device/log/tracePlanner.log
TPC_installation_directory/device/log/msgPlanner.log

These logs are turned on by default and provide detailed information about the problem.

If you encounter problems during the planning process, and the problem is related to replication, you can investigate the Tivoli Storage Productivity Center for Replication client API log. This log is available at:

TPC_installation_directory/device/rm<n>.log

Where <n> is the job run number. You can also investigate the Tivoli Storage Productivity Center for Replication server logs at:

TPC_installation_directory/eWas/profiles/CSM/logs/CSM/csmTrace.log

Plan execution

When running a plan generated through SAN Planner, the SAN Planner job log provides detailed information about the job. You can look at these log files from the GUI or they can also be found in:

```
TPC_installation_directory/data/log/user.planner.xxx.log
```

For more detailed information about all the debug level details, you can look at the following log file:

```
TPC_installation_directory/data/log/TracePlanner.log
```

At plan execution Tivoli Storage Productivity Center for Replication client APIs are called to perform replication operations. To investigate any replication execution issues, consult the following logs:

```
TPC_installation_directory/data/log/rm<n>.log
```

Where <n> is the job log number.

Multipath subsystem device driver

To check the multipath subsystem device driver (SDD) used by the SAN Planner, run the appropriate following command:

Linux SDD

Check to see if the SDD is installed by running the following command:

```
lsmod|grep sdd
```

To install the SDD, run the following command:

```
sdd start
```

Check to see if any device is manually excluded by running the following file:

```
ls -la /etc/vpath.conf
```

If a device is excluded, uncomment the entry in the file and run **cfgvpath**.

AIX SDD

Check to see if any device is manually excluded by running the following command:

```
ls -la /etc/vpexclude
```

If a device is excluded, delete the entry from the file. Run **cfallvpath**.

AIX SDDPCM

Check to see if any underlying path status is in the missing or disabled state by running the command **lspath**. Run the following command for the corresponding adapters:

```
rmdev -d1 fscsiX -R
```

Run **cfgmgr**.

HP-UX and Solaris SDD

Check to see if any device is manually excluded by running the following command:

```
ls -la /etc/vpathmanualexcl.cfg
```

If a device has been excluded, delete the entry from the file. Run **cfgvpath**.

Common usability problems

Some of the common usability problems are:

SAN Planner planning operation fails with a timeout error

SAN Planner planning runs through the GUI. If planning takes longer than the GUI timeout value of 45 minutes, the result is an internal error from the timeout. For example, this timeout can occur when trying to get a plan with hundreds of volumes requiring replication. The alternative is to create more than one plan with a limited number of volumes in each plan. If the GUI does not timeout, but takes a long time to generate the plan, an additional GUI can be started to continue working with Tivoli Storage Productivity Center while the plan is being generated.

SAN Planner plan execution fails with some changes made

SAN Planner plan execution does not have a rollback feature. Depending upon where the plan execution failed, you might or might not want to undo the changes made by planner. Examine the job log to see what changes have been made and you can decide on whether to undo any of those changes or not.

Specifying the RAID level for SAN Volume Controller virtual disks planning fails with a 'not enough space available' message

For SAN Volume Controller virtual disks planning, the RAID level input is for identifying the RAID level of the backend subsystems. If you specify a RAID level other than **auto**, then you must have all of the backend subsystems added to Tivoli Storage Productivity Center. SAN Planner uses their information from the Tivoli Storage Productivity Center database to find the RAID level. If you do not have all the backend subsystems configured into Tivoli Storage Productivity Center, you see the message that there is not enough space available.

When the planning operation fails with a 'not enough space available' message, the available space displayed is less than the actual available space

In the message, the available space displayed is the amount of available space in the storage pools that are eligible for the plan recommendation based on the user requirements. If no storage pools meet the requirements, then the available space displayed is zero. You can determine the eligible storage pools by selecting the **Show Recommended Pools** option on the **Specify Advanced Capacity Plan** step of the SAN Planner wizard.

There is also a possibility that the size of the volume requested does not span a complete extent. For example, if the extent size on the subsystem is 1 GB, but the size of each volume is 512 MB, then there is a 512 MB extent overhead. The available space displayed in the message only accounts for the size of the volumes to be created within the eligible storage pools.

New zone set is created when 'use active zone set' is selected in the SAN Planner wizard

If the fabric selected to perform the zone configuration does not have an active zone set at plan execution time, a new zone set is created. You can use Fabric Manager to determine whether the recommended fabrics have active zone sets.

New zones are created when selecting zone planning but there are existing zones for the recommended ports in the plan

The purpose of zone planning is to specify new zones for the recommended ports. If zones for the recommended ports exist in the selected fabrics, then zone planning must not be selected. You can use

Fabric Manager to determine if zone planning is necessary for a host and a subsystem or SAN Volume Controller. If the fabric selected to perform the zone configuration does not have an active zone set at plan execution time, then a new zone set is created. You can use Fabric Manager to determine whether the recommended fabrics have active zone sets.

Paths with no zones are recommended by SAN Planner when zone planning is not specified

SAN Planner does not require a fabric agent in order to recommend a plan. Use the fabric agent to discover and change the zones and zone sets during plan execution. If a fabric agent does not exist, then the zones might not be populated and cannot be relied upon during the planning operation. If zones do not exist between the host ports and the subsystem or SAN Volume Controller ports, then new zones are needed. You can create new zones either through zone planning, or if no fabric agent exists, create zones through another zone creation utility before running SAN Planner.

In SAN Volume Controller planning, SAN Planner reports no common fabrics between the host and SAN Volume Controller but a common fabric exists

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

SAN Planner recommends a replication plan between two DS8000 systems, but there is no connectivity between the two systems

SAN Planner can create a replication plan, but it does not check replication connectivity during planning. If you want SAN Planner to run a replication plan, then you must ensure that there is connectivity between the two systems.

SAN Planner planning with replication operation for SAN Volume Controller fails with 'unable to find suitable placement for replication volumes' error

Planning with replication for SAN Volume Controller is different to planning with replication for other storage systems. If there is no connectivity between SAN Volume Controller systems, SAN Planner cannot create a replication plan and throws the 'unable to find suitable placement for replication volumes' error. If you want SAN Planner to generate a replication plan, then you must ensure that there is connectivity between the systems.

SAN Planner job reports that Global Mirror plan executed successfully, with replication session created and started successfully. However, in IBM Tivoli Storage Productivity Center for Replication, the session status indicates errors

Tivoli Storage Productivity Center for Replication runs Global Mirror sessions in two phases:

1. The long-distance copy relationships (Global Copy or H1-H2) are created and the initial copy is run.
2. The master copy relationship is created and run (Global Mirror or H1-J2).

SAN Planner reports the success of the plan based on the first phase. The second phase in a normal environment might take a number of hours.

Configuration History and Analysis

This topic 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 and provides information about 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 condition might happen when an entity gets removed manually from the database or by the removed resource retention function, and 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 items 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 are 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 limitation is done to prevent overloading the topology viewer with numerous 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, Storage Resource 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 or Storage Resource agents to collect the host information (operating system type and version). This information is required for the evaluation of configuration analysis policies 3 and 9.
- In-band Fabric agents or Storage Resource agents to collect the attribute information from the host information (HBA and host port details). This information is required for the evaluation of configuration analysis policies 2, 3, 4, 8, and 9.
- CIM agents for tape to collect the tape information. The agents are required 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 conditions when you encounter a problem with Configuration History and Analysis:

- Configuration Checker job logs are located here:
*TPC_installation_directory/data/log/*config_analysis*.log*
- Message log for the Device Server. The log file is at this location:
TPC_installation_directory/device/log/msgTPCDeviceServer.log

The message prefix is "HWNCC."

- Device Server trace log.

To turn on tracing 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:
/opt/IBM/TPC/device/bin/aix (for AIX)
/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 data is logged in `traceTPCDeviceServer.log`.

- Data Server trace log. For information about turning on tracing for the Data Server, see “Tracing the servers and agents” on page 10. Select `DEBUG_MAX` when turning on tracing. The log file is at this location:

`TPC_installation_directory/data/log/ServerTrace*.log`

Data Path Explorer

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

The Data Path Explorer is a type of view in the topology viewer. Data Path Explorer combines several of the usual topology views into a visualization that allows users to follow connectivity from one end point in the topology view to another. Storage administrators can use this view to debug connectivity and performance problems in the storage environment. Data Path Explorer displays the different data paths (SAN access paths or I/O paths) for each host disk. You can 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 displaying 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 previously?

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_installation_directory/data/log/server_*.log`

To look at the log file, follow these steps:

1. Open the Tivoli Storage Productivity Center GUI and click **Administrative Services** → **Services** → **Data Server** → **Server**.
 2. For the server, right-click on **Server**. Click **View Log**. The log is displayed in the right pane.
 3. For the GUI, right-click on **GUI**. Click **View Log**. the log is displayed in the right pane.
- Server trace log. For information about how to turn on tracing, see “Tracing the servers and agents” on page 10. When turning on tracing, specify DEBUG_MAX. The log file is located here:

```
TPC_installation_directory/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.

```
"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 statement:

```
"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 topic provides information about what to look for when diagnosing problems with the SMI-S fabric probes.

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 Tivoli Storage Productivity Center Storage Resource agent (or in-band Fabric agent) or out-of-band Fabric agent.

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

This topic 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 are 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 might 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 gathers 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 gathers 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 action lets you 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 an 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* for the SMI Agent at the Brocade website. For information about the Brocade SMI agent for EOS switches (also known as McDATA), see *Brocade SMI Agent for EOS Products User Guide* at the Brocade website. See the appropriate version of the SMI Agent you are using. Also see the documentation for the *DCFM Integrated SMI Agent* for any problems encountered with this new type of SMI Agent.

For information about the Brocade DCFM SMI Agent, see the *Brocade Data Center Fabric Manager User Manual* that supports your version of DCFM.

Things to note

- 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 actions occur:
 - Rerun the CIMOM discovery.
 - When a backend fabric agent assignment is made, Tivoli Storage Productivity Center checks to see which fabric agents are available. The fabric agents are sorted and the top agent in the list is run. The other agents in the list are run only if the fabric probe fails with the previous fabric agent.

NAS and NetApp

This topic provides information for troubleshooting NAS and NetApp problems.

NAS probes

The best way to troubleshoot NAS probe issues is to check the log and trace files for the Storage Resource agent that is performing the probe job.

The log and trace files are found under the following directory:

```
SRA_directory/log/TPC_server_name
```

These files use the following naming convention:

```
schedule_name.schedule_number_naprobe.job_run_number.log  
schedule_name.schedule_number_naprobe.job_run_number.trace
```

NetApp ImportQuota

Check the following items:

- Make sure the NetApp discovery and probe jobs have collected the basic NetApp filer information.
- If the ImportQuota job state is "No filers found," verify that there is still a Storage Resource agent configured to probe and scan the filer.
- If numeric IDs are reported with the quotas for Windows-SIDs instead of user names, make sure to add an additional Storage Resource agent for the respective platform type.
- If you see an alert generated but does not show up in Tivoli Storage Productivity Center, verify that you have a current file system scan for the NetApp volume in question.

Note: Resolving all user names that might be involved with quotas is a two-step process requiring two agents of different operating system types. For example, you must configure an agent for Windows and UNIX. As a result, the corresponding agent log files are also found on different systems.

The log and trace files found on the primary Storage Resource agent (where the ImportQuota job runs) has the following naming convention:

```
schedule_name.schedule_number_netapp_quota.job_run_number.log  
schedule_name.schedule_number_netpp_quota.job_run_number.trace
```

The log and trace files found on the auxiliary Storage Resource agent (for resolving additional user IDs, usually a Windows Storage Resource agent) has the following naming convention:

```
netapp_resolveids_run_number.log
netapp_resolveids_run_number.trace
```

VMware ESX

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

Tivoli Storage Productivity Center 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 the ESX Servers for the data center. You would install a Data agent on each virtual machine that you want to monitor.

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 Tivoli Storage Productivity Center.
- The information from VMWare is equivalent of a read-only Data agent (no scripts, no FS extension).

To get detailed information for space and capacity reports, you must perform the following steps:

- 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

Some common user errors are:

- 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

Some of the dependencies are:

- VMWare ESX 3.0.1 and Virtual Center 2.0.1
- The hypervisor and Data agent or Storage Resource agent probes are required for complete reporting.
- Hypervisors and VirtualCenters have https communication turned on and http communication turned off by default.
- If you are using the https protocol, make sure that you import the VMWare certificate into the truststore before selecting the https protocol:

- The truststore must be created under *TPC_installation_directory/device/conf/vmware.jks*.
- The VMWare certificate for the hypervisor can be found on the hypervisor under the location */etc/vmware/ssl/ruicert*.
- The VMWare certificate for the VirtualCenter can be found in the following location:

```
\Documents and Settings\All Users\Application Data\
VMware\VMware VirtualCenter\SSL\ruicert
```

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

```
TPC_installation_directory/jre/bin/keytool
```

- The command to import the truststore is:

```
keytool -import -file ruicert -alias alias
-keystore vmware.jks
```

Diagnosing the problem

When troubleshooting a VMWare problem, note the following items:

- Look at the job logs for discovery and probe create jobs which might indicate that tracing information is available.
- Look for trace messages in the general trace file:

```
TPC_installation_directory/device/log/traceTPCDeviceServer.log
```

- Set the trace level.

To turn on tracing 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:

```
/opt/IBM/TPC/device/bin/aix (for AIX)
/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 -filterkey INFO
```

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

The output data is logged in *traceTPCDeviceServer.log*.

Reporting groups

This topic 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 are displayed in different monitoring groups. Because a resource can be displayed 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 the problem

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

- For AIX or Linux, the memory dump is displayed in the shell window
- For Windows, the memory dump is displayed in the DOS window

You must run the GUI manually using the following command:

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

You must also perform the following actions:

- Turn tracing on in the GUI. For information about turning tracing on, see “Tracing the servers and agents” on page 10.
- Run **repocopy** of the IBM Tivoli Storage Productivity Center database. For information about how to run **repocopy**, see “Repository Copy tool” on page 5.
- Run the **service** tool for both the servers and common agents. For information about how to run **service** tool, see “Service tool: collecting information” on page 7.
- Get DB2 logs (db2diag.log).
- Get output data from the **db2 list applications** command. To get output data, run these commands:

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

Monitored Computer Storage Space reports

This topic 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 list 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 information 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 is listed when Tivoli Storage Productivity Center has not probed the storage subsystem from where the storage originates. If Tivoli Storage Productivity Center has probed the storage subsystem, this condition indicates a problem with the way the storage subsystem is retrieved by the probe of the host. This condition 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 displayed 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 Tivoli Storage Productivity Center. This field is a user-defined storage subsystem name to represent the storage subsystem that was probed.

Unknown Storage Subsystem

If Tivoli Storage Productivity Center is unable to identify the storage subsystem, the storage subsystem name is set to unknown. This condition 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 listed in the All Disk Storage reports or the Disk Storage On Storage Subsystem reports. This column indicates that Tivoli Storage Productivity Center has been 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. The rows that indicate that 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 this condition 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 condition requires a reprobe of the host.
- The storage subsystem storage was not correctly identified by the host probe. This condition 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. Tivoli Storage Productivity Center was not able to uniquely identify the disk storage across more than one host. This

condition 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 Tivoli Storage Productivity Center was unable to query the device to extract a serial number that would uniquely identify the disk. This condition is known 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 must 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 condition occurs for network-attached storage, remote file system mounts, and temporary file systems on the processor cache. These reports only report on these types of file systems. If there are logical volumes or file systems that are not the expected types, then this condition 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.

Diagnosing 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`. After you have configured tracing, the probe on the host system must be run again. If there is a problem with the correlation of the storage subsystem volumes or the storage subsystem name, then the Tivoli Storage Productivity Center server logs must 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

Tivoli Storage Productivity Center FlashCopy® support labels volumes as having a FlashCopy property of "source", "target", or "none" as appropriate. This action 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 topic provides information about what to look for when diagnosing problems with FlashCopy support.

Tivoli Storage Productivity Center does not provide the following information or actions:

- Additional information about FlashCopy relationships
- The Tivoli Storage Productivity Center user cannot create or manipulate FlashCopy relationships
- A report (or any information) that links together the source and target volumes of a particular FlashCopy relationship

Diagnosing the problem

Most issues arise from one of two causes:

Missing data

This condition occurs when the subsystem does not report the FlashCopy data in a manner that Tivoli Storage Productivity Center expects. This action is most likely to occur in the event of the release of new microcode or CIM agents for a subsystem.

Note: The DS8000, SAN Volume Controller, and XIV storage systems do not require CIM agents. These storage systems use native interfaces. Check the following conditions:

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

Improperly created FlashCopy relationships

This condition occurs when the user did not create a persistent FlashCopy relationship. If the relationship is temporary, chances are it the relationship is gone by the time Tivoli Storage Productivity Center's probe requests the information. Verify that the subsystem's management software still reports the FlashCopy relationship.

Element Management

The Element Management is integrated with the Tivoli Storage Productivity Center GUI. You can manage multiple DS element managers within the Tivoli Storage Productivity Center GUI.

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 10.

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.ElementManagerMgmtTrclLogger -filterkey INFO
```

The filter 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_installation_directory\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_installation_directory/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 can then view the JSP log entries in the GUI panel.

McData Intrepid 10000

This topic 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 Tivoli Storage 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 Tivoli Storage Productivity Center, the enumeration approach is used when the monitored ports are over 20% 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 timeout.
- Increase the collection interval time from the default of 5 minutes to 10 minutes, for example. (The McData switch is also set to 5 minutes.)
- 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 must be increased to 1 GB when monitoring more than 300 ports.

Increasing the Java heap size for McData CIMOM

If you are managing medium or large fabrics, increase the heap size of the JVM to 1 GB (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 by running the following command:

```
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 **McdData_ManagedSwitch** instance in the /root/McdData namespace.

If you want to check which virtual switches for the McData 10000 show up in Tivoli Storage Productivity Center, check the **McdData_Switch** instance in the /root/McdData 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_installation_directory/device/log/tracePerfMgr.log  
TPC_installation_directory/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 data is shown.

```
2007-04-0119:15:00,109-07:00Co1Thread8 PmFabricCache  
isEnumerationPreferred  
INFO Selected Enumeration; CIMOM https://xxx.xxx.xxx.xxx:5989;  
TimeSlot 3; 267 / 298.0 = 0.89593295880014981; MinPortRatio 0.4,  
MaxPortLimit 256  
...  
2007-04-0119:15:00,156-07:00Co1Thread8 CounterDataServiceSwitch  
collectFCPortAndRateStats  
INFO Using Enumeration approach  
...
```

tpctool

Tpctool is a stand-alone Java client and connects to the Device server only. **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.

Tpctool is installed in the following default directories:

```
C:\Program Files\IBM\TPC\cli (for Windows)  
/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 action 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. 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 is updated in the database and the Device server is notified.
- The configuration file for all the Fabric agents must 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**

A common error message for **tpctool** is:

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_installation_directory\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_installation_directory\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 is filtered for the specified filter:

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

- To remove a specified filter:

```
cd TPC_installation_directory\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_installation_directory\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_installation_directory\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_installation_directory\device\bin\w32-ix86  
srmcp -u user_ID -p password ConfigService setPw  
new_host_password
```

Fabric-specific problems

This topic 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. See *IBM Tivoli Storage Productivity Center Messages* for the error codes. The error codes have two parts: command failure code and native failure code.

- If the appropriate agent is available for zone control. The capabilities available depend on the agent.
- If the agent is connected to the SAN and if the fabric is configured appropriately for zoning.
- If the Storage Resource agent is connected to a McData fabric, and that agent is expected to be used for zone control, check to see if Open Systems Management Server (OSMS) and host control is turned on for the McData switches in the fabric.
- For zone configuration failures with a McData switch: If a fabric is in Interop Mode (open mode), only port WWN zoning is supported. This condition applies whether it is a heterogeneous fabric or contains switches from only a single vendor. Attempts to use other zone member types fail in this configuration. Use port WWN zoning or consider moving the fabric out of Interop Mode (open mode) if it is not required.

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 devices using the native interface 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: 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 items:

- Look for the switch management tools:
 - Firmware
 - Topology
 - Current[®] zoning configuration
 - Transaction lock (cfgTransShow and cfgTransAbort).
- Look for 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 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 7.
2. Run **repocopy** to capture the 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 double quotation marks.
 - f. Take the default values for the connection properties.
 - g. Compress 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 items:
 - 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 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 can 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 Tivoli Storage Productivity Center GUI.

Fabrics not discovered by the Storage Resource agent

Check the computer probe log to identify any additional information that is helpful to isolating the problem. Check the following items:

- If the HBA is connected to the fabric. Confirm through operating system commands or the switch element manager or both that the HBA is connected to the fabric.
- If the HBA is connected to a McData fabric, check to see if Open Systems Management Server (OSMS) is turned on.
- If the Storage Resource agent is deployed on an AIX system. Check to see if the system has the minimum level of AIX maintenance level and patches. See http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=/com.ibm.tpc_V42.doc/fqz0_r_sw_requirements_os.html.
- If the Storage Resource agent is running on a platform for which the Storage Resource agent fabric functions are not supported. See http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=/com.ibm.tpc_V42.doc/fqz0_r_sw_requirements_os.html.
- If the Storage Resource agent has been disabled for fabric functions. To enable an agent, complete the following steps:
 1. Open the Tivoli Storage Productivity Center GUI.
 2. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
 3. In the content pane, select an agent. Under the Select Action menu, click **Enable**.

Storage Resource agents running on AIX

Your fabric probe or zone control job can fail on AIX with error message: AGT0430I.

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

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

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

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

Fabric events

Some of the common reasons for not getting notification of external fabric events or not having fabric probes run automatically in response to events are as follows:

- Switches are not configured to send SNMP traps to the server.
- The fabric for which an event was received or fabrics discovered by the Storage Resource agent are not included in any probes defined by the user. An automatic probe of fabrics is done only for fabrics that are included in some probe definition.
- The switch is configured for a trap level that results in the trap being filtered and not sent from the switch to the server. The trap sent from the switch to the server is filtered by the server because of default filtering rules or filtering rules (or both) configured by the user.

To configure tracing to check if the traps are filtered, set the trace level parameter `dsn.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_installation_directory\data\log\traceTPCDeviceServer.log`.

HBAs not discovered by a Storage Resource agent

The Storage Resource agents do not use RNID and SCSI inquiries for remote HBA and device identification. You need a Storage Resource agent on each system for which HBA information is needed. You need a native interface storage system or CIM agent to identify the storage devices.

Check for any messages in the computer probe log for information about what might be the problem. Other things to check are:

- If the supported HBA driver and API are installed. In some cases, the HBA driver includes the API. In other cases, the HBA driver and API need to be installed separately. Contact your HBA vendor for information about the driver and API.
- If the HBA and API are at a level supported by Tivoli Storage Productivity Center. See http://www-01.ibm.com/support/docview.wss?rs=40&uid=swg21386446&loc=en_US&cs=utf-8&lang=en.

- There are some platform restrictions for Storage Resource agent fabric functions. For information about platforms, see http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=/com.ibm.tpc_V42.doc/fqz0_r_sw_requirements_os.html.
- Check to see if the HBA is configured appropriately. For information about how to configure the HBA, contact your HBA vendor.

The HBA to port relationship is a logical concept. Sometimes one physical card might have more than one HBA port and vice versa. Multiple ports in the same HBA share the same node WWN. Check the node and port WWN names in the topology viewer.

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_installation_directory/data/log/`

Fabric alerts

If you are not getting a fabric alert when you expected one, check if the information displayed in Tivoli Storage Productivity Center reflects the change. Check if the appropriate types of alerts have been defined for the fabric.

Check the following logs for error messages:

`TPC_installation_directory/data/log/server_XXXX.log`
`TPC_installation_directory/data/log/TPCD_XXXX.log`

Fabric probe failures

If you are getting a fabric probe failure, check the following items:

- Look at the fabric probe job log and Device server message and trace log files. If the error message is in processing before the agent is called, the issue is with the agent assignment.
- Check for communication errors when the agent is called, any error response from the agent, or no response from the agent.
- Check the agent log files for error messages.
- Check to see if the Data agent or Fabric agent is configured to communicate with the server and that the server is registered with the Agent Manager. To register the server with the Agent Manager, in the navigation tree, click **Administrative Services** → **Configuration** → **Agent Manager Registration**.

CIM indications with Data Center Fabric Manager

Here is a list of some things to check when you are having issues with CIM indications from the Data Center Fabric Manager (DCFM):

CIM indication not received

- Verify that SNMP registration is successful by the DCFM discovery dialog
- Verify if the configured SNMP port is blocked by a firewall

- Verify if the MIB capability for trap forwarding and tracking changes are configured properly

CIM indication is delivered in reverse order

If a node moves from fabric 1 to fabric 2, you might see these issues:

- The DCFM asset collection gets triggered in fabric 2 first.
- The SNMP registration fails for fabric 1 and goes through for fabric 2.

To get the DCFM/IA logs, follow these steps:

1. Open the DCFM Server Management Console.
2. Go to the **Technical Support Information** tab.
3. Click **Browse** and select the log destination. Click **capture**.

CIM agents

This topic provides information for troubleshooting CIM client or CIM agent problems.

XML requests and responses are written to the CIM client trace file that is in the following directory:

```
C:\Program Files\IBM\TPC\device\log\cim-xml-trace.txt
```

The **cimxmlTracing** configuration parameter in table T_RES_CONFIG_DATA must be set to true to collect tracing information.

To set the CIM client tracing parameter, follow these steps:

1. Check to see what the **cimxmlTracing** attribute is set to in the DB2 table T_RES_CONFIG_DATA. Run the following command:

```
tpctool getdscfg -user user_ID -pwd password -url localhost:9550
-property cimxmlTracing
```

This command displays the following information:

```
Property      Context Value
=====
cimxmlTracing CIM      false
```

If the value is false, set the attribute to true:

```
tpctool setdscfg -user user_ID -pwd password -url localhost:9550
-context CIM -property cimxmlTracing true
```

Run the **getdscfg** command again to display the setting of the value. The value is displayed:

```
Property      Context Value
=====
cimxmlTracing CIM      true
```

2. After modifying the attribute, stop and restart the Device server.

Chapter 4. Troubleshooting Tivoli Integrated Portal

Tivoli Storage Productivity Center includes 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 Tivoli Storage Productivity Center single sign-on environment using LTPA tokens. This topic provides information about 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 Tivoli Integrated Portal authentication service to authenticate its users against 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 Tivoli Integrated Portal, follow these steps:

1. Open the Tivoli Integrated Portal console.
2. Click **Troubleshooting** → **Logs and Trace**.
3. In the right pane, select a server (**server1**). Click 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 Tivoli Integrated Portal server.

Turn tracing on for the token service infrastructure

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

To turn on tracing 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 wanted package to debug.
3. Restart the Tivoli Integrated Portal server.

Chapter 5. Troubleshooting DB2 and the database

Tivoli Storage Productivity Center stores information in a DB2 database. This topic provides information for troubleshooting the DB2 database and for daily operations for DB2.

DB2 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 can also 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. You can use the health monitor 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.

For the best DB2 database performance, monitor the table spaces for Tivoli Storage Productivity Center. You can monitor the table spaces through the DB2 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 might still have some space left for other data containers or their file systems.

To get to the DB2 Health Center on Windows, click **Start** → **All Programs** → **IBM DB2** → **Monitoring Tools** → **Health Center**. To get to the DB2 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 are set for the Monitor Heap Utilization:

- Default is Yes
- Evaluate is Yes (indicates that evaluation is turned on)

- 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 are set as follows:

- Default is Yes
- Evaluate is Yes (indicates that evaluation is turned on)
- 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.

Tivoli Storage Productivity Center uses the default configuration of IBM the DB2 health monitor. This health monitor configuration is sufficient in most cases. If you want to get earlier warnings or include more parameters to monitor, you can do so.

For more information about the DB2 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 7.
2. Run **repocopy** to capture the Tivoli Storage Productivity Center database. For more information about **repocopy**, see “Repository Copy tool” on page 5. 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 quotation marks.
 - f. Take the default values for the connection properties.
 - g. Compress 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`

Determining when a DB2 table is accessible

To determine whether the DB2 table is accessible, follow these steps:

1. Ensure that all table space 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 **System** → **Instances** → **DB2** → **Databases**.
2. Select a database and go to **Tables**.
3. Select a table from the right pane. (The right pane contains all the tables under the database that you selected.)
4. Right-click the table that you selected, and click **Sample Contents** from the menu that is displayed.

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 6. Troubleshooting performance and memory problems

This topic provides information about how and when to tune Tivoli Storage Productivity Center. This topic also provides troubleshooting tips and information about common user errors.

This information applies to the Data Manager.

From a server perspective, you do not have to tune 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 only shows up in the selection list if it supports performance data collection. This situation can occur if the CIM Agent is not at the appropriate 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 condition 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 SAN Volume Controller subsystems, this error message can occur during performance data collection. This condition usually indicates a problem with mismatching timestamps for the I/O statistics dump files on the SAN Volume Controller 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 check the agent by clicking: **Administrative Services** → **Agents** → **CIMOM**. Select **Test CIMOM Connection**.

Why does my CIM Agent run slowly?

Try stopping and restarting the Tivoli Storage Productivity Center services. If this action does not help, increase the value of `com.ibm.tpc.perf.ConteConnectTimeOut` in the following file:

```
TPC_installation_directory\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 numerous volumes. This condition causes a delay in collecting the first performance sample. To work around this problem, place the CIM Agent and Tivoli Storage Productivity Center servers on faster systems. 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 the switch is not listed in the storage subsystems list that can be monitored, perform the following actions:

- 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 subprofile.
- 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 action is required for the switch performance monitor even if you are not using out-of-band Fabric agents.

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 status of a performance monitor, expand **Disk Manager > Monitoring > Subsystem Performance Monitors** in the navigation tree, select the performance monitor, and click **View Job History** in the content pane. The Job Management panel is displayed. Use the **Jobs for Selected Schedule** section on this panel to view the status of a performance monitor and access its log files.

You can also run the GUI from a command prompt window:

```
TPC_installation_directory\jre\bin\java.exe -Xmx512m  
-classpath TPC_installation_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_installation_directory\data\log\server_xxx.log  
TPC_installation_directory\device\log\tracePerfMgrxxx.log
```

There are no GUI logs, so you must use the Service tool to collect trace data. For information about the Service tool, see “Service tool: collecting information” on page 7. You can gather trace data for review by running the Service tool:

```
TPC_installation_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 calls show more detailed information. For example, you can 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 is listed in 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. Investigate the cause of this condition before increasing the memory allocation for either the server or the agent.

Note: If you must substantially increase agent memory for one agent or several, then you must increase the memory on the server at **least** 25% greater than that of the largest agent memory setting. This action is required so that the server can 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 causes this condition can be related to the following actions:

- 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 actions listed.

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 memory for the agent.

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.

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. Tivoli Storage Productivity Center tries 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 this value is not set, the default cutoff value for the enumeration algorithm is 20%. If that value is less than 20% for the ports for a CIMOM targeted for performance data collection, the association algorithm is used instead. The association algorithm is used if the **MaxSwitchPortLimit** is not exceeded (see **MaxSwitchPortLimit**). This action 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 this value is not set, the default cutoff value for the association algorithm is 256 ports, so that performance data is not collected for more than 256 ports simultaneously using the association algorithm. Instead, the enumeration algorithm is 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_installation_directory/device/conf/pm.conf`

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

```
#com.ibm.tpc.perf.MinSwitchPortRatio = 0.2  
#com.ibm.tpc.perf.MaxSwitchPortLimit = 256
```

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

Some commonly used parameters

AbbreviatedProbe parameter

When Tivoli Storage Productivity Center probes the system hardware, it attempts to gather the same set of information for all disk devices. This action generates nonfatal errors in the probe log for certain disk devices that do not support some of the data gathering commands for the probe (like getting a list of disk defects). This action can also cause errors to be generated in the logs for the system and trigger alerts on other systems management consoles, which might be undesirable. To alleviate this problem, the **abbreviatedProbe** parameter might be set on the server. It is a global setting and effects 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 can be sent to any disk drives anywhere are:

- INQUIRY
- READ CAPACITY

This means that the following information is unavailable 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 NAS file systems for the device. 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 exported file systems for the device. 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 action is done so that a situation cannot occur where the product double counts information about a file system because it is treating two export jobs to the same file system as separate file systems. If for some reason the discovery criteria cannot be met, then it can be overridden by setting the **saveNonRoot** parameter. This action then discovers 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"
```

The instance name for the server. On installing the Tivoli Storage Productivity Center server for the first time, the instance name is automatically set to the host name for the server. If you move Tivoli Storage Productivity Center to a new host and copy the repository to a new DBMS using a

backup/restore or database memory dump, this parameter must reflect the original instance name (which always is the first host name for the server). Otherwise, you cannot see any of the old data on the new server. After the new server is started, the server detects that its host name has changed and broadcasts this change to the agent systems. The agents then change this entry in their agent configuration files.

maxConnections=500

The maximum connections that the server can handle simultaneously. After this number of connections is reached, the rejected agent waits for four hours and tries to make the connection again. This number can be increased. On Solaris, this action requires 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 Tivoli Storage Productivity Center. Other operating systems might require tuning as well.

Note: GUI connections override this setting.

routerThreads=1

All incoming connections are handled by the routing thread for the server before they are rejected or handed to another service provider (agent, GUI, or scheduler). Connections to this process are not logged or limited by Tivoli Storage Productivity Center. This action results in allowing many more connections to the server than the **maxConnection** parameter would allow in large environments. This action occurs because the routing thread allows all connections to be queued before it hands off the connection 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 three threads helps by not allowing many *unclassified* connections to queue up.

scheduler.config file

maxSubmitThreads=3

Controls the number of threads that the scheduler service provider can create 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 see "wait time exceeded" errors. This action 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 waits 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 does not submit the request to 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 SNMP Enterprise designation as assigned in IANA:

<http://www.iana.org/assignments/enterprise-numbers>

Do not set any of the other parameters unless told to do so by support. Setting some of these parameters might 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 are located in the `TPC_installation_directory/data/config` directory, where `TPC_installation_directory` represents the directory where you installed the product. When you change the configuration file for the server, you must stop and restart the server before those changes take effect.

Chapter 7. Troubleshooting the tape library

This topic provides information about 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 IBM TotalStorage 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 condition assumes that the tape library supports the SMI-S 1.1. profile for tape libraries.

Note:

- The media changers displayed by Tivoli Storage Productivity Center for the TS3500 tape libraries are logical partitions of the given library, and not physical accessors.
- Depending on the size of the tape library and network latency between the SMI-S agent host and the tape library, probes of the tape library might take a long time and might fail because of timeouts.
- The SMI-S agent does not display an exception to the SMI-S client such as Tivoli Storage Productivity Center when communication between the SMI-S agent and the tape library fails. Instead, empty result sets are returned. This limits Tivoli Storage Productivity Center's capabilities in detecting connection failures relating to tape libraries.
- In this release, support for the IBM 3494 Tape Libraries is limited to discovery and in-place starting of the ETL Specialist. This condition 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 timeouts. Instead, spread the libraries across multiple probe jobs with different start times.
- After the tape libraries are registered, and then a change is made to the IBM SMI-S Agent for Tape, a condition can occur where not all of the tape cartridges are returned to the CIM client. To resolve this situation, restart the IBM SMI-S Agent for Tape; see the documentation for the SMI-S Agent for instructions on how to do this.

The same issue might 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 8. Network connectivity

The following information describes network connectivity commands that you can use to identify and resolve network-related problems.

There are some commands that are helpful in identifying network connectivity problems:

- **netstat**
- **nslookup**
- **ping**
- **telnet**
- **tracert**

These commands are described in this topic.

netstat command

The **netstat** command queries TCP/IP about the network status of the local system. This command is useful for debugging purposes.

Syntax

```
►► --netstat --a --e --n --o --p protocol --r --s --interval seconds ◀◀
```

Parameters

- a Displays all connections and listening ports.
- e Displays Ethernet statistics. This parameter can be combined with the **-s** parameter.
- n Displays addresses and port numbers in numeric form.
- o Displays the owning process ID associated with each connection.
- p *protocol*
Shows connections for the specified protocol. You can specify one of these protocols:
 - IP (only if specified with the **-s** parameter)
 - IPv6 (only if specified with the **-s** parameter)
 - ICMP (only if specified with the **-s** parameter)
 - ICMPv6 (only if specified with the **-s** parameter)
 - TCP
 - TCPv6
 - UDP
 - 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 might be used to specify a subset of the default.

-interval*seconds*

Displays the selected statistics at intervals, pausing the specified number of seconds between each display.

Press Ctrl+C to stop displaying statistics. If this parameter is omitted, the **netstat** command displays the current configuration information once.

Description

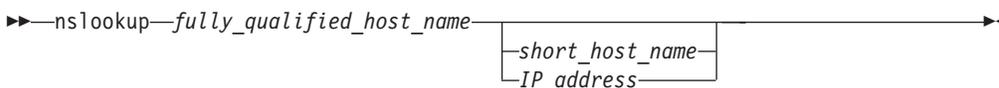
This command displays the following information:

- Active TCP connections on the local system
- State of all TCP/IP servers on the local system and the sockets used by them
- Devices and links used by TCP/IP
- IP routing tables (gateway tables) in use by the local system

nslookup command

The **nslookup** command locates information about network nodes, examines the contents of a name server database, and establishes the accessibility of name servers.

Syntax



Parameter

fully_qualified_host_name , *short_host_name*, *IP_address*

Specify a fully qualified host name, short name, or IP address.

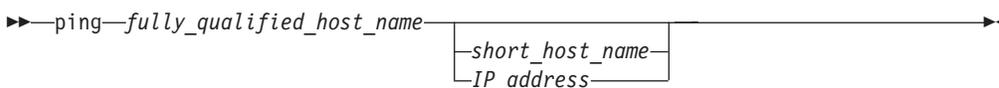
Description

The **nslookup** command is a DNS utility for querying 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.

ping command

The **ping** command tests, manages, and measures network performance. It can be used to isolate network failures.

Syntax



Parameters

fully_qualified_host_name, *short_host_name*, *IP_address*

Specifies the target host.

Description

This 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 data, including round-trip time.

Note: Most UNIX systems provide a similar utility called **traceroute**.

telnet command

The **telnet** command contacts a specific port.

Syntax

```
►► telnet [host_name | ip_address] [port]
```

Parameters

host_name | *ip_address*
Specifies the target host.

port
Specifies the port.

Description

Use the **telnet** command to contact a specific port to perform its scheduling activities and to establish communication with the server and agent. If **telnet** is started without parameters, it enters command mode as indicated by its prompt (telnet>). In this mode, **telnet** accepts and runs the commands. If **telnet** is started with parameters, it performs an open command with those parameters.

In there is an error, or if the connection is closed by the remote host, **telnet** returns a value of 1. Otherwise, **telnet** returns a value of zero (0).

Error messages

The following error messages might be displayed by the **telnet** command:

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 IP address. Contact the system administrator to check whether there is an entry for the remote host in the hosts database.

?Invalid command

A command that is not valid was typed in the **telnet** command mode.

tracert command (Windows only)

The **tracert** command displays the network route to a specific host and identifies faulty gateways.

Syntax

```
►► --tracert --d --h maximum_hops --j host_list --w timeout --target_name ►►
```

Parameters

-d Does 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.

Description

Ensure that you are connected to a network. Set the prompt to: C:\windows\>

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 provides about each location to which the packet is sent before it is received.

Appendix A. Accessibility features for Tivoli Storage Productivity Center

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

The following list includes the major accessibility features in IBM Tivoli Storage Productivity Center:

- Tivoli Storage Productivity Center functions are available using the keyboard for navigation instead of the mouse. You can use keys or key combinations to perform operations that can also be done using a mouse. However, you must use the mouse to navigate the Topology Viewer and report graphs. Standard operating system keystrokes are used for standard operating system operations.
- You can use screen readers to read the user interface.
- The user interface communicates all information independently of color.
- The *IBM Tivoli Storage Productivity Center Information Center*, and its related publications are accessibility-enabled and include the following accessibility features:
 - The information center is provided in XHTML 1.0 format, which is viewable in most Web browsers. XHTML allows you to view documentation according to the display preferences set in your browser. With XHTML, you can use screen readers and other assistive technologies.
 - All documentation for Tivoli Storage Productivity Center is available in Adobe Portable Document Format (PDF) using the Adobe Acrobat Reader. You can access the PDFs from the Printable PDFs topic in the information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp>.
 - All images are provided with alternative text, so that users with vision impairments can understand the contents of the images.

Keyboard navigation

This product uses standard Microsoft Windows navigation keys.

Displaying the interface

Click **Preferences** → **Look and Feel** to select how to display the Tivoli Storage Productivity Center user interface. To do so, complete the following steps:

1. Start the Tivoli Storage Productivity Center user interface.
2. Click **Preferences** → **Look and Feel** to change the visual appearance of the user interface to best suit your visual needs:
 - Windows Classic
 - Windows
 - CDE/Motif
 - Metal

Limitations for sight-impaired users on IBM Tivoli Storage Productivity Center for Replication

There are some limitations in the product graphical user interface (GUI) that require special attention before you start using Tivoli Storage Productivity Center for Replication with a screen reader:

Dojo components are not read by all screen readers.

The JAWS screen reader does not read some Dojo components on Internet Explorer 7. Use the command-line interface instead of the GUI with JAWS on Internet Explorer 7.

IBM and accessibility

See the IBM Human Ability and Accessibility Center Web site for more information about the commitment that IBM has to accessibility.

Appendix B. Accessibility features for IBM Tivoli Storage Productivity Center for Replication

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

IBM strives to provide products with usable access for everyone, regardless of age or ability.

These are the major accessibility features in IBM Tivoli Storage Productivity Center for Replication:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

See the IBM Human Ability and Accessibility Center website at www.ibm.com/able for more information about the commitment that IBM has to accessibility.

Accessibility and keyboard shortcuts in the information center

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. Using the major accessibility features in this product, users can perform these tasks:

- Use assistive technologies, such as screen-reader software and digital speech synthesizer, to hear what is displayed on the screen. Consult the product documentation of the assistive technology for details on using those technologies with this product.
- Operate specific or equivalent features using only the keyboard.
- Magnify what is displayed on the screen.

In addition, the documentation was modified to include the following features to aid accessibility:

- All documentation is available in HTML formats to give the maximum opportunity for users to apply screen-reader software technology.
- All images in the documentation are provided with alternative text so that users with vision impairments can understand the contents of the images.

Use the following key combinations to navigate the interface by keyboard:

- To go directly to the Topic pane (the right-hand side), press Alt+K, and then press Tab.
- In the Topic pane, to go to the next link, press Tab.
- To go directly to the Search Results view in the left hand side, press Alt+R, and then press Enter or Up Arrow to enter the view.
- To go directly to the Navigation (Table of Contents) view in the left-hand side, press Alt+C, and then press Enter or Up Arrow to enter the view.
- To expand and collapse a node in the navigation tree, press the Right and Left Arrow.

- To move to the next topic node, press the Down Arrow or Tab.
- To move to the previous topic node, press the Up Arrow or Shift+Tab.
- To go to the next link, button, or topic node from inside on of the views, press Tab.
- To scroll all the way up or down in a pane, press Home or End.
- To go back, press Alt+Left Arrow; to go forward, press Alt+Right Arrow.
- To go to the next pane, press F6.
- To move to the previous pane, press Shift+F6.
- To print the active pane, press Ctrl+P.

Limitations for sight-impaired users

There are some limitations in the product graphical user interface (GUI) that require special attention before you start using the product with a screen reader:

Pull-down menus are positioned directly above or before the radio button that activates it.

Currently, there are several GUI panels that have a pull-down menu positioned directly above or before the radio button that activates it. Using a screen reader that processes the fields and controls of a panel sequentially, you might select the radio button, and then not know that the associated pull-down menu has been activated. The screen reader would process the inactive pull-down menu first, and then move on to the radio button next, but not advise you that the pull-down menu has been activated if you were to select the radio button.

As you encounter pull-down menus in the following GUI panels, keep in mind that there are radio buttons associated with them that activate the pull-down menu:

- Administration
- ESS/DS Paths
- Sessions
- Session Details
- Storage Systems

Tables do not have summary attributes and do not use the header or row ID attributes when reading a single cell.

Another limitation that a user requiring assistive technology might encounter concerns the table objects. These tables do not have summary attributes, nor do these tables use the header or row ID attributes when reading a single cell within. The table summary is the description adopted from the Main content section of the panel. The table row and column number is still read along with cell data and therefore you can infer the column header and row ID.

Some screen readers do not correctly read the abbreviation for Enterprise Storage Server.

Some screen readers do not correctly read the abbreviation for Enterprise Storage Server, which consists of the first letter of each word pronounced individually. Instead, the screen readers recognize it as a word. If you can add dictionary words with your screen reader, replace the capitalized character sequence with the sequence E space S space S.

Typically, this abbreviation is used in the combination form of ESS/DS. It refers to the Enterprise Storage Server 800, the DS6000, or the DS8000. Interpret any references to ESS/DS as belonging to those products.

Some visual artifacts might persist in the web browser after disabling the cascading style sheet (CSS).

If you disable the CSS, some visual artifacts, such as decorative elements, might persist in the Firefox and Internet Explorer GUIs. These artifacts do not affect performance. If they become too distracting, consider using the command-line interface instead.

The confirmation dialog box appears and immediate focus is given to Yes button.

When the confirmation dialog box appears, focus is given to the **Yes** button; therefore, the screen reader reads “Yes” but does not read the confirmation text. This might occur in the following places: performing an action on a session, removing connections to storage systems, clicking the About link, and creating a high-availability connection. To read the confirmation text before selecting the **Yes**, **No**, or **OK** button, use your screen reader's ability to navigate headings to view the previous heading before the button.

Dojo components are not read by all screen readers.

The JAWS screen reader does not read some Dojo components on Internet Explorer 7. Use the command-line interface instead of the GUI with JAWS on Internet Explorer 7.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
U.S.A*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE: This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

DIA Data Integrity Assurance, Storewiz, Storwize, and the Storwize logo are trademarks or registered trademarks of Storwize Inc., an IBM Company.

Adobe is a registered trademark or trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Intel Xeon, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Glossary

This glossary includes terms and definitions for IBM Tivoli Storage Productivity Center.

The following cross-references are used in this glossary:

- *See* refers the reader from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- *See also* refers the reader to a related or contrasting term.

To view glossaries for other IBM products, go to: <http://www.ibm.com/software/globalization/terminology/>.

active management server

A management server from which the storage environment can be monitored and managed. The active server replicates its database to the standby server.

administrative role

A job function that identifies the tasks that a user can perform and the resources to which a user has access. A user can be assigned one or more roles.

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.
- Permitting only CKD volumes to be added to the HyperSwap session.

business continuity

The capability of a business to withstand outages and to operate mission-critical services normally and without interruption in accordance with predefined service-level agreements.

business continuity solution

A combination of hardware, software, and services that ensures business continuity. The business continuity solution must address data, operational environment, applications, the application hosting environment, and the user interface.

common agent package (CAP)

An XML-based policy file that defines how Agent Management Services monitors a Storage Resource agent. Tivoli Storage Productivity Center creates a CAP file with each instance of a Storage Resource agent.

CIM See *Common Information Model*.

CIM agent

The code that consists of common building blocks that can be used instead of proprietary software or device-specific programming interfaces to manage devices that are compliant with the Common Information Model (CIM).

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 See *count key data*.

cluster

1. In SAN Volume Controller or Storwize V7000, a pair of nodes that provides a single configuration and service interface.
2. In IBM System Storage[™] DS8000[®], a partition capable of performing all DS8000 series functions. With two clusters

in the DS8000 storage unit, any operational cluster can take over the processing of a failing cluster.

Common Information Model (CIM)

An implementation-neutral, object-oriented schema for describing network management information. The Distributed Management Task Force (DMTF) develops and maintains CIM specifications.

comma-separated value (CSV)

A common text file format where a comma is used to separate values. See also *CSV file*.

community name

The part of a Simple Network Management Protocol (SNMP) message that represents a password-like name and that is used to authenticate the SNMP message.

consistency group

A group of copy relationships between virtual volumes or data sets that are maintained with the same time reference so that all copies are consistent in time.

copy set

The set of source volumes or target volumes involved in a FlashCopy operation.

count key data (CKD)

In mainframe computing, a data-record format employing self-defining record formats in which each record is represented by up to three fields: a count field identifying the record and specifying its format, an optional key field that can be used to identify the data area contents, and an optional data field that typically contains the user data.

CSV See *comma-separated value*.

CSV file

A comma-separated value text file, commonly used to exchange files between database systems that use different formats.

data collection

The process of obtaining performance and availability monitoring data and providing that data to a metric evaluator. Examples of data collectors include DNS

probes, web page analyzers, or database analyzers. See also *discovery*.

data exposure

The time between the point at which the data is written to primary storage and its replication in secondary storage.

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 *schedule*.

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 storage area network (SAN) switch.

Domain Name System (DNS)

The distributed database system used to map domain names to IP addresses.

DNS See *Domain Name System*.

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 An occurrence of significance to a task or system. Events can include completion or failure of an operation, a user action, or the change in state of a process.

extended count key data (ECKD™)

An extension of the count key data (CKD) architecture.

fabric A complex network of hubs, switches, adapter endpoints, and connecting cables that support a communication protocol between devices. For example, fibre channel uses a fabric to connect devices.

failback

In high availability disaster recovery, the process of restarting the original primary system and returning it to its status of

primary system after a failover has occurred. See also *failover*.

failover

An automatic operation that switches to a redundant or standby system in the event of a software, hardware, or network interruption. See also *fallback*.

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.

fixed-block (FB) volumes

An architecture for logical devices that specifies the format of the logical data units on the device. The logical data unit is a block. All blocks on the device are the same size (fixed size); the subsystem can access them independently. This format is the required format of the logical data units for host systems that attach with a Small Computer Systems Interface (SCSI) or fibre-channel interface.

FlashCopy

An optional feature of the Storage System DS family that can make an instant copy of data, that is, a point-in-time copy of a volume.

Global Copy

An optional capability of the IBM System Storage 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 two-site extended-distance remote copy. Data that is written by the host to the storage unit at the local site is automatically maintained at the remote site. See also *Metro Mirror* and *remote mirror and copy*.

globally unique identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system.

Hardware Management Console (HMC)

A system that provides server management and virtualization (partition) management for storage systems.

heat map

A color-coded data chart where colors are used to differentiate values in a data set.

high availability

Pertaining to a clustered system that is reconfigured when node or daemon failures occur, so that workloads can be redistributed to the remaining nodes in the cluster.

high availability solution

A combination of hardware, software, and services that fully automates the recovery process and does not disrupt user activity. High availability solutions must provide an immediate recovery point with a fast recovery time.

host

A computer that is connected to a network and that provides an access point to that network. The host can be a client, a server, or both a client and server simultaneously.

host volume

A volume that represents the volume functional role from an application point of view. The host volume can be connected to a host or server. It receives read, write, and update application I/O, depending on the site to which the application is writing. Host volumes are abbreviated as Hx.

HyperSwap

A z/OS[®] replication method that manages CKD volumes in Peer-to-Peer Remote Copy (PPRC) relationships and ensures that data consistency is not violated.

in-band discovery

The process of discovering information about the storage area network (SAN), including topology and attribute data, through the fibre-channel data paths. Contrast with *out-of-band discovery*.

intermediate volume

A volume that receives data from the

primary host volume during a replication with practice session. Depending on the replication method being used, data on this volume might not be consistent. Intermediate volumes are abbreviated as Ix.

Java Database Connectivity (JDBC)

An industry standard for database-independent connectivity between the Java platform and a wide range of databases. The JDBC interface provides a call level interface for SQL-based and XQuery-based database access.

JDBC See *Java Database Connectivity*.

job A unit of work within a run for a schedule. One run can have multiple jobs. For example, a system probe schedule has one job created per system in its definition.

job scheduler

A component of the Data server that deploys all monitoring activities. The job scheduler controls when monitoring jobs are run by agents.

journal volume

A volume that holds a consistent copy of data until a new consistent copy is formed. The journal volume restores the last consistent point during a recovery. Journal volumes are abbreviated as Jx

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier used to differentiate devices, each of which is a logical unit (LU).

LUN See *logical unit number*.

managed disk (MDisk)

A Small Computer System Interface (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.

Management Information Base (MIB)

In the Simple Network Management

Protocol (SNMP), a database of objects that can be queried or set by a network management system.

management server

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

management server relationship

A connection between two Tivoli Storage Productivity Center for Replication servers. One server acts as the active server and replicates data necessary for the standby server to take over control of the Tivoli Storage Productivity Center for Replication environment if necessary.

MDisk

See *managed disk*.

Metro Mirror

A method of synchronous remote replication that operates over distances up to 300 kilometers. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Changes made to the source data are propagated to the target before the write finishes posting. See also *remote mirror and copy* and *Global Copy*.

Metro Global Mirror

A three-site, high availability, disaster recovery solution. Metro Global Mirror uses synchronous replication to mirror data between a local site and an intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site.

MIB See *Management Information Base*.

native interface

An interface that is specific to a system or subsystem.

out-of-band discovery

The process of discovering storage area network (SAN) information, including topology and device data, without using the fibre-channel data paths. A common mechanism for out-of-band discovery is the use of Simple Network Management Protocol (SNMP) Management Information Base (MIB) queries, which are

- invoked over an IP network. Contrast with *in-band discovery*.
- ping** A job that tracks the availability of assets and that is performed by an agent. Several ping jobs can be used to monitor the availability of any computer or subset of computers in the network. See also *schedule*.
- pool** A division of main or auxiliary storage.
- practice volume**
Volumes that can be used to test disaster-recovery actions while maintaining disaster recovery capability. Practice volumes are available in Metro Mirror Failover and Failback sessions, Global Mirror (two directions) Failover and Failback sessions, and Metro Global Mirror with Practice sessions.
- primordial pool**
Storage capacity that is deallocated on a storage device. Storage pools are created by allocating storage capacity from primordial pools.
- probe** A job that itemizes and creates an inventory of assets, such as computers, controllers, disk drives, file systems, and logical units, and that is performed by an agent. Several probe jobs can be used on any computer or subset of computers. See also *schedule*.
- RAID** See *Redundant Array of Independent Disks*.
- Redundant Array of Independent Disks (RAID)**
A collection of two or more physical disk drives that present to the host an image of one or more logical disk drives. In the event of a physical device failure, the data can be read or regenerated from the other disk drives in the array due to data redundancy.
- recovery point objective**
The maximum amount of data loss that can be tolerated during a service interruption.
- remote console**
A computer with access to the server component of Tivoli Storage Productivity Center that is installed on another computer.
- remote mirror and copy**
A feature of a storage server that constantly updates a secondary copy of a logical volume to match changes made to a primary logical volume. The primary and secondary volumes can be on the same storage server or on separate storage servers. See also *Metro Mirror* and *Global Mirror*.
- role** See *administrative role* or *volume role*.
- role pair**
The association of two roles in a replication session that take part in a copy relationship. For example, in a Metro Mirror session, the role pair can be the association between the volume roles of Host1 and Host2.
- run** A single invocation of a schedule. A schedule that collects data multiple times has multiple runs. For example, a scan schedule that collects data daily has seven runs per week.
- SAN** See *storage area network*.
- scan** A job that monitors the usage and consumption of storage and the constraints. See also *schedule*.
- schedule**
A definition of a monitoring task that collects data about storage resources. Monitoring tasks include scans, probes, pings, discovery, and performance monitors. For example, define a probe schedule to collect data about storage systems; define a scan schedule to collect file-level information about hosts.
- SCSI** See *Small Computer Systems Interface*.
- session**
A collection of volumes within a logical subsystem that are managed together during the creation of consistent copies of data. All volumes in a session must transfer their data successfully to the remote site before the increment can be called complete.
- Simple Network Management Protocol (SNMP)**
A set of protocols for monitoring systems and devices in complex networks. Information about managed devices is defined and stored in a Management Information Base (MIB).
- site awareness**
Association of a location (site) with each storage system and each session's volume role. When a copy set is added, only the

storage systems with locations that match those locations of the volumes' roles can be selected. This constraint prevents the reversing of a hardware relationship or the selection of a volume at the wrong location.

- site 1** The location of the primary storage that contains the source data.
- site 2** The location of the secondary storage that receives the data updates from the primary storage.
- site 3** In Metro Global Mirror, the location of the tertiary storage that receives data updates from the secondary storage.

site switching
See *failover*.

Small Computer System Interface (SCSI)
An ANSI-standard electronic interface that allows personal computers to communicate with peripheral hardware, such as disk drives, tape drives, CD-ROM drives, printers, and scanners faster and more flexibly than previous interfaces.

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*.

SNMP
See *Simple Network Management Protocol*.

standby management server
A management server that is a backup for the active server. The replication environment cannot be monitored or managed from the standby server.

storage area network (SAN)
A dedicated storage network tailored to a specific environment, combining servers, storage products, networking products, software, and services.

storage connection
The connection between a storage system and the management server, either directly or through a management console or z/OS connection.

Storage Management Initiative Specification (SMI-S)

A design specification developed by the Storage Networking Industry Association (SNIA) that specifies a secure and reliable interface with which storage management systems (SMSs) can identify, classify, monitor, and control physical and logical resources in a storage area network (SAN). The interface integrates the various devices to be managed in a SAN and the tools used to manage them.

Storage Networking Industry Association (SNIA)

An alliance of computer vendors and universities that focus on developing and promoting industry standards for storage networks.

storage pool
A grouping of storage space consisting of volumes, logical unit numbers (LUNs), or addresses that share a common set of administrative characteristics

storage resource group
A named collection of logically-related storage entities. Storage entities include fabrics, switches, computers, storage systems, tape libraries, and other storage resource groups. For example, a storage resource group that represents a business critical email application might include all the hosts that participate in the email application cluster, the storage systems that provide storage to the application, and the switches and fabrics through which data of the application is configured to travel.

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 or logical mapping of the location of networking components or nodes within a network. Use the topology viewer to display a topology graphically, showing storage devices and their interconnections.

track space-efficient (TSE) volumes
A volume in which storage space is allocated on an as-needed basis by using space on the target volume only when tracks are copied from the source volume

to the target volume. FlashCopy sessions permit TSE volumes to be targets of the flash.

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. Subsystems from various vendors can be used. The pool can be split into virtual disks (VDisks) that are visible to the host systems that use them.

virtual storage area network (VSAN)

A fabric within the storage area network (SAN).

volume

A discrete unit of storage on disk, tape, or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control.

volume role

The function that the volume plays in the copy set (such as host volume, journal volume, or intermediate volume) and the location of the volume. The role is composed of the intended use and, for Global Mirror and Metro Mirror, the site location of the volume. For example, the H1 role is made up of host-attached volumes that are located at the primary site.

volume role pair

The association of two volume roles in a replication session. For example, in a Metro Mirror session, the volume role pair can be the association between host volumes at the primary site and host volumes at the secondary site (H1->H2).

VSAN See *virtual storage area network*.

zoning

In fibre-channel environments, the grouping of multiple ports to form a virtual, private, storage network. Ports that are members of a zone can communicate with each other, but are isolated from ports in other zones.

zone alias

A collection of one or more zone members. A zone alias can be added to one or more zone members.

Note: Not all devices support zone aliases. Check with the device manufacturer to determine if zone aliases are supported for that particular device.

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

A

- accessibility features 89
- accessibility features for users with disabilities 91
- agent.config 22
- agents
 - tracing 10

C

- CIM agents
 - for fabric 48
- commands
 - netstat 85
 - nslookup 86
 - ping 86
 - repocopy 5
 - telnet 87
 - tracert 88
- configuration files
 - agent.config 22
 - default locations 19
 - encryption 19
 - for IBM Tivoli Storage Productivity Center 19
 - for tracing 15
 - scheduler.config 21
 - server.config 19
 - TPCD.config 21

D

- Data agent
 - tracing 11
- Data server
 - tracing 10
- Device server
 - tracing 10, 12
- disabilities, accessibility features for users with 91

E

- encryption 19

F

- fabric
 - troubleshooting probes 48
- Fabric agent
 - tracing 13

G

- glossary 99
- GUI
 - tracing 10

I

- IBM Tivoli Storage Productivity Center
 - log files 22

L

- log files
 - audit logs 24
 - for IBM Tivoli Storage Productivity Center 22
 - for Tivoli Integrated Portal 22
 - for Tivoli Storage Productivity Center 23

N

- native.log.config file 16
- netstat command 85
- nslookup command 86

P

- performance monitor
 - changing configuration of switch 78
- ping command 86
- pull-down menus, limitations of to the sight-impaired 91

R

- repocopy command 5
 - exporting data 5
 - importing data 6
- repocopy tool
 - exporting data 5
- Repository Copy tool 5
 - importing data 6

S

- scheduler
 - tracing 10
- scheduler.config 21
- server.config 19
- Service tool
 - for agents 9
 - for servers 8
- sight-impaired 91
- SNMP traps
 - debugging 13
- Storage Resource agent
 - tracing 11
- switch
 - changing configuration for performance monitor 78

T

- telnet command 87
- Tivoli Integrated Portal
 - installation log files 22
- Tivoli Storage Productivity Center
 - configuration files 19
 - installation log files 23
- tools
 - repository copy 5
 - tracing 10
- TPCD.config 21
- tracert command 88
- tracing
 - agents 10
 - configuration files 15
 - Data agent 11
 - Data server 10
 - Device server 10, 12
 - Fabric agent 13
 - GUI 10
 - scheduler 10
 - Storage Resource agent 11
- tracing log files
 - Java files 14
 - native files 14
- trademarks 97
- troubleshooting
 - CIM agents 67
 - configuration history and analysis 45
 - Configure Devices wizard 28
 - Data Path Explorer 47
 - database monitoring 31
 - DB2 71
 - discovery 26
 - Element Management 56
 - fabric probes 48
 - fabric-specific problems 61
 - file system scans 32
 - FlashCopy 55
 - McData Intrepid 57
 - Monitored Computer Storage Space reports 53
 - monitoring service 29
 - NAS and NetApp 50
 - performance and memory problems 75
 - performance monitoring 31
 - reporting groups 53
 - rollup reports 38
 - SAN Planner 39
 - single sign-on 35
 - Storage Optimizer 36
 - tape library 83
 - Tivoli Integrated Portal 69
 - topology viewer 34
 - tpctool 59
 - VMware ESX 51



Program Number: 5608-E14, 5608-WB1, 5608-WB2, 5608-WB3, 5608-WC0,
5608-WC3, 5608-WC4

Printed in USA

GC27-2342-02

